

Coevolutionary Opinion Formation Games

Kshipra Bhawalkar^{*}
Dept. of Computer Science
Stanford University
kshipra@cs.stanford.edu

Sreenivas Gollapudi
Microsoft Research
Silicon Valley
sreenig@microsoft.com

Kamesh Munagala[†]
Dept. of Computer Science
Duke University
kamesh@cs.duke.edu

ABSTRACT

We present game-theoretic models of opinion formation in social networks where opinions themselves *co-evolve* with friendships. In these models, nodes form their opinions by maximizing agreements with friends weighted by the strength of the relationships, which in turn depend on difference in opinion with the respective friends. We define a social cost of this process by generalizing recent work of Bindel *et al.*, FOCS 2011. We tightly bound the price of anarchy of the resulting dynamics via local smoothness arguments, and characterize it as a function of how much nodes value their own (intrinsic) opinion, as well as how strongly they weigh links to friends with whom they agree more.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Nonnumerical Algorithms and Problems; J.4 [Computer Applications]: [Social and Behavioral Sciences]

General Terms

Algorithms, Economics, Theory

Keywords

Price of Anarchy, Games, Opinions

1. INTRODUCTION

The exponential growth in the popularity of the online social networks such as Facebook, Twitter has led to a lot of renewed research in understanding basic sociological phenomena such as opinion and consensus formation. The earli-

^{*}Supported by a Stanford Graduate Fellowship and NSF grant CCF-1016885

[†]Supported by an Alfred P. Sloan Research Fellowship, a gift from Cisco, and by NSF via grants CCF-0745761, CCF-1008065, and IIS-0964560. Part of this work was done while the author was visiting Twitter, Inc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'13, June 1–4, 2013, Palo Alto, California, USA.

Copyright 2013 ACM 978-1-4503-2029-0/13/06 ...\$15.00.

est work in this domain comes from the sociology literature. One notable example in this setting is the *DeGroot model* [6] which studies, in a fixed network, how consensus is formed when individual opinions are updated using the average of the neighborhood. These models were enriched in subsequent work to capture bias agents have for similar opinions. In the popular *bounded confidence* model of Hegselmann and Krause [11] (HK model), opinions lie on a real line. Suppose agent i has opinion z_i . For fixed ϵ , the confidence region of an agent is captured by the set

$$S_i(\vec{z}) = \{j \mid |z_j - z_i| \leq \epsilon\}$$

Each agent i iteratively updates its opinion by averaging the opinions of agents in $S_i(\vec{z})$. In the DW model of Weisbuch *et al.* [22], two random nodes i and j meet and update their opinions z_i and z_j to the average of the two opinions if these originally satisfied $|z_i - z_j| \leq \epsilon$. (See also Dandekar *et al.* [5] for a related model that they term “biased assimilation”.) The key difference in these two models is that the updates in the DW model are more symmetric compared to the HK model in terms of how pairs of agents impact each other’s opinions.

Recent work [13, 12, 7] extends these models to define stochastic processes over social networks, where opinions and friendships evolve dynamically. Holme and Newman [13] term these processes *coevolution* of opinions and friendships. In their model, opinions are discrete (usually binary). At each step, with probability α , a node breaks an existing link chosen at random and forms a link to a node chosen uniformly at random from the nodes of same opinion; with probability $1 - \alpha$, the node changes the opinion of the other end-point of a randomly chosen edge to that of its own. This entire line of work characterizes the conditions under which the network either reaches consensus or partitions into components with disparate opinions.

In a departure from this entire line of work, other studies focused on the importance of going beyond understanding when absolute consensus happens, and studying the *social cost* of the outcomes that emerge [15, 10, 3]. The social cost of an outcome is the sum of the cost of disagreement for all participating agents. Generally, these studies note that consensus rarely emerges in reality, and even if it does, different consensus outcomes have different levels of “desirability” given that the cost of disagreement for different agents can vary. (This point is also implicit in the recent work of Golub and Jackson [9] and Acemoglu *et al.* [1].) For instance, Bindel, Kleinberg, and Oren [3] adopt a model originally proposed by Friedkin and Johnsen [8], who extended

the DeGroot model to include both disagreement and consensus by associating with each node i an innate opinion s_i in addition to its expressed opinion. In this model, given a fixed social network, players choose an expressed opinion to minimize the cost of disagreement with their intrinsic opinion and expressed opinions of their neighbors. The cost of disagreement is modeled as a quadratic function. In contrast with bounded confidence models, the space of equilibrium outcomes in the Bindel *et al.* model is richer due to the diversity of intrinsic opinions. Bindel *et al.* analyze the quality of equilibria by considering the *Price of Anarchy* (PoA) which compares the social cost of the equilibrium with the optimal social cost.

1.1 Our Contributions

The question we ask is the following: *If agents with fixed intrinsic opinions prefer other agents based on their opinions and confidence therein (so that the social network is not fixed, but is itself weighted by opinions), how does the social cost at equilibrium behave?* In effect, we study bounded confidence type models in the framework introduced by Bindel *et al.* and characterize the price of anarchy of the resulting equilibria. This helps us understand how the social outcomes that emerge are influenced by not only intrinsic opinions and the social graph but also by the confidence agents place in close by opinions in weighting friendships. We study the resulting *coevolutionary games* on opinions and friendships in the presence of both symmetric and asymmetric influence models and provide price of anarchy bounds for both models. A well-known example of a symmetric network is Facebook wherein actions (e.g., friendship request, liking a user action, etc.) are observed by all participating agents whereas activities in Twitter (e.g., following a user, re-tweeting a user’s tweet, etc.) illustrate the asymmetric nature of the underlying network.

Symmetric Coevolutionary Games.

In Section 3, we first consider the case where the network is fixed and undirected and the cost of any agent has two contributing factors, *viz.*, the difference in the expressed opinions along all incidence edges, i.e., the agents neighborhood, and the difference between the expressed and the intrinsic opinion of the agent. In other words, if (s_i, z_i) denote the intrinsic and expressed opinions of agent i and $N(i)$ denotes the neighbors of i in the network, the cost to agent i is $\sum_{j \in N(i)} f_{ij}(z_i - z_j) + w_i g_i(z_i - s_i)$. The *weighting functions* f and g capture different costs of disagreement. Moreover, they can also capture the setting where the network is formed probabilistically and the link (i, j) is formed based on the difference in the expressed opinions of the agents. This is similar in spirit to the DW model [22], where nodes average close by opinions.

We present a *tight* PoA bound for *all* strictly convex weighting functions (Section 3.1). As a corollary, this yields a PoA of at most 2 for all convex functions. We also provide a general lower bound construction that has PoA matching the upper bound (see Section 3.2). The technically interesting aspect is that for most weighting functions, the upper bound we obtain involves a set of inequalities capturing the worst case behavior of the *derivatives* of the weighting functions, and this, though simple to state, often does not have a closed form. Despite this difficulty, we show that the functional form of the inequalities themselves can be used to carefully

construct a tight example. For the special class of weighting functions $f(x) = g(x) = |x|^\alpha$ with $\alpha > 1$, we obtain a closed form bound.

Since the bounds we obtain are tight, they allow us to explore how inefficiency of equilibria varies as the weighting functions are varied. We show that as nodes give more weight to having neighbors with close by expressed opinions, *i.e.*, the weighting functions become less convex, the PoA increases and equilibria with worse social cost are possible.

We finally show in Section 3.4 that in the limiting case when the weighting functions are not strictly convex, the PoA can be unbounded. This happens even for linear functions which correspond to node i expressing the *median* of $\{z_j | j \in N(i)\}$, and its own intrinsic opinion s_i (weighted by w_i). Our observations are reminiscent of the classic *Schelling segregation* [21] which shows that even when the social graph and intrinsic opinions start with only a slight bias, the amplification of this bias results in an unbounded polarization at equilibrium.

Asymmetric K-NN Games.

In a natural extension, we study the case where the network is not fixed (see Section 4). Each agent i gets connected to set $N(i)$ of K other agents whose expressed opinions z_j are closest to its intrinsic opinion s_i . The node then expresses opinion z_i to minimize its cost $\sum_{j \in N(i)} (z_i - z_j)^2 + \rho K (z_i - s_i)^2$. Clearly, this results in an asymmetric game whose underlying model is similar to the bounded confidence HK model [11] or the Holme-Newman model [13]. We term this game as the K -NN game; it is a special case of a more general class of games where the weight that agent i puts on agent j ’s opinion in its averaging process is $f_{ij}(s_i, z_{-i})$, where f_{ij} is non-increasing in $|z_j - s_i|$, and non-decreasing in $|z_k - s_i|$ for $k \neq j$ (see Section 4.3).

We show that for $\rho > 1$, the PoA of the K -NN game is at most a constant, where the constant improves as ρ increases. In other words, the social outcomes become better when nodes are “narrow minded” and give larger weight to their own opinion. Though this is a trivial statement in the limit as $\rho \rightarrow \infty$, the interesting point we show is that the PoA is in fact bounded even when agents place roughly equal weight on their own opinion and the combined opinion of their neighbors.

The constant PoA bound is in contrast to the unbounded PoA results in Bindel *et al.* for the directed graph case. In fact, we show that if nodes can choose their neighbors based on closeness of opinions, the PoA can indeed be bounded. We complete the picture by showing that for small ρ , the PoA is at least $1/\rho^2$, so the PoA does indeed deteriorate as nodes become more “broad minded”.

Local Smoothness.

All of our bounds for both the symmetric and the asymmetric models are established using the local smoothness framework of Roughgarden and Schoppman [20]. The framework implies that our bounds extend to the PoA of mixed Nash and correlated equilibria [2], which are always guaranteed to exist. This is useful since we show for the asymmetric K -NN game, a pure Nash equilibrium is not guaranteed to exist. In our context, correlated equilibria can also capture the influence of common signals such as news sources that correlate the expression strategies of the agents.

Finally, we note the local smoothness framework was introduced to prove exact price of anarchy bounds for atomic splittable congestion games. Our work describes another model where the framework can be used to obtain tight bounds.

1.2 Related Work

The closest related work to our paper is that of Bindel et al. [3]. They study the symmetric coevolutionary model when functions f, g are quadratic. Their PoA bound of $9/8$ can be obtained as a corollary of our results. We note that their analysis was specific for quadratic functions, and did not extend in any natural way to our setting. In contrast, our local smoothness approach is simpler and easy to generalize. It also implies the same PoA bound for correlated equilibria as well.

We refer the reader to the excellent text of Jackson [14] for a treatment of various opinion formation models, in particular variants of the bounded confidence models. There is also a large body of work (see [4, 16] and references therein) that has focused on characterizing the convergence of bounded confidence dynamics to either *absolute consensus* or some clustering (polarization). These works show in various settings a sharp transition between absolute consensus and clustering as the value of the confidence bound is varied. Most of these results are obtained by simulation, since the models are complex to analyze due to their non-linear nature (a fact acknowledged early on in [11]). Recent work by Chazelle [4] frames these processes as a special case of what they term “one-dimensional influence systems”, and presents a technique for analyzing convergence time.

2. PRELIMINARIES

2.1 Coevolutionary Games

There are n players each with an intrinsic opinion s_i . Each player expresses an opinion z_i which is not necessarily equal to s_i . The player’s goal is to minimize its cost $C_i(\mathbf{z})$ which is a function of the players intrinsic opinion s_i and expressed opinions of all players (denoted \mathbf{z}). The social cost denoted $C(\mathbf{z})$ is sum of the players costs $\sum_i C_i(\mathbf{z})$.

We consider two different fairly general manifestations of this model termed the *symmetric co-evolutionary game* and the *K-nearest neighbor (K-NN) game*. In the symmetric coevolutionary game, a player i ’s cost function is given by

$$C_i(z_i, \mathbf{z}_{-i}) = \sum_{j \neq i} f_{ij}(z_i - z_j) + w_i g_i(z_i - s_i)$$

Here, \mathbf{z}_{-i} denotes the opinions expressed by all of the other players. Functions f_{ij} and g_i are real valued functions and remain fixed. We assume $f_{ij} = f_{ji}$ which makes the game symmetric in the pairs of players. A simple realization of this is the class of games studied by Bindel et al. [3]. In their game $g_i(z) = z^2$ and $f_{ij}(z) = w_{ij}z^2$ where w_{ij} indicates the weight put on the edge connecting players i and j .

In the K -NN game, each agent forms exactly K friends. Given expressed opinion vector \mathbf{z} , each node i forms directed links to the K nodes with smallest $|z_j - s_i|$ (breaking ties in a consistent fashion). Denote this set of friends $S(\mathbf{z}, i)$. The cost of i is:

$$C_i(z_i, \mathbf{z}_{-i}) = \sum_{j \in S(\mathbf{z}, i)} (z_j - z_i)^2 + \alpha K (z_i - s_i)^2$$

2.2 Equilibria and Local Smoothness

We analyze the equilibrium outcomes of these games. The quality of the equilibrium outcomes is assessed in comparison with the outcome that minimizes social cost. Suppose each player’s intrinsic opinion is fixed. The optimal social cost in this case is well-defined. We consider three different solution concepts each one progressively weaker than the previous. The PoA with respect to these solution concepts is the ratio of the cost of the worst equilibrium outcome to the optimal social outcome.

Notions of Equilibria.

In a pure Nash equilibrium (NE), each player’s strategy is a z_i so that for any other z , we have $C_i(z_i, \mathbf{z}_{-i}) \leq C_i(z, \mathbf{z}_{-i})$. In a *mixed* NE, a player’s strategy is a distribution σ_i over expressed opinions, so that for each z_i in the support of σ_i , we have: $\mathbb{E}_{\mathbf{z}_{-i} \sim \sigma_{-i}} [C_i(z_i, \mathbf{z}_{-i})] \leq \mathbb{E}_{\mathbf{z}_{-i} \sim \sigma_{-i}} [C_i(z, \mathbf{z}_{-i})]$. Here σ_{-i} denotes the joint distribution of other players strategies.

A correlated equilibrium σ is a distribution such that for each player i and each opinion z_i in the support of σ , $\mathbb{E}_{\mathbf{z}_{-i} \sim \sigma_{-i} | z_i} [C_i(z_i, \mathbf{z}_{-i})] \leq \mathbb{E}_{\mathbf{z}_{-i} \sim \sigma_{-i} | z_i} [C_i(z, \mathbf{z}_{-i})]$. Here $\sigma_{-i} | z_i$ denoting the distribution σ conditioned on z_i .

Since a pure Nash equilibrium is a mixed Nash equilibrium, and any mixed Nash equilibrium is a correlated equilibrium, an upper bound on the PoA of correlated equilibria also bounds the pure PoA and mixed PoA. Additionally any lower bound example with a bad pure Nash equilibrium also has a bad mixed and correlated equilibrium.

Local Smoothness.

We use the local smoothness technique developed by Roughgarden and Schoppmann [20] to prove Price of Anarchy bounds.¹ This technique can be applied only when the cost function C_i is continuously differentiable with respect to player i ’s strategy (in our case z_i).

One has to prove an inequality of the following form. For fixed profile \mathbf{o} and values $\mu < 1$, $\lambda > 0$; for every \mathbf{z} ,

$$\sum_i C_i(z_i, \mathbf{z}_{-i}) + (o_i - z_i) \frac{\partial}{\partial z_i} C_i(z_i, \mathbf{z}_{-i}) \leq \lambda C(\mathbf{o}) + \mu C(\mathbf{z}). \quad (1)$$

An extension theorem from [20] gives us the following.

THEOREM 2.1. *Let σ denote a correlated equilibrium. If equation (1) holds for any outcome \mathbf{z} with respect to a fixed outcome \mathbf{o} , then the ratio of $\mathbb{E}_{\mathbf{z} \sim \sigma} [C(\mathbf{z})]$ to $C(\mathbf{o})$ is at most $\lambda/(1 - \mu)$. In particular, when \mathbf{o} denotes the optimal outcome, the correlated PoA is at most $\lambda/(1 - \mu)$.*

A PoA bound obtained using the local smoothness technique is referred as *robust PoA*.

3. THE SYMMETRIC GAME

In this section, we consider the symmetric coevolutionary game. Recall that an agent i ’s cost function is given by

$$C_i(z_i, \mathbf{z}_{-i}) = \sum_{j \neq i} f_{ij}(z_i - z_j) + w_i g_i(z_i - s_i)$$

¹Note that the local smoothness technique is slightly different from the smoothness technique of [19]. We apply the local smoothness technique here since the smoothness technique does not yield tight bounds.

Here, \mathbf{z}_{-i} denotes the opinions expressed by all of the other agents.

We assume $f_{ij} = f_{ji}$ which makes the game symmetric in pairs of players. We further assume that functions f_{ij} and g_i are convex and continuously differentiable. We also assume that they are symmetric i.e. $f_{ij}(x) = f_{ij}(-x)$, $g_i(x) = g_i(-x)$ and $g_i(0) = 0$.

Under these assumptions, we provide an exact PoA bound for this model. The upper bound will be obtained by an application of the local smoothness inequalities (Section 3.1). We show that this bound is always at most 2. However, this bound will not have a closed form for most functions. The challenge then is to show a matching lower bound example. We use the functional form of the inequalities in the upper bound to argue the existence of exact equations of a certain type, which we use to carefully construct our lower bound example (Section 3.2). We finally show that for $f(x) = g(x) = |x|^\alpha$, we can compute a closed form expression that is exact (Section 3.3). We use this closed form to develop an interpretation for how the PoA varies as α varies, and finally show that if f, g are either not convex or not differentiable, the PoA is unbounded (Section 3.4).

As a first step, we briefly understand when a pure Nash equilibrium can be guaranteed to exist in this model.

Equilibrium existence.

In this game, when the cost functions are convex we can easily prove that a pure Nash equilibrium exists. This follows from the fact this is a potential game with the following potential function $\phi(\mathbf{z}) = \sum_i w_i g_i(z_i - s_i) + \sum_{i < j} f_{ij}(z_i - z_j)$, and the fact that all the weighting functions are convex and bounded from below. Further, the pure Nash equilibrium in this game is unique. Moreover, since each player always has a unique best response any mixed Nash equilibrium is also a pure Nash equilibrium. Finally, a result of Neyman [17] shows that the unique correlated equilibrium also coincides with the pure Nash equilibrium, so that the robust PoA coincides with pure Nash PoA in this case.

3.1 Robust PoA Upper Bound

We obtain a general characterization of the PoA in terms of the weighting functions. Suppose we are given sets \mathcal{F} and \mathcal{G} from which the weighing function f_{ij} and g_i for each i, j are chosen. We present a functional form for the PoA upper bound; this form does not require f, g to be convex. We first identify a set of tuples of constants (λ, μ) for which inequality (1) can be proved.

For a fixed triple (x, y, f) of reals $x, y \geq 0$ and function $f \in \mathcal{F}$, let $\mathcal{H}_{x,y,f}$ denote the set of (λ, μ) satisfying $f(x) + \frac{(y-x)}{2}f'(x) \leq \lambda f(y) + \mu f(x)$. Note that this is a half-plane in the (λ, μ) space. The boundary of this half-plane, denoted $\partial\mathcal{H}_{x,y,f}$, is the set of (λ, μ) satisfying $f(x) + (y-x)/2 \cdot f'(x) = \lambda f(y) + \mu f(x)$. Similarly for triple (u, v, g) of reals $u, v \geq 0$ and function $g \in \mathcal{G}$, we denote by $\mathcal{H}_{u,v,g}$ the set of (λ, μ) satisfying $g(u) + (v-u)g'(u) \leq \lambda g(v) + \mu g(u)$. The boundary of this half-plane is denoted $\partial\mathcal{H}_{u,v,g}$.

To compute the PoA upper bound, we define the following sets.

$$\mathcal{A}_1 := \{(\lambda, \mu) : (\lambda, \mu) \in \mathcal{H}_{x,y,f} \forall f \in \mathcal{F}, x, y \geq 0\} \quad (2)$$

$$\mathcal{A}_2 := \{(\lambda, \mu) : (\lambda, \mu) \in \mathcal{H}_{u,v,g} \forall g \in \mathcal{G}, u, v \geq 0\} \quad (3)$$

The sets $\mathcal{A}_1, \mathcal{A}_2$ are convex regions in the (λ, μ) plane formed from the intersection of many half-planes. We next prove

the local smoothness inequality (1) for λ, μ that lie in sets $\mathcal{A}_1, \mathcal{A}_2$. The inequality is proved by breaking it up into many inequalities each of which is of the form of the constraints that define sets $\mathcal{A}_1, \mathcal{A}_2$.

LEMMA 3.1. *Let (λ, μ) in $\mathcal{A}_1 \cap \mathcal{A}_2$. Then for a fixed outcome \mathbf{o} and any \mathbf{z} ,*

$$\begin{aligned} \sum_i C_i(z_i, \mathbf{z}_{-i}) + (o_i - z_i) \frac{\partial}{\partial z_i} C_i(z_i, \mathbf{z}_{-i}) \\ \leq \lambda C(\mathbf{o}) + \mu C(\mathbf{z}). \end{aligned}$$

PROOF. We begin with the left hand side.

$$\begin{aligned} \sum_i C_i(z_i, \mathbf{z}_{-i}) + (o_i - z_i) \frac{\partial}{\partial z_i} C_i(z_i, \mathbf{z}_{-i}) \\ = \sum_i \left[\sum_{j \neq i} f_{ij}(z_i - z_j) + w_i g_i(z_i - s_i) \right] \\ + (o_i - z_i) \left[\sum_{j \neq i} f'_{ij}(z_i - z_j) + w_i g'_i(z_i - s_i) \right] \end{aligned}$$

Regrouping terms corresponding to each pair (i, j) we obtain,

$$\begin{aligned} \sum_{i \neq j} [2f_{ij}(z_i - z_j) + (o_i - z_i)f'_{ij}(z_i - z_j) + (o_j - z_j)f'_{ij}(z_j - z_i)] \\ + \sum_i w_i [g_i(z_i - s_i) + (o_i - z_i)g'_i(z_i - s_i)] \\ = \sum_{i \neq j} 2 \left[f_{ij}(z_i - z_j) + \frac{1}{2}[(o_i - o_j) - (z_i - z_j)]f'_{ij}(z_i - z_j) \right] \\ + \sum_i w_i [g_i(z_i - s_i) + ((o_i - s_i) - (z_i - s_i))g'_i(z_i - s_i)] \\ = 2 \sum_{i \neq j} B_{ij} + \sum_i w_i A_i, \end{aligned}$$

where,

$$B_{ij} = f_{ij}(z_i - z_j) + \frac{1}{2}[(o_i - o_j) - (z_i - z_j)]f'_{ij}(z_i - z_j), \text{ and } A_i = g_i(z_i - s_i) + ((o_i - s_i) - (z_i - s_i))g'_i(z_i - s_i).$$

First, let us bound B_{ij} . Recall the definition of the set \mathcal{A}_1 . For any (λ, μ) in \mathcal{A}_1 ,

$$\begin{aligned} B_{ij} = f_{ij}(z_i - z_j) + \frac{1}{2}[(o_i - o_j) - (z_i - z_j)]f'_{ij}(z_i - z_j) \\ \leq \lambda f_{ij}(o_i - o_j) + \mu f_{ij}(z_i - z_j). \end{aligned}$$

This follows from the fact that (λ, μ) lie in the half-plane $\mathcal{H}_{|z_i - z_j|, |o_i - o_j|, f_{ij}}$. Similarly for any (λ, μ) in \mathcal{A}_2 ,

$$\begin{aligned} A_i = g_i(z_i - s_i) + ((o_i - s_i) - (z_i - s_i))g'_i(z_i - s_i) \\ \leq \lambda g_i(o_i - s_i) + \mu g_i(z_i - s_i), \end{aligned}$$

which follows from the fact that (λ, μ) lie in the half-plane $\mathcal{H}_{|z_i - s_i|, |o_i - s_i|, g_i}$.

We can then conclude that for any $(\lambda, \mu) \in \mathcal{A}_1 \cap \mathcal{A}_2$,

$$\sum_i C_i(z_i, \mathbf{z}_{-i}) + (o_i - z_i) \frac{\partial}{\partial z_i} C_i(z_i, \mathbf{z}_{-i}) \leq \lambda C(\mathbf{o}) + \mu C(\mathbf{z}).$$

□

With this lemma and the extension Theorem 2.1 we can conclude that for any $(\lambda, \mu) \in \mathcal{A}_1 \cap \mathcal{A}_2$, $\lambda/(1 - \mu)$ is an upper bound on the correlated POA. We pick the best upper bound and denote it by ζ .

DEFINITION 3.2. For sets of functions \mathcal{F}, \mathcal{G} let sets $\mathcal{A}_1, \mathcal{A}_2$ be defined as in equations 2.

$$\zeta(\mathcal{F}, \mathcal{G}) = \inf\{\lambda/(1-\mu) : (\lambda, \mu) \in \mathcal{A}_1 \cap \mathcal{A}_2, \mu < 1\}. \quad (4)$$

THEOREM 3.3. For weighting function classes $(\mathcal{F}, \mathcal{G})$, the (robust) PoA is at most $\zeta(\mathcal{F}, \mathcal{G})$ given by Definition 3.2.

As a corollary, $\zeta(\mathcal{F}, \mathcal{G}) \leq 2$ when the functions in \mathcal{F} and \mathcal{G} are convex and differentiable.

COROLLARY 3.4. When functions in \mathcal{F} and \mathcal{G} are convex and differentiable, the robust POA is at most 2.

PROOF. Consider a function $f \in \mathcal{F}$. Since f is convex we have, $f(x) + (y-x)f'(x) \leq f(y)$. We then see that

$$f(x) + (y-x)/2f'(x) \leq 1/2f(y) + 1/2f(x) \leq f(y) + 1/2f(x)$$

and the tuple $(1, 1/2)$ belongs to \mathcal{A}_1 . In fact, even $(1/2, 1/2) \in \mathcal{A}_1$.

Similarly for $g \in \mathcal{G}$, convexity implies $g(x) + (y-x)g'(x) \leq g(y)$. This implies

$$g(x) + (y-x)g'(x) \leq g(y) \leq g(y) + tg(x)$$

i.e. $(1, t)$ is in \mathcal{A}_2 for any $t \geq 0$. We conclude that $(1, 1/2) \in \mathcal{A}_1 \cap \mathcal{A}_2$. For $\lambda = 1$ and $\mu = 1/2$, $\lambda/(1-\mu) = 2$. \square

This bound can also be obtained from the observation that the potential $\phi(\mathbf{z})$ for any strategy \mathbf{z} is within a factor of 2 of the social cost $\sum_i C_i(\mathbf{z})$, so that the minimizer of ϕ (Nash equilibrium) has social cost within factor 2 of the optimal social cost. Since the unique correlated equilibrium coincides with the pure Nash equilibrium [17], the robust PoA bound is also factor 2.

3.2 Tight Lower Bounds

In this section, we provide a PoA lower bound that matches the upper bound constructed above. This part will require f, g to be convex, differentiable, and symmetric.

We will provide a lower bound for sets of functions \mathcal{F}, \mathcal{G} that are closed under two operations: *scaling* i.e., $f \in \mathcal{F}$ implies $\alpha f \in \mathcal{F}$; and *dilation* i.e., $f(x) \in \mathcal{F}$ implies $f(\alpha x) \in \mathcal{F}$ for any α . We call a set of functions a *functional family* if they are related to each other through scaling and/or dilation. An example of a functional family is polynomials of degree at most d .

Note that the operations of scaling and dilation do not introduce new constraints in the definitions of sets \mathcal{A}_1 and \mathcal{A}_2 , which means the PoA bound will be the same for any functions from a functional family (and hence Definition 3.2 applies to a family). For illustration, consider function $f_1(x) = f(\alpha x)$ in set \mathcal{F} and the half-plane $\mathcal{H}_{x_1, y_1, f_1}$:

$$f_1(x_1) + (y_1 - x_1)f_1'(x_1) \leq \lambda f_1(y_1) + \mu f_1(x_1)$$

Note that if $f_1'(x) = \alpha f'(x)$, then the above constraint is equivalent to:

$$f(\alpha x_1) + (\alpha y_1 - \alpha x_1)f'(\alpha x_1) \leq \lambda f(\alpha y_1) + \mu f(\alpha x_1).$$

Setting $x = \alpha x_1$ and $y = \alpha y_1$ gives us the half-plane $\mathcal{H}_{x, y, f}$. Also note that $y/x = y_1/x_1$. Similar claim follows for scaling using the property that for $f_1(x) = \alpha f(x)$, $f_1'(x) = \alpha f'(x)$. We can also verify that no new constraints are introduced in set \mathcal{A}_2 if set \mathcal{G} is assumed to be closed under dilation. We consider sets \mathcal{F}, \mathcal{G} that contain a countable number of families of functions.

Next, we will prove the following lower bound theorem.

THEOREM 3.5. Let \mathcal{F}, \mathcal{G} denote a set of countable families of functions. Let $\zeta(\mathcal{F}, \mathcal{G})$ be a price of anarchy upper bound as defined in Definition 3.2. Then there exist an instance using weighing functions from \mathcal{F}, \mathcal{G} with price of anarchy arbitrarily close to the upper bound $\zeta(\mathcal{F}, \mathcal{G})$.

Step 1: Tight Equations in $\mathcal{A}_1 \cap \mathcal{A}_2$.

First step in constructing the lower bound and hence proving the theorem is analyzing the set of constraints that define $\mathcal{A}_1, \mathcal{A}_2$ and in turn the upper bound $\zeta(\mathcal{F}, \mathcal{G})$. Insights into the structure of these constraints can be used to construct the lower bound example. Since the cost functions are continuous it is sufficient to focus on the case when values of x, y are rational. Any constraint corresponding to irrational values of x or y can be arbitrarily closely approximated using a constraint corresponding to nearby rational values. Since the set of rational numbers is countable and we assume that we have a countable set of weighing function families, the set of constraints that define sets $\mathcal{A}_1, \mathcal{A}_2$ are countable.

We consider an arbitrary order on these constraints. Let \mathcal{C}_n denote the feasible region formed from the first n constraints. Let $\zeta_n = \inf\{\lambda/(1-\mu) : (\lambda, \mu) \in \mathcal{C}_n\}$ and (λ_n, μ_n) denote the values that define ζ_n . Note that we consider an infimum over a set that is bounded from below. Hence, the minimum is attained and further, (λ_n, μ_n) are well-defined. The next lemma establishes a few more properties of (λ_n, μ_n) .

LEMMA 3.6. For $\mathcal{C}_n, (\lambda_n, \mu_n)$ and ζ_n be defined as above. If $\zeta_n > 1$ then there exists values x_1, y_1, x_2, y_2 and functions $f \in \mathcal{F}, g \in \mathcal{G}$ such that

$$f(x_1) + \frac{(y_1 - x_1)}{2} f'(x_1) = \lambda_n f(y_1) + \mu_n f(x_1) \quad \text{and}$$

$$g(x_2) + (y_2 - x_2)g'(x_2) = \lambda_n g(y_2) + \mu_n g(x_2).$$

Moreover if $w_1 = x_1/y_1$ and $w_2 = x_2/y_2$ then $(w_1 - 1)(w_2 - 1) \leq 0$.

PROOF. Recall that ζ_n is the minimum value of $\lambda/(1-\mu)$ over $(\lambda, \mu) \in \mathcal{C}_n$. The set \mathcal{C}_n is a convex region formed from the intersection of a number of half-planes. It is easy to see that the minimum must be on the boundary of this region.

Consider a constraint corresponding to function $f \in \mathcal{F}$. We can rewrite the constraint as

$$\frac{\lambda}{1-\mu} \geq \frac{f(x)}{f(y)} + \frac{(y-x)f'(x)}{2(1-\mu)f(y)}.$$

The function f is non-decreasing on the positive domain. Hence, $f'(x) \geq 0$. Note that along the boundary, the way $\lambda/(1-\mu)$ varies as μ increases is governed by the sign on $(y-x)$. In particular, the value of $\lambda/(1-\mu)$ is minimized as μ approaches 0 if $y > x$ and as μ approaches 1 if $y < x$. Moreover, as noted in the proof of Corollary 3.4 $(1/2, 1/2)$ lies in this half-plane. This in particular implies that the minimum value of $\lambda/(1-\mu)$ on this half-plane boundary is less than 1. Also since any other half-plane corresponding to $f_2 \in \mathcal{F}$ also contains $(1/2, 1/2)$, the two half-planes must intersect at a point (λ, μ) such that $\lambda/(1-\mu) < 1$.

Similarly, consider a constraint corresponding to a function $g \in \mathcal{G}$. We can rewrite that constraint as,

$$\frac{\lambda}{(1-\mu)} \geq \frac{g(x)}{g(y)} + \frac{(y-x)g'(x)}{(1-\mu)g(y)}.$$

The value of $\lambda/(1-\mu)$ is minimized along this half plane's boundary as μ approaches 0 if $y > x$ and as μ approaches 1

if $y < x$. On any single line the value $\lambda/(1-\mu)$ is monotone in μ . Therefore, the minimum value is obtained at the intersection of two lines. Also from the proof of Corollary 3.4, $(1,0)$ lies in this half-plane. We can conclude that on this half-plane boundary the minimum value of $\lambda/(1-\mu) < 1$. And any two half-planes corresponding to functions in \mathcal{G} intersect in (λ, μ) such that $\lambda/(1-\mu) < 1$.

Based on the analysis above, we can conclude the following. First, the minimal value of $\lambda/(1-\mu)$ must occur either at the intersection of two half-plane boundaries or on one half-plane boundary as μ approaches 0 or 1. But since the minimum value of $\lambda/(1-\mu)$ is less than 1 along any half-plane boundary and we are considering the case when $\zeta_n > 1$, the latter does not happen. We also noted above that any two half planes corresponding to two functions in \mathcal{F} or two functions in \mathcal{G} intersect at a point with $\lambda/(1-\mu) < 1$. Thus if $\zeta_n > 1$, (λ_n, μ_n) must lie at the intersection of two half-planes boundaries that correspond to one function from set \mathcal{F} and one function from set \mathcal{G} . Suppose they correspond to triples x_1, y_1, f and x_2, y_2, g , respectively. Note that we can always scale or dilate the function and change x, y proportionately to represent the same half-plane.

We also noted earlier that for each of the type of constraints, if $x > y$, then the minimum is obtained along the boundary as μ approaches 1 and is obtained as μ approaches 0 if $x < y$. Since the minimum value on the set, ζ_n , is obtained at the intersection of these two half-plane boundaries, it must be true that for one of them (the left one in particular) $x > y$ and for other $x < y$. We set $w_1 = x_1/y_1$ and $w_2 = x_2/y_2$ and conclude $(w_1 - 1)(w_2 - 1) \leq 0$. \square

Step 2: Lower Bound Instance.

The next step is constructing the lower bound. We will prove the following theorem, which will complete the proof of Theorem 3.5.

THEOREM 3.7. *Let triples x_1, y_1, f and x_2, y_2, g with $w_1 = x_1/y_1$ and $w_2 = x_2/y_2$ be such that*

$$\begin{aligned} f(x_1) + \frac{(y_1 - x_1)}{2} f'(x_1) &= \lambda_n f(y_1) + \mu_n f(x_1), \\ g(x_2) + (y_2 - x_2) g'(x_2) &= \lambda_n g(y_2) + \mu_n g(x_2), \end{aligned}$$

and $(1-w_1)(1-w_2) \leq 0$. Then we can construct an instance of the consensus game using just scaled or dilated versions f and g such that its pure Nash PoA is at least $\lambda_n/(1-\mu_n)$.

To prove this theorem we first describe the construction of the instance and then verify that it has the desired properties.

EXAMPLE 3.1. *Consider a game with three players having intrinsic opinions $s_1 = 0, s_2 = 1, s_3 = 2$, respectively.*

The Nash equilibrium outcome \mathbf{z} will be as follows,

$$z_1 = \frac{(1-w_1)w_2}{w_2-w_1}, z_2 = 1, z_3 = 2 - \frac{(1-w_1)w_2}{w_2-w_1}. \quad (5)$$

The optimal outcome is as follows,

$$o_1 = \frac{1-w_1}{w_2-w_1}, o_2 = 1, o_3 = 2 - \frac{1-w_1}{w_2-w_1}. \quad (6)$$

We will choose functions f_1, g_1 such that

$$\begin{aligned} f_1(z_2 - z_1) + ((z_2 - z_1) - (o_2 - o_1)) f_1'(z_2 - z_1) \\ = \lambda_n f_1(o_2 - o_1) + \mu_n f_1(z_2 - z_1) \\ g_1(z_1 - s_1) + (z_1 - o_1) g_1'(z_1 - s_1) \\ = \lambda_n g_1(o_1 - s_1) + \mu_n g_1(z_1 - s_1) \end{aligned} \quad (7)$$

Notice that $(z_2 - z_1)/(o_2 - o_1) = w_1 = y_1/x_1$ and $(z_1 - s_1)/(o_1 - s_1) = w_2 = y_2/x_2$. So we can identify f_1, g_1 which will be dilated versions of f, g .

Choose w as follows,

$$w = g_1' \left(\frac{(1-w_1)w_2}{w_2-w_1} \right) / f_1' \left(\frac{w_1(w_2-1)}{w_2-w_1} \right).$$

Note that as long as g_1 and f_1 are increasing functions and $(1-w_1)(w_2-1) \geq 0$ then $w \geq 0$.

We define a set of neighbors $N(i)$ for each player i . $N(1) = \{2\}, N(2) = \{1, 3\}, N(3) = \{2\}$. Finally the cost functions for each player would be as follows:

$$C_i(s_i, \mathbf{z}) = w \sum_{j \in N(i)} f_1(z_i - z_j) + g_1(z_i - s_i)$$

In the following lemmas we verify various properties of the above example.

LEMMA 3.8. *Outcome \mathbf{z} from equation 5 is a pure Nash equilibrium*

PROOF. For player 1, its cost in the Nash equilibrium is $w f_1(1-z_1) + g_1(z_1)$ which is minimized when $g_1'(z_1) - w f_1'(1-z_1) = 0$. If $z_1 = (1-w_1)w_2/(w_2-w_1)$ then $(1-z_1) = w_1(w_2-1)/(w_2-w_1)$. We see that the condition is then satisfied by the choice of w .

Similarly for player 3, its cost in the Nash equilibrium outcome is $w f_1(z_3-1) + g_1(2-z_3)$ which is minimized when $w f_1'(z_3-1) - g_1'(2-z_3) = 0$. If $z_3 = 2 - (1-w_1)w_2/(w_2-w_1)$ then $z_3-1 = w_1(w_2-1)/(w_2-w_1)$ and $2-z_3 = (1-w_1)w_2/(w_2-w_1)$. The condition is once more satisfied by the choice of w .

Before we proceed to player 2, let's prove that $g_1'(0) = 0$. Recall that g_1 is symmetric i.e. $g_1(x) = g_1(-x)$. Then $g_1'(x) = -g_1'(-x)$. And at 0, $g_1'(0) = -g_1'(0)$. We conclude that $g_1'(0) = 0$.

For player 2, the Nash equilibrium cost is $w f_1(z_2 - z_1) + w f_1(z_3 - z_2) + g_1(z_2 - 1)$, which is minimized when $w f_1'(z_2 - z_1) - w f_1'(z_3 - z_2) + g_1'(z_2 - 1) = 0$. For $z_2 = 1$, $g_1'(z_2 - 1) = g_1'(0) = 0$. By choice of z_1 and z_3 , $z_3 - 1 = 1 - z_1$. Hence, $f_1'(z_2 - z_1) = f_1'(z_3 - z_2)$. \square

Next we verify that the costs of the two outcomes have the desired ratio. Since the (real) optimal outcome will have cost even lower than the reported outcome \mathbf{o} , we can conclude that the PoA is at least $\lambda_n/(1-\mu_n)$.

LEMMA 3.9. *The ratio of cost of the two outcomes \mathbf{z} and \mathbf{o} from equations (5) and (6) is $\lambda_n/(1-\mu_n)$.*

PROOF. Recall that we chose f_1, g_1 to satisfy equations (7). Note that we have defined z_1, z_2, z_3 such that $z_2 - z_1 = z_3 - z_2$ and $z_1 - s_1 = s_3 - z_3$. Similarly o_1, o_2, o_3 are such that $o_2 - o_1 = o_3 - o_2$ and $o_1 - s_1 = o_2 - s_2$. Finally, $o_2 = z_2 = s_2$ hence $o_2 - s_2 = z_2 - s_2 = 0$.

Thus for each pair $(1, 2), (2, 3)$ we have

$$\begin{aligned} 2f_1(z_i - z_j) + ((o_i - o_j) - (z_i - z_j)) f_1'(z_i - z_j) \\ = \lambda_n 2f_1(o_i - o_j) + \mu_n 2f_1(z_i - z_j). \end{aligned}$$

and for players 1 and 3,

$$\begin{aligned} g_1(z_i - s_i) + (o_i - z_i) g_1'(z_i - s_i) \\ = \lambda_n g_1(o_i - s_i) + \mu_n g_1(z_i - s_i). \end{aligned}$$

On the other hand for player 2, $g_1(0) + 0g'_1(0) = \lambda_n g_1(0) + \mu_n g_1(0)$ since $g_1(0) = 0$. For each player i , $\sum_{j \in N(i)} w f'_1(z_i - z_j) + g'_1(z_i - s_i) = 0$. Recall that this is what we verified in the proof of Lemma 3.8. Thus if we combine first expression multiplied by w for each adjacent pair (i, j) and second expression for all players i and simplify, we get

$$\begin{aligned} & \sum_i \left[\sum_{j \in N(i)} f_1(z_i - z_j) + w g_1(z_i - s_i) \right] \\ &= \lambda_n \sum_i \left[\sum_{j \in N(i)} f_1(o_i - o_j) + w g_1(o_i - s_i) \right] \\ &+ \mu_n \sum_i \left[\sum_{j \in N(i)} f_1(z_i - z_j) + w g_1(z_i - s_i) \right], \end{aligned}$$

i.e. $c(\mathbf{z}) = \lambda_n c(\mathbf{o}) + \mu_n c(\mathbf{z})$ and the PoA = $c(\mathbf{z})/c(\mathbf{o})$ is then $\lambda_n/(1 - \mu_n)$. \square

This completes the proof of Theorem 3.7, and hence Theorem 3.5.

3.3 Some Closed-form Bounds

We consider the case when the functions $f(x)$ and $g(x)$ are $|x|^\alpha$ for $\alpha > 1$. We calculate numerical price of anarchy bounds for these cost functions. Instead of computing the optimal λ, μ , we construct a feasible pair. We prove that the bound is tight by constructing a lower bound example with the same price of anarchy.

LEMMA 3.10. *The price of anarchy in the symmetric co-evolutionary game when $f_{ij}(x) = |x|^\alpha$ and $g_i(x) = |x|^\alpha$ for all i, j is at most $\frac{(\alpha-1)^{(\alpha-1)}}{\alpha^\alpha} \cdot \frac{(2^\alpha/(\alpha-1)-1)^\alpha}{2^{\alpha/(\alpha-1)-2}}$.*

PROOF. Our bounds work over the non-negative domain, so we pretend $f(x) = g(x) = x^\alpha$. Let us set $\beta = \alpha/(\alpha - 1)$. We will show that when both cost functions are x^α choosing $\lambda = \beta^{-(\alpha-1)} \cdot (2^\beta - 1)^{\alpha-1}/2$ and $\mu = 1 - \frac{\alpha}{2} + \frac{\alpha}{2}[1/(2^\beta - 1)]$, establishes the desired bound.

The main inequality we use in this proof is ²

$$(\alpha - 1)A^\alpha + B^\alpha \geq \alpha B \cdot A^{\alpha-1}. \quad (8)$$

This inequality is tight when $A = B$. For the function $f(x) = x^\alpha$, we have to show, $x^\alpha + (y - x)/2\alpha x^{\alpha-1} \leq \lambda y^\alpha + \mu x^\alpha$. Recall that $\beta = \alpha/(\alpha - 1)$. Then, for the choice of λ, μ mentioned earlier, it is equivalent to showing

$$\alpha(2^\beta - 1)yx^{\alpha-1} \leq \beta \cdot (\alpha - 1)x^\alpha + \beta^{-(\alpha-1)} \cdot (2^\beta - 1)^\alpha \cdot y^\alpha. \quad (9)$$

This can be established by choosing $A = x \cdot \beta^{1/\alpha}$ and $B = y \cdot \beta^{-(\alpha-1)/\alpha} \cdot (2^\beta - 1)$ in inequality 8.

Similarly for the second constraints, we have to show $x^\alpha + (y - x)\alpha x^{\alpha-1} \leq \lambda y^\alpha + \mu x^\alpha$, which is equivalent to showing

$$2\alpha(2^\beta - 1)yx^{\alpha-1} \leq \beta \cdot 2^\beta \cdot (\alpha - 1)x^\alpha + \beta^{-(\alpha-1)} \cdot (2^\beta - 1)^\alpha \cdot y^\alpha. \quad (10)$$

This can be established by choosing $A = x \cdot \beta^{1/\alpha} \cdot 2^{1/(\alpha-1)}$ and $B = y \cdot \beta^{-(\alpha-1)/\alpha} \cdot (2^\beta - 1)$ in inequality 8. Thus (λ, μ) are in the set and the PoA is at most $\lambda/(1 - \mu)$. \square

See Figure 1 for approximate values of the price of anarchy for various values of α .

²This can be easily established for $\alpha \geq 1$ using Jensen's inequality: for concave function f , $\delta < 1$, $\delta f(x) + (1 - \delta)f(y) \leq f(\delta x + (1 - \delta)y)$. We choose $f(x) = \ln x$, $x = A^\alpha$, $y = B^\alpha$, and $\delta = (\alpha - 1)/\alpha$.

Next we use the construction from Example 3.1 to show that the upper bound established above is tight.

LEMMA 3.11. *The upper bound established in Lemma 3.10 is tight.*

PROOF. We use Theorem 3.7 to prove this result. Recall, we used β to denote $\alpha/(\alpha - 1)$, which is at least 1 as long as $\alpha \geq 1$. Equation (9) is tight if $y = \beta/(2^\beta - 1)x$ i.e. $y/x = w_1 = \beta/(2^\beta - 1)$. Moreover, $w_1 \leq 1$ as long as $\beta \geq 1$. Equation (10) is tight if $y = \beta/(2^\beta - 1) \cdot 2^{\beta-1} \cdot x$. Thus $w_2 = \beta \cdot 2^{\beta-1}/(2^\beta - 1)$, which is at least 1 for $\beta \geq 1$.

Thus $(w_1 - 1)(w_2 - 1) \leq 0$ and the construction from example 3.1 can be used. \square

3.4 Discussion

The results in the previous section shed light on how the equilibria behave as nodes give more weight to nodes with close-by opinions. If the cost function is $f(x) = g(x) = |x|^\alpha$, this can be written as $f(x) = |x|^2/|x|^{2-\alpha}$. If $f(x) = |x|^2$, at a pure Nash equilibrium, each node simply takes the average of the opinions of its neighbors and its intrinsic opinion. If $f(x) = |x|^\alpha$, this can be interpreted as a node i using the link to j for averaging with probability $|z_i - z_j|^{\alpha-2}$. As α reduces, node i gives more importance to nodes with close-by opinions, and we see that this increases the PoA. Therefore, bad social outcomes are more likely to arise if nodes only consider opinions of similarly biased nodes.

We also note that if the cost functions f and g are not strictly convex, the PoA can be unbounded. This occurs even when $f(x) = g(x) = |x|$, which corresponds to each node i expressing the *median* of the opinions of its neighbors and its own opinion. We present two examples where consensus and polarization result, and both of these have high social cost.

1. The graph is a clique; each node has $s_i = 0$. All nodes expressing $z_i = 1$ is a NE with social cost strictly positive. However, the optimal outcome has $z_i = 0$ with social cost 0.
2. There are two sets of nodes V_1 and V_2 of equal size. The induced graph on V_1 (resp. V_2) is a d -regular graph. There is a $(d - 2)$ -regular bipartite graph between the set V_1 and set V_2 . All intrinsic opinions are $s = 0$; again the optimal social cost is 0. It is easy to check that any outcome where the nodes in V_1 express the same opinion z_1 and the nodes in V_2 express the same opinion z_2 is a Nash equilibrium. This outcome is polarized with social cost strictly positive.

The last example above is reminiscent of *Schelling segregation* [21] - each node only has slightly more neighbors in V_1 than in V_2 , and simply prefers to express the majority opinion. Even in this setting, the opinions can polarize dramatically with unbounded PoA.

4. ASYMMETRIC K -NN MODEL

As mentioned earlier, this model allows a node to base its set of friends on their expressed opinions. In particular, each agent forms K friends. Given expressed opinion vector \mathbf{z} , each node i forms directed links to the K agents with smallest $|z_j - s_i|$ (breaking ties in a consistent fashion). Denote this set of friends $S(\mathbf{z}, i)$. Where the context of the

α	POA
1.001	1.986
1.008	1.919
1.062	1.647
1.5	1.188
2	$9/8 = 1.125$
4	1.083
8	1.071
16	1.066
32	1.064

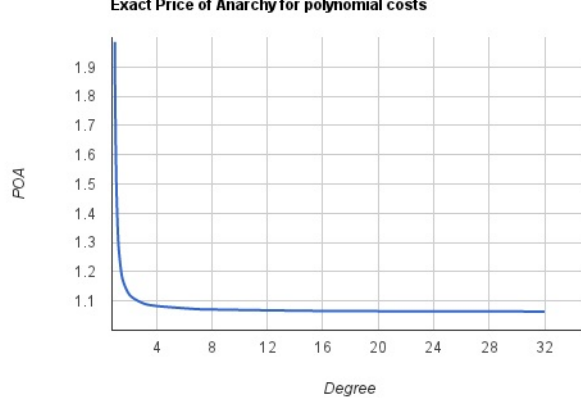


Figure 1: Exact POA bounds in the symmetric model for cost function $f(x) = |x|^\alpha$.

expressed opinion is clear, we will denote this set by $S(i)$. The cost of i is:

$$C_i(z_i, \mathbf{z}_{-i}) = \sum_{j \in S(i)} (z_j - z_i)^2 + \rho K (z_i - s_i)^2$$

Here, the set $S(\mathbf{z}, i)$ is fixed given \mathbf{z}_{-i} and does not depend on z_i . Note that though the cost $C_i(z_i, \mathbf{z}_{-i})$ is discontinuous in \mathbf{z}_{-i} , it is smooth in z_i for fixed \mathbf{z}_{-i} , which is sufficient for local smoothness arguments we use below. We also note that in a pure Nash equilibrium, given a setting of \mathbf{z}_{-i} , agent i sets z_i to be the weighted average of the opinions in $S(\mathbf{z}, i)$ and s_i as follows:

$$z_i = \frac{1}{1 + \rho} \left(\frac{\sum_{j \in S(\mathbf{z}, i)} z_j}{K} + \rho s_i \right)$$

We show (Section 4.1) that for $\rho = 1 + \epsilon$ where $\epsilon > 0$, the robust PoA of this game is at most $\frac{(7+\epsilon)(2+\epsilon)}{\epsilon(1+\epsilon)}$. This implies the PoA of correlated and mixed Nash equilibria is at most a constant *regardless* of the value of K , assuming $\rho > 1$. The PoA reduces as ρ increases and tends to 1 as $\rho \rightarrow \infty$. This shows that as nodes give larger and larger weight to their intrinsic opinion, the PoA improves, and is in fact bounded even when nodes place roughly the same weight on their own opinion as all their neighbors' put together.

In Section 4.2, we show that a pure NE is not guaranteed to exist for this game, showing the need for a local smoothness analysis. We also complement our upper bound result by showing that for $\rho < 1$, the PoA is at least $1/\rho^2$; in other words, the PoA is not bounded as ρ reduces. Finally, in Section 4.3, we present a generalization of this game and show that the generalization admits a pure NE.

4.1 Robust PoA Bound

The optimal solution chooses \mathbf{o} to minimize $\sum_i C_i(o_i, o_{-i})$. We will first show that for any ρ , setting $z_i = s_i$ is a $\frac{\rho+6}{\rho}$ approximation to OPT , the optimal solution.

Let $S^*(i)$ denote the K closest o_j to s_i in OPT (where we exclude $i = j$). Similarly, let $Q(i)$ denote the K closest s_j to s_i (where we again exclude $i = j$). The following lemma is immediate since the s_j lie on a line.

LEMMA 4.1. *Each j appears in at most $2K$ sets $Q(i)$.*

The following lemma bounds OPT , with the bound being exact as $\rho \rightarrow \infty$.

LEMMA 4.2. *For $\rho \geq 0$, we have*

$$OPT \geq \frac{\rho}{\rho + 6} \sum_i \sum_{j \in Q(i)} (s_j - s_i)^2$$

PROOF. We have

$$C_i(o_i, o_{-i}) = \sum_{j \in S^*(i)} (o_j - o_i)^2 + \rho K (o_i - s_i)^2$$

Fix any $j \in S^*(i)$. There are two mutually exclusive and exhaustive cases:

Case 1. Suppose there is no $s_k \in Q(i)$ such that s_k is on the same side of s_i as o_j and such that s_k is further away from s_i than o_j . In this case, there exists a unique $k' \in Q(i)$ so that $(o_j - s_i)^2 \geq (s_{k'}^2 - s_i)^2$.

CLAIM 4.3.

$$(o_j - o_i)^2 + \frac{\rho}{3} (o_i - s_i)^2 \geq \frac{\rho}{\rho + 3} (o_j - s_i)^2 \geq \frac{\rho}{\rho + 6} (s_{k'}^2 - s_i)^2$$

PROOF. Set $|o_j - o_i| = r$, and $|o_i - s_i| = 1$, so that $|o_j - s_i| = r + 1$. Then the LHS is $r^2 + \rho/3 \geq \frac{\rho}{\rho+3} (r+1)^2$ for all $r, \rho \geq 0$. \square

Case 2. In the other case, observe that for any j , there are at most $2K$ i for which o_j is closer to s_i than some $k \in Q(i)$ and s_k and o_j both lie on the same side of s_i . This follows from Lemma 4.1. Now split the term $2K \frac{\rho}{3} (o_j - s_j)^2$ in the cost C_j into $2K$ equal parts and assign one part to each such i to be added to the term $(o_j - o_i)^2 + \frac{\rho}{3} (o_i - s_i)^2$.

CLAIM 4.4.

$$(o_j - o_i)^2 + \frac{\rho}{3} (o_j - s_j)^2 + \frac{\rho}{3} (o_i - s_i)^2 \geq \frac{\rho}{\rho + 6} (s_i - s_j)^2$$

PROOF. Set $|o_j - o_i| = r$, and $|o_j - s_j| = |o_i - s_i| = 1$, so that $|s_i - s_j| = r + 2$. Then the LHS is $r^2 + 2\rho/3 \geq \frac{\rho}{\rho+6} (r+2)^2$ for all $r, \rho \geq 0$. \square

Either $j \in Q(i)$ or we can replace j with a unique $k \in Q(i)$ so that $(s_i - s_j)^2 \geq (s_i - s_k)^2$

Therefore, in either case, summing over all i , it is easy to check that we have:

$$\begin{aligned} OPT &\geq \sum_i \sum_{j \in S^*(i)} \left((o_j - o_i)^2 + \frac{\rho}{3}(o_i - s_i)^2 + \frac{\rho}{3}(o_j - s_j)^2 \right) \\ &\geq \frac{\rho}{\rho + 6} \sum_{j \in Q(i)} (s_j - s_i)^2 \end{aligned}$$

This completes the proof of Lemma 4.2. We will now prove the following theorem:

THEOREM 4.5. *For $\rho = (1 + \epsilon)$ where $\epsilon > 0$, the robust PoA of the K -NN game is at most $\frac{(7+\epsilon)(2+\epsilon)}{\epsilon(1+\epsilon)}$.*

PROOF. From Lemma 4.2, the vector $\mathbf{z} = \vec{s}$ is a $\frac{7+\epsilon}{1+\epsilon}$ approximation to OPT . The local smoothness argument hinges on showing the following for any opinion vector \mathbf{z} .

$$\sum_i C_i(z_i, \mathbf{z}_{-i}) + (s_i - z_i) \frac{\partial}{\partial z_i} C_i(z_i, \mathbf{z}_{-i}) \leq \lambda C(\vec{s}) + \mu C(\mathbf{z})$$

If the above holds for $\mu < 1$, the robust price of anarchy will be $\frac{7+\epsilon}{1+\epsilon} \frac{\lambda}{1-\mu}$, where the factor $\frac{7+\epsilon}{1+\epsilon}$ is because \vec{s} is a $\frac{7+\epsilon}{1+\epsilon}$ approximation. We will set $\mu = 0$ so that the PoA bound is $\frac{7+\epsilon}{1+\epsilon} \lambda$.

Fix an expressed opinion vector z and let $S(i) = S(i, \mathbf{z}_{-i})$ be the K agents j whose z_j are closest to s_i . We need to show:

$$\begin{aligned} &\sum_i \left(\rho K (z_i - s_i)^2 + \sum_{j \in S(i)} (z_i - z_j)^2 \right) \\ &+ \sum_i \left(2(s_i - z_i) \left(\rho K (z_i - s_i) + \sum_{j \in S(i)} (z_i - z_j) \right) \right) \\ &\leq \sum_i \lambda \sum_{j \in Q(i)} (s_j - s_i)^2 \end{aligned}$$

For a fixed $j \in S(i)$, $(z_i - z_j)^2 + (z_i - s_i)^2 + 2(s_i - z_i)(z_i - z_j) = (s_i - z_j)^2$. Therefore, it is sufficient to show:

$$\begin{aligned} &\sum_i \sum_{j \in S(i)} (s_i - z_j)^2 \\ &\leq \sum_i \left((\rho + 1)K (s_i - z_i)^2 + \lambda \sum_{j \in Q(i)} (s_j - s_i)^2 \right) \end{aligned}$$

Observe that since $S(i)$ are the K closest z_j to s_i , we have $\sum_{j \in S(i)} (s_i - z_j)^2 \leq \sum_{j \in Q(i)} (s_i - z_j)^2$. Next we will use the following inequality $(a+b)^2 \leq (d^2+1)a^2 + (1/d^2+1)b^2$ for any $a, b, d \geq 0$. This inequality is equivalent to $(da - b/d)^2 \geq 0$ and is hence true. Now substituting $a = (s_i - s_j)$, $b = (s_j - z_j)$ and $d^2 = (\rho - 1)/2$, we get,

$$\begin{aligned} &\sum_{j \in Q(i)} (s_i - z_j)^2 \\ &\leq \sum_{j \in Q(i)} \left(\left(1 + \frac{2}{\rho - 1} \right) (s_j - s_i)^2 + \frac{(\rho + 1)}{2} \cdot (s_j - z_j)^2 \right) \end{aligned}$$

Summing the above inequality over all i and observing that each j appears in at most $2K$ $Q(i)$ using Lemma 4.1:

$$\begin{aligned} &\sum_i \sum_{j \in S(i)} (s_i - z_j)^2 \\ &\leq \sum_i \left(\left(1 + \frac{2}{\rho - 1} \right) \sum_{j \in Q(i)} (s_j - s_i)^2 + K(\rho + 1)(s_i - z_i)^2 \right) \end{aligned}$$

This implies for $\rho = 1 + \epsilon$, $\lambda = (1 + 2/\epsilon)$. This completes the proof. \square

4.2 Lower Bounds

We will now present some lower bounds for the case where $K = 1$, which will complement the robust PoA bound presented above. We will show that the K -NN game need not admit to a pure Nash equilibrium. Furthermore, for $\rho < 1$, the PoA is at least $\frac{1}{\rho^2}$. In other words, the PoA deteriorates for small ρ , implying our upper bound is close to best possible.

We first simplify notation. For expressed opinions \mathbf{z} , let $\sigma(i) = \operatorname{argmin}_{j \neq i} |z_j - s_i|$ where ties are broken arbitrarily but consistently. Define

$$C_i(z_i, \mathbf{z}_{-i}) = (z_i - z_{\sigma(i)})^2 + \rho(z_i - s_i)^2$$

We call this the *nearest neighbor game*. Note that in a pure Nash equilibrium, given \mathbf{z}_{-i} , agent i simply sets $z_i = (z_{\sigma(i)} + \rho s_i)/(1 + \rho)$.

PROPOSITION 4.1. *In the instance with three players when players' intrinsic opinions are $s_1 = 0$, $s_2 = 1/2$, and $s_3 = 1$, and weight $\rho = 1$, a pure NE does not exist when player can express any real number in $[0, 1]$.*

PROOF. Suppose a pure Nash equilibrium exists. Let a, b , and c denote the expressed opinions of the three players in this equilibrium, respectively. First note that c cannot be less than both a and b , as it is the average of one of them with 1. Similarly a cannot be bigger than both b and c . It is also easy to check that b must necessarily lie between a and c , else one of them has a feasible deviation, which forces the ordering to be $a < b < c$.

If $a < b < c$, the first player points to the second player and the third player also points to the second player. Hence $a = b/2$ and $c = (1 + b)/2$. There are two cases. Suppose player 2 points at player 1. Then $b = (1/2 + a)/2 = (1/2 + b/2)/2$. Solving this we get that $b = 1/3$ and $a = 1/6$. But then $c = 2/3$ and player 2 should point at player 3 instead. On the other hand if player 2 does point to player 3, then $b = (1/2 + c)/2 = (1 + b/2)/2$ and $b = 2/3$. In this case, $a = 1/3$ and $c = 5/6$ and player 2 should point to player 1 instead. Thus an equilibrium does not exist. \square

In fact, the above proof shows a pure Nash equilibrium need not exist even when there are three players and seven possible opinions. Our local smoothness proof circumvents this impossibility result to bound the price of anarchy of mixed Nash and correlated equilibria, both of which are always guaranteed to exist.

We next present a lower bound on the PoA for small values of ρ which motivates the need for considering $\rho \geq 1$ for presenting our upper bound above.

PROPOSITION 4.2. *For $\rho \leq 1$, the (robust) PoA of the nearest neighbor game is at least $\frac{1}{\rho^2}$.*

PROOF. Consider the scenario where $s_1 = s_2 = 0$, and $s_5 = s_6 = 1$. For $x < 1/2$, let $s_3 = x$ and $s_4 = 1 - x$. Define $\delta = |s_4 - s_3| = 1 - 2x$.

It is clear that $z_1 = z_2 = 0$ and $z_5 = z_6 = 1$ in any pure NE. If agents 3 and 4 point to each other, it is easy to check that $OPT \leq 2\rho\delta^2/(1 + \rho)$.

We now exhibit a NE where agent 3 points to 1 and 4 points to 5. If this happens, $z_3 = \rho x/(1 + \rho)$ and $z_4 = 1 - \rho x/(1 + \rho)$. Since this is an NE, we need $|s_3 - z_1| \geq |z_4 - s_3|$ which implies $\delta \geq \rho x/(1 + \rho)$. Choose x so that $\delta = \rho x/(1 + \rho)$ which makes the NE feasible. The cost of the NE is exactly $2x^2\rho/(1 + \rho)$, so that the PoA is at least $(x/\delta)^2 \geq 1/\rho^2$. \square

4.3 Generalized Asymmetric Games

The K -NN game is a special case of a more general *asymmetric coevolution* game, which we define next. Each agent i has an intrinsic opinion s_i . His strategy involves expressing an opinion z_i which could be different from s_i . The vector of expressed opinions \mathbf{z}_{-i} together with s_i determines the strength of agent i 's friendship with other agents. In particular, for $i \neq j$, define $d_{ij}^j = |z_j - s_i|$ as the distance of agent j 's expressed opinion from agent i 's true opinion. Let $q_{ij}(\mathbf{z}) = F_i(d_{ij}^j, d_{-i,-j}^i)$ where F_i is a continuous function that is decreasing in the first coordinate and increasing in the remaining coordinates. This captures the strength of i 's interaction with j (which is asymmetric). This strength increases as d_{ij} decreases, and increases as $d_{i,j'}$ decreases for $j' \neq j$. Furthermore, for any \vec{d} , the values $\{q_{ij}(\mathbf{z})\}$ for each i lie within some polyhedral constraints that define feasible friendship formation. As an example, if each agent i is required to have exactly K_i friends, which is captured by $\sum_j q_{ij}(\mathbf{z}) = K_i$.

Define the cost of agent i as

$$C_i(z_i, \mathbf{z}_{-i}) = \sum_{j \neq i} (z_i - z_j)^2 q_{ij}(\mathbf{z}) + \rho_i (z_i - s_i)^2.$$

This defines the *asymmetric coevolution* game.

THEOREM 4.6. *The asymmetric coevolution game admits to a pure strategy Nash equilibrium when the cost functions C_i are continuous.*

PROOF. Note that the quantity $q_{ij}(\mathbf{z})$ is independent of z_i and only depends on s_i and \mathbf{z}_{-i} . Therefore, the function $(z_i - z_j)^2 q_{ij}(\mathbf{z})$ is continuous in \mathbf{z} and convex in z_i . This implies $C_i(z_i, \mathbf{z}_{-i})$ is continuous in \mathbf{z} and convex in z_i . This game is therefore a *concave game* and admits to a Nash equilibrium [18]. \square

We note that for the K -NN game, the cost C_i is not continuous in \mathbf{z}_{-i} so that this game violates the premise of the above theorem and need not admit to a pure NE. An interesting open question is to decide the computational complexity of NE when C_i are continuous.

5. REFERENCES

- [1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. *Review of Economic Studies*, 78(4):1201–1236, 2011.
- [2] R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- [3] D. Bindel, J. M. Kleinberg, and S. Oren. How bad is forming your own opinion? In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 57 – 66, 2011.
- [4] B. Chazelle. Natural algorithms and influence systems. *Communications of the ACM*, 2012.
- [5] P. Dandekar, A. Goel, and D. Lee. Biased assimilation, homophily, and the dynamics of polarization. In *Workshop on Internet and Network Economics (WINE)*, 2012.
- [6] M. H. DeGroot. Reaching a consensus. *J. American Statistical Association*, 69:118–121, 1974.
- [7] R. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. . Mucha, F. Shi, D. Sivakoff, J. E. S. Socolar, and C. Varghese. Graph fission in an evolving voter model. *Proc. of National Academy of Sciences*, 109(10):3682–3687, 2012.
- [8] N. E. Friedkin and E. C. Johnsen. Social influence and opinions. *J. Mathematical Sociology*, 15(3-4):193–205, 1990.
- [9] B. Golub and M. O. Jackson. Naive learning in social networks: Convergence, influence, and the wisdom of crowds. *Humanities*, page Working Paper 64, 2007.
- [10] J. Hagenbach and F. Koessler. Strategic communication networks. *Review of Economic Studies*, 77(3):1072–1099, 2010.
- [11] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Artificial Societies and Social Simulations*, 5(3), 2002.
- [12] A. D. Henry, P. Pralat, and C. Q. Zhang. Emergence of segregation in evolving social networks. *Proc. of National Academy of Sciences*, 108:8605–8610, 2011.
- [13] P. Holme and M. E. J. Newman. Nonequilibrium phase transitions in the coevolution of networks and opinions. *Phys Rev E Stat Nonlin Soft Matter Phys*, 74:056108, 2006.
- [14] M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008.
- [15] T. Kuran and W. H. Sandholm. Cultural integration and its discontents. *Review of Economic Studies*, 75(1):201–228, 2008.
- [16] J. Lorenz. Heterogeneous bounds of confidence: Meet, discuss and find consensus! *Complexity*, 15(4):43–52, 2010.
- [17] A. Neyman. Correlated equilibrium and potential games. *Int. J. Game Theory*, 26(2):223–227, 1997.
- [18] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.
- [19] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, pages 513–522, 2009.
- [20] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *SODA*, pages 255–267, 2011.
- [21] T. C. Schelling. Models of segregation. *American Economic Review*, 59(2):488–493, 1969.
- [22] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal. Meet, discuss and segregate! *Complexity*, 7(3):55–63, 2002.