

Bounding the Inefficiency of Compromise

Ioannis Caragiannis
 University of Patras
 caragian@ceid.upatras.gr

Panagiotis Kanellopoulos
 CTI & University of Patras
 kanellop@ceid.upatras.gr

Alexandros A. Voudouris
 University of Patras
 voudouris@ceid.upatras.gr

Abstract

Social networks on the Internet have seen an enormous growth recently and play a crucial role in different aspects of today’s life. They have facilitated information dissemination in ways that have been beneficial for their users but they are often used strategically in order to spread information that only serves the objectives of particular users. These properties have inspired a revision of classical opinion formation models from sociology using game-theoretic notions and tools. We follow the same modeling approach, focusing on scenarios where the opinion expressed by each user is a compromise between her internal belief and the opinions of a small number of neighbors among her social acquaintances. We formulate simple games that capture this behavior and quantify the inefficiency of equilibria using the well-known notion of the price of anarchy. Our results indicate that compromise comes at a cost that strongly depends on the neighborhood size.

1 Introduction

Opinion formation has been the subject of much research in sociology, economics, physics, and epidemiology. Due to the widespread adoption of the Internet and the subsequent blossoming of social networks, it has recently attracted the interest of researchers in AI (e.g., see [Auletta *et al.*, 2016; Schwind *et al.*, 2015; Tsang and Larson, 2014]) and CS at large (e.g., see [Bindel *et al.*, 2015; Mossel and Tamuz, 2014; Olshevsky and Tsitsiklis, 2009]).

An influential model that captures the adoption of opinions in a social context has been proposed by Friedkin and Johnsen [1990]. According to this, each individual has an internal belief on an issue and publicly expresses an opinion; internal beliefs and public opinions are modeled as real numbers. In particular, the opinion an individual expresses follows by averaging between her internal belief and the opinions expressed by her social acquaintances. Recently, Bindel *et al.* [2015] show that this behavior can be explained through a game-theoretic lens: averaging between the internal belief of an individual and the opinions in her social circle is simply a *strategy* that minimizes an implicit cost for the individual.

Bindel *et al.* [2015] use a quadratic function to define this cost. This function is equal to the total squared distance of the opinion expressed by the individual from her belief and the opinions expressed in her social circle. In a sense, this behavior leads to opinions that follow the majority of her social acquaintances.

Bindel *et al.* [2015] consider a static snapshot of the social network. In contrast, Bhawalkar *et al.* [2013] implicitly assume that the opinion of an individual depends on a small number of acquaintances only, her *neighbors*. So, in their model, opinion formation co-evolves with the neighborhood for each individual: her neighborhood consists of those who have opinions similar to her belief. Then, the opinion expressed is assumed to minimize the same cost function used by Bindel *et al.* [2015], taking into account the neighborhood instead of the whole social circle.

We follow the co-evolutionary model of [Bhawalkar *et al.*, 2013], but we deviate from their cost definition and instead consider individuals that seek to *compromise* with their neighbors. Hence, we assume that each individual aims to minimize the *maximum* distance of her expressed opinion from her belief and each of her neighbors’ opinion. Like [Bhawalkar *et al.*, 2013], we assume that opinion formation co-evolves with the social network. The neighborhood of each individual consists of the k players with the closest opinions to her belief. Naturally, these modeling decisions lead to the definition of strategic games, which we call *k-compromising opinion formation* (or, simply, *k-COF*) games. Each individual is a (cost-minimizing) player with the opinion expressed as her strategy.

Technical contribution. We study questions related to the existence, computational complexity, and quality of equilibria in *k-COF* games. We show that there exist simple 1-COF games that do not admit pure Nash equilibria and, furthermore, that even in games where equilibria exist, their quality may be suboptimal, i.e., the *price of stability* (defined in [Anshelevich *et al.*, 2008]) is strictly greater than 1. We also show that there is a representation of each 1-COF game as a directed acyclic graph, in which every pure Nash equilibrium corresponds to a path between two designated nodes. Hence, the problems of computing the best or worst pure Nash equilibrium (or even of computing whether such an equilibrium exists) are equivalent to simple path computations that can be performed in polynomial time. These results appear in

Section 3. In Sections 4 and 5, we present upper and lower bounds on the *price of anarchy* (introduced by Koutsoupias and Papadimitriou [1999]) of k -COF games that suggest a linear dependence on k . Our upper bound on the price of anarchy exploits, in a non-trivial way, linear programming duality in order to lower-bound the optimal social cost. For the fundamental case of 1-COF games, we obtain a tight bound of 3.

Related work. DeGroot [1974] proposed a framework that models the opinion formation process, where each individual has a real number as opinion and updates it based on a weighted averaging procedure. Subsequently, Friedkin and Johnsen [1990] refined the model by assuming that each individual has a private belief and expresses a (possibly different) public opinion that depends on her belief and the opinions of people to whom she has social ties. More recently, Bindel *et al.* [2015] studied this model and proved that, again for the setting where beliefs and opinions are reals, the repeated averaging process leads to an opinion vector that can be thought of as the unique equilibrium in a corresponding opinion formation game.

Deviating from the assumption that opinions depend on the whole social circle, Bhawalkar *et al.* [2013] consider co-evolutionary opinion formation games, where as opinions evolve so does the neighborhood of each person. This model is conceptually similar to previous ones that have been studied by Hegselmann and Krause [2002] and Holme and Newman [2006]. Both Bindel *et al.* [2015] and Bhawalkar *et al.* [2013] show constant bounds on the price of anarchy of the games they study. In contrast, the modified cost function we use in order to model compromise yields considerably higher price of anarchy.

A series of recent papers from the EconCS community consider discrete models with binary opinions. Chierichetti *et al.* [2013] consider discrete preference games, where beliefs and opinions are binary and study questions related to the price of stability. For these games, Auletta *et al.* [2015] characterize the social networks where the belief of the minority can emerge as the opinion of the majority. Auletta *et al.* [2016] generalize discrete preference games so that players are not only interested in agreeing with their neighbors and more complex constraints can be used to represent the players' preferences. Bilò *et al.* [2016] extend co-evolutionary formation games to the discrete setting. Other models assume that opinion updates depend on the entire social circle of each individual, who consults a small random subset of social acquaintances; see the recent paper by Fotakis *et al.* [2016] and the survey of Mossel and Tamuz [2014].

In spite of the extensive related literature in many different disciplines, we believe that our model captures the tendency to compromise more accurately.

2 Preliminaries

A compromising opinion formation game defined by the k nearest neighbors (henceforth, called k -COF game) is played by a set of n players whose beliefs lie on the line of real numbers. Let $\mathbf{s} = (s_1, s_2, \dots, s_n) \in \mathbb{R}^n$ be the vector containing the players' beliefs such that $s_i \leq s_{i+1}$ for each $i \in [n - 1]$.

Let $\mathbf{z} = (z_1, z_2, \dots, z_n) \in \mathbb{R}^n$ be a vector containing the (deterministic or randomized) opinions expressed by the players; these opinions define a state of the game. We denote by \mathbf{z}_{-i} the opinion vector obtained by removing z_i from \mathbf{z} . In an attempt to simplify notation, we omit k from all relevant definitions.

Given vector \mathbf{z} (or a realization of it in case \mathbf{z} contains randomized opinions), we define the neighborhood $N_i(\mathbf{z}, \mathbf{s})$ of player i to be the set of k players whose opinions are the closest to the belief of player i breaking ties arbitrarily (but consistently). For each player i , we define $I_i(\mathbf{z}, \mathbf{s})$ as the shortest interval of the real line that includes the following points: the belief s_i , the opinion z_i , and the opinion z_j for each player $j \in N_i(\mathbf{z}, \mathbf{s})$. Furthermore, let $\ell_i(\mathbf{z}, \mathbf{s})$ and $r_i(\mathbf{z}, \mathbf{s})$ be the players with the leftmost and rightmost point in $I_i(\mathbf{z}, \mathbf{s})$, respectively. For example, $\ell_i(\mathbf{z}, \mathbf{s})$ can be equal to either player i or some player $j \in N_i(\mathbf{z}, \mathbf{s})$, depending on whether the leftmost point of $I_i(\mathbf{z}, \mathbf{s})$ is s_i , z_i , or z_j . In the following, we present the relevant definitions for the case of possibly randomized opinion vectors; clearly, these can be simplified whenever \mathbf{z} consists entirely of deterministic opinions.

Given a k -COF game with belief vector \mathbf{s} and the opinion vector \mathbf{z} , the cost of player i is defined as

$$\begin{aligned} \mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] &= \mathbb{E} \left[\max_{j \in N_i(\mathbf{z}, \mathbf{s})} \left\{ |z_i - s_i|, |z_j - z_i| \right\} \right] \\ &= \mathbb{E} \left[\max \left\{ |z_i - s_i|, |z_{r_i(\mathbf{z}, \mathbf{s})} - z_i|, |z_i - z_{\ell_i(\mathbf{z}, \mathbf{s})}| \right\} \right]. \end{aligned}$$

For the special case of 1-COF games, we denote by $\sigma_i(\mathbf{z}, \mathbf{s})$ (or $\sigma(i)$ when \mathbf{z} and \mathbf{s} are clear from context) the player (other than i) whose opinion is closest to the belief s_i of player i , i.e., $\sigma(i)$ is the only member of $N_i(\mathbf{z}, \mathbf{s})$. In this case, the cost of player i can be simplified as

$$\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] = \mathbb{E} \left[\max \left\{ |z_i - s_i|, |z_{\sigma_i(\mathbf{z}, \mathbf{s})} - z_i| \right\} \right].$$

We say that an opinion vector \mathbf{z} consisting entirely of deterministic opinions is a *pure Nash equilibrium* if no player i has any incentive to unilaterally deviate to a deterministic opinion z'_i in order to decrease her cost, i.e.,

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq \text{cost}_i((z'_i, \mathbf{z}_{-i}), \mathbf{s}),$$

where by (z'_i, \mathbf{z}_{-i}) we denote the opinion vector in which player i chooses the opinion z'_i and all other players choose the opinions they have according to vector \mathbf{z} . Similarly, a possibly randomized opinion vector \mathbf{z} is a *mixed Nash equilibrium* if for any player i and any deviating deterministic opinion z'_i we have

$$\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] \leq \mathbb{E}_{\mathbf{z}_{-i}}[\text{cost}_i((z'_i, \mathbf{z}_{-i}), \mathbf{s})].$$

The *social cost* of the opinion vector \mathbf{z} is defined as the total cost experienced by all players, i.e.,

$$\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})] = \sum_{i=1}^n \mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})].$$

Let \mathbf{z}^* be a deterministic opinion vector that minimizes the social cost; we will refer to it as the *optimal* opinion vector.

The *price of anarchy* (PoA) of a k -COF game over pure Nash equilibria is defined as the ratio between the social cost of its *worst* (in terms of the social cost) pure Nash equilibrium and the optimal social cost, i.e.,

$$\text{PoA}(\mathbf{s}) = \sup_{\mathbf{z} \in \text{PNE}} \frac{\text{SC}(\mathbf{z}, \mathbf{s})}{\text{SC}(\mathbf{z}^*, \mathbf{s})},$$

where PNE denotes the set of pure Nash equilibria. The *price of stability* (PoS) over pure Nash equilibria is defined as the ratio between the social cost of the *best* pure equilibrium (in terms of social cost) and the optimal social cost, i.e.,

$$\text{PoS}(\mathbf{s}) = \inf_{\mathbf{z} \in \text{PNE}} \frac{\text{SC}(\mathbf{z}, \mathbf{s})}{\text{SC}(\mathbf{z}^*, \mathbf{s})}.$$

Similarly, the price of anarchy and the price of stability over mixed Nash equilibria are defined as

$$\text{PoA}(\mathbf{s}) = \sup_{\mathbf{z} \in \text{MNE}} \frac{\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})]}{\text{SC}(\mathbf{z}^*, \mathbf{s})}, \quad \text{PoS}(\mathbf{s}) = \inf_{\mathbf{z} \in \text{MNE}} \frac{\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})]}{\text{SC}(\mathbf{z}^*, \mathbf{s})}$$

where MNE denotes the set of mixed Nash equilibria.

Then, the price of anarchy and the price of stability of k -COF games, for a fixed k , are defined as the supremum of $\text{PoA}(\mathbf{s})$ and $\text{PoS}(\mathbf{s})$ over all belief vectors \mathbf{s} .

We now state useful properties of pure Nash equilibria.

Lemma 1. *In any pure Nash equilibrium \mathbf{z} of a k -COF game with belief vector \mathbf{s} ,*

- z_i lies in the middle of interval $I_i(\mathbf{z}, \mathbf{s})$, for each player i ;
- $z_i \leq z_{i+1}$, for any $i \in [n-1]$;
- $N_i(\mathbf{z}, \mathbf{s}) = \{j, \dots, j+k\} \setminus \{i\}$ with $i-k \leq j \leq i$;
- $s_{\ell_i(\mathbf{z}, \mathbf{s})} \leq z_i \leq s_{r_i(\mathbf{z}, \mathbf{s})}$, for each player i .

The first property follows trivially by the definition of the cost function. Below, we give the proof of the second property of Lemma 1; the remaining proofs (as well as many proofs in the following) are omitted due to lack of space.

Proof of Lemma 1b. To simplify notation, we abbreviate $\ell_i(\mathbf{z}, \mathbf{s})$ and $r_i(\mathbf{z}, \mathbf{s})$ by $\ell(i)$ and $r(i)$, respectively. For the sake of contradiction, let us assume that $z_{i+1} < z_i$ for a pair of players i and $i+1$. Then, it cannot be the case that the leftmost endpoint of the interval $I_i(\mathbf{z}, \mathbf{s})$ of player i is at the left of (or coincides with) the leftmost endpoint of interval $I_{i+1}(\mathbf{z}, \mathbf{s})$ of player $i+1$ and the rightmost endpoint of $I_i(\mathbf{z}, \mathbf{s})$ is at the left of (or coincides with) the rightmost endpoint of $I_{i+1}(\mathbf{z}, \mathbf{s})$. In other words, it cannot be the case that $\min\{s_i, z_{\ell(i)}\} \leq \min\{s_{i+1}, z_{\ell(i+1)}\}$ and $\max\{s_i, z_{r(i)}\} \leq \max\{s_{i+1}, z_{r(i+1)}\}$ hold simultaneously. Since, by the first property, points z_i and z_{i+1} lie in the middle of the corresponding intervals, we would have $z_i \leq z_{i+1}$, contradicting our assumption.

So, at least one of the two inequalities between the interval endpoints above must not hold. In the following, we assume that $\min\{s_i, z_{\ell(i)}\} > \min\{s_{i+1}, z_{\ell(i+1)}\}$ (the case where $\max\{s_i, z_{r(i)}\} > \max\{s_{i+1}, z_{r(i+1)}\}$ is symmetric). Our assumption implies that $z_{\ell(i+1)} < s_i \leq s_{i+1}$, and, subsequently, that $z_{\ell(i+1)} < z_{\ell(i)}$. Hence, the opinion at the leftmost endpoint of interval $I_{i+1}(\mathbf{z}, \mathbf{s})$ cannot belong to player $i+1$, i.e., $\ell(i+1) \neq i+1$; see also Figure 1.




Figure 1: Example of the argument used in the proof of Lemma 1b. An arrow connects the belief of a player to her opinion.

Since $\ell(i+1)$ does not belong to $I_i(\mathbf{z}, \mathbf{s})$, there are at least k players different than $i+1$ and i that have opinions at distance strictly less than $s_i - z_{\ell(i+1)}$ from belief s_i . All these players are also at distance strictly less than $s_{i+1} - z_{\ell(i+1)}$ from belief s_{i+1} . This contradicts the fact that the opinion of player $\ell(i+1)$ is among the k closest opinions to s_{i+1} . \square

3 Existence, Complexity, and Quality of Equilibria


In this section, we focus entirely on 1-COF games. We first warm up with a negative statement: pure Nash equilibria may not exist. The construction in the proof of the next statement is inspired by [Bhawalkar *et al.*, 2013].

Theorem 2. *There exists a 1-COF game with no pure Nash equilibria.*

We continue to present a polynomial-time algorithm that determines whether a 1-COF game admits pure Nash equilibria, and, in case it does, allows us to compute the best and worst pure Nash equilibrium with respect to the social cost. We do so by establishing a correspondence between pure Nash equilibria and source-sink paths in a suitably defined directed acyclic graph.

Assume that we are given neighborhood information according to which each player i has either player $i-1$ or player $i+1$ as neighbor. From Lemma 1c, such a neighborhood structure is necessary in a pure Nash equilibrium. We claim that this information is enough in order to decide whether there is a consistent opinion vector that is a pure Nash equilibrium or not. All we have to do is to use Lemma 1a and obtain n equations that relate the opinion of each player to her belief and her neighbor's opinion. These equations have a unique solution which can then be verified whether it indeed satisfies the neighborhood conditions or not. So, the main idea of our algorithm is to cleverly search among all possible neighborhood structures that are not excluded by Lemma 1c for one that defines a pure Nash equilibrium.

For integers $1 \leq a \leq b < c \leq n$, let us define the *segment* $C(a, b, c)$ to be the set of players $\{a, a+1, \dots, c\}$ together with the following neighborhood information for them: $\sigma(p) = p+1$ for $p = a, \dots, b$ and $\sigma(p) = p-1$ for $p = b+1, \dots, c$. See Figure 2 for an example. It can be easily seen that the neighborhood information for all players at a pure Nash equilibrium can always be decomposed into disjoint segments. Importantly, given the neighborhood information in segment $C(a, b, c)$ and the beliefs of its players, the opinions they could have in any pure Nash equilibrium that contains this segment is uniquely defined using Lemma


 Figure 2: Graphical representation of a segment $C(1, 2, 4)$.

1a (as $z_i = (s_i + z_{\sigma(i)})/2$). This opinion vector is not necessarily consistent to the given neighborhood structure. We call segment $C(a, b, c)$ *legit* if $a \neq 2$, $c \neq n - 1$ (so that it can be part of a decomposition) and the uniquely defined opinions are consistent to the neighborhood information of it.

A decomposition of neighborhood information for all players will consist of consecutive segments $C(a_1, b_1, c_1)$, $C(a_2, b_2, c_2)$, ..., $C(a_t, b_t, c_t)$ so that $a_1 = 1$, $c_t = n$, $a_\ell = c_{\ell-1} + 1$ for $\ell = 2, \dots, t$. Such a decomposition will yield a pure Nash equilibrium if it consists of legit segments and, furthermore, the uniquely defined opinions of players in consecutive segments is consistent to the neighborhood information in both of them.

Now, consider the directed graph G that has two special nodes designated as the source and the sink, and a node for each legit segment $C(a, b, c)$. Note that G has $\mathcal{O}(n^3)$ nodes. The source node is connected to all segment nodes $C(1, b, c)$ while all segment nodes $C(a, b, n)$ are connected to the sink. An edge from segment node $C(a, b, c)$ to segment node $C(a', b', c')$ exists if $a' = c + 1$ and the uniquely defined opinions of players in the two segments are consistent to the neighborhood information in both of them. By the definition of segments and of its edges, G is acyclic.

Based on the discussion above, there is a bijection between pure Nash equilibria and source-sink paths in G . In addition, we can assign a weight to each node of G that is equal to the total cost of the players in the corresponding segment. Then, the total weight of a source-sink path is equal to the social cost of the corresponding pure Nash equilibrium.


Hence, standard algorithms for computing shortest or longest paths in directed acyclic graphs can be used not only to detect whether a pure Nash equilibrium exists, but also to compute the equilibrium of best or worst social cost.

Theorem 3. *Given a 1-COF game, deciding whether a pure Nash equilibrium exists can be done in polynomial time. Furthermore, computing a pure Nash equilibrium of highest or lowest social cost can be done in polynomial time as well.*

Example 1. Consider a 1-COF game with four players with belief vector $\mathbf{s} = (0, 3, 4, 7)$. According to the discussion above, there are 10 segments of the form $C(a, b, c)$ with $1 \leq a \leq b < c \leq 4$, but it can be shown that only 3 of them are legit; these are $C(1, 1, 2)$, $C(3, 3, 4)$, and $C(1, 2, 4)$. For example, segment $C(1, 1, 4)$ corresponds to the opinion vector $(1, 2, 3, 5)$ which is not consistent to the neighborhood information $\sigma(2) = 1$ in the segment. The resulting directed acyclic graph G , appears in Figure 3 and implies that there exist two pure Nash equilibria for this 1-COF game, namely the opinion vectors $(1, 2, 5, 6)$ and $(\frac{5}{3}, \frac{10}{3}, \frac{11}{3}, \frac{16}{3})$.

We conclude this section by stating that even the best equilibrium can be inefficient. Inefficiency of (worst) equilibria will be the subject of the upcoming two sections.

Theorem 4. *The price of stability of 1-COF games is at least 17/15.*


 Figure 3: The directed acyclic graph G for Example 1.

4 Upper Bounds on the Price of Anarchy

In the proof of the upper bound on the price of anarchy of k -COF games, we relate the social cost of any deterministic opinion vector, including the optimal one, to a quantity that depends only on the beliefs of the players and can be thought of as the cost of the truthful opinion vector (in which the opinion of every player is equal to her belief).

Consider an n -player k -COF game with belief vector $\mathbf{s} = (s_1, \dots, s_n)$. For player i , we denote by $\ell^*(i)$ and $r^*(i)$ the integers in $[n]$ that are such that $\ell^*(i) \leq i \leq r^*(i)$, $r^*(i) - \ell^*(i) = k$, and $|s_{r^*(i)} - s_{\ell^*(i)}|$ is minimized. The proof of the next lemma exploits linear programming and duality.

Lemma 5. *Consider a k -COF game and let $\mathbf{s} = (s_1, \dots, s_n)$ denote the belief vector and \mathbf{z} be any deterministic opinion vector. Then, $SC(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2(k+1)} \sum_{i=1}^n |s_{r^*(i)} - s_{\ell^*(i)}|$.*

Furthermore, we can show that for any pure Nash equilibrium \mathbf{z} of a k -COF game, we have that $SC(\mathbf{z}, \mathbf{s}) \leq 2 \sum_{i=1}^n |s_{r^*(i)} - s_{\ell^*(i)}|$. By combining this statement with Lemma 5, we obtain the following theorem.

Theorem 6. *The price of anarchy of k -COF games over pure Nash equilibria is at most $4(k+1)$.*

For the case of 1-COF games we can prove a stronger statement using a similar proof roadmap but simpler (and shorter) arguments. We denote by $\eta(i)$ the player (other than i) that minimizes the distance $|s_i - s_{\eta(i)}|$; note that $\eta(i) \in \{i-1, i+1\}$.

Lemma 7. *Consider a 1-COF game and let $\mathbf{s} = (s_1, \dots, s_n)$ denote the belief vector and \mathbf{z} be any deterministic opinion vector. Then, $SC(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|$.*

Proof. We call a player i with $z_i \notin [s_{i-1}, s_{i+1}]$ a *kangaroo* player and associate the quantity excess_i with her. If $z_i \in [s_j, s_{j+1}]$ for some $j > i$, we say that the players in the set $C_i = \{i+1, \dots, j\}$ are *covered* by player i and define $\text{excess}_i = z_i - s_j$. Otherwise, if $z_i \in [s_{j-1}, s_j]$ for some $j < i$, we say that the players in the set $C_i = \{j, \dots, i-1\}$ are covered by player i and define $\text{excess}_i = s_j - z_i$.

Let \mathcal{K} be the set of kangaroo players and \mathcal{C} the set of players that are covered by a kangaroo; these need not be disjoint. We now partition the players not in $\mathcal{K} \cup \mathcal{C}$ into the set L of *large* players i with $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3}(|s_i - s_{\eta(i)}|)$ and the set S that contains the remaining *small* players. See Figure 4.

We proceed to prove five useful claims.

Claim 8. *Let $i \in S$ such that $\sigma(i) \in \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)} \geq \frac{1}{3}|s_i - s_{\eta(i)}|$.*

Proof. We assume that $\sigma(i) > i$ (the other case is symmetric). If $z_{\sigma(i)} > s_{\sigma(i)}$, then $\text{cost}_i(\mathbf{z}, \mathbf{s}) = \max\{|s_i - z_i|, |z_i -$



Figure 4: Kangaroos, covered, large, and small players: $1 \in \mathcal{K}$, $2 \in \mathcal{K} \cap \mathcal{C}$, $3 \in \mathcal{C}$, $4 \in \mathcal{S}$, and $5 \in \mathcal{L}$.

$z_{\sigma(i)}\} \geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{2}(s_{\sigma(i)} - s_i) \geq \frac{1}{3}|s_i - s_{\eta(i)}|$, which contradicts the fact that i is a small player. Hence, $z_{\sigma(i)} \in [s_i, s_{\sigma(i)}]$, otherwise player i would be covered. Let j be the player with the leftmost belief that is covered by player $\sigma(i)$. Then, $\text{excess}_{\sigma(i)} = s_j - z_{\sigma(i)}$. We have $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)} = \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} + s_j - z_{\sigma(i)} \geq \frac{1}{2}(z_{\sigma(i)} - s_i) + \frac{1}{2}(s_j - z_{\sigma(i)}) \geq \frac{1}{3}(s_j - s_i) \geq \frac{1}{3}|s_i - s_{\eta(i)}|$. \square

Claim 9. Let $i \in \mathcal{S}$ such that $\sigma(i) \in \mathcal{L}$ or $\sigma(i) \in \mathcal{C} \setminus \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3}(|s_i - s_{\eta(i)}| + |s_{\sigma(i)} - s_{\eta(\sigma(i))}|)$.

Proof. We assume that $\sigma(i) > i$ (the other case is symmetric). If $z_{\sigma(i)} > s_{\sigma(i)}$, then $\text{cost}_i(\mathbf{z}, \mathbf{s}) = \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} \geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{2}(s_{\sigma(i)} - s_i) \geq \frac{1}{3}|s_i - s_{\eta(i)}|$, which contradicts the fact that i is a small player. Hence, $z_{\sigma(i)} \in [s_i, s_{\sigma(i)}]$, otherwise player i would be covered. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) = \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} + \max\{|s_{\sigma(i)} - z_{\sigma(i)}|, |z_{\sigma(i)} - z_{\sigma(\sigma(i))}|\} \geq z_{\sigma(i)} - z_i + s_{\sigma(i)} - z_{\sigma(i)} = s_{\sigma(i)} - z_i$. Since i is small, we have $z_i < s_i + \frac{1}{3}(s_{\sigma(i)} - s_i)$ and we get $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) \geq \frac{2}{3}(s_{\sigma(i)} - s_i) \geq \frac{1}{3}|s_i - s_{\eta(i)}| + \frac{1}{3}|s_{\sigma(i)} - s_{\eta(\sigma(i))}|$. \square

Claim 10. Let $i \in \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i \geq \frac{1}{3}(|s_i - s_{\eta(i)}| + \sum_{j \in \mathcal{C}_i} |s_j - s_{\eta(j)}|)$.

Proof. We assume that $z_i > s_i$ (the other case is symmetric). Let ℓ be the player with the rightmost belief that is covered by i . Then, $\text{excess}_i = z_i - s_\ell$. We have $\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i = \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} - (z_i - s_\ell) \geq s_\ell - s_i = \sum_{j=i}^{\ell-1} (s_{j+1} - s_j) \geq \frac{1}{3}(|s_i - s_{\eta(i)}| + \sum_{j \in \mathcal{C}_i} |s_j - s_{\eta(j)}|)$. \square

Let $N(S) = \{j \in [n] : \sigma(i) = j \text{ for } i \in S\}$.

Claim 11. $N(S)$ does not contain small players.

Proof. Assume otherwise that $\sigma(i) \in \mathcal{S}$ for some player $i \in S$. Without loss of generality $\sigma(i) > i$. If $z_{\sigma(i)} \geq s_{\sigma(i)}$, then $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}|z_{\sigma(i)} - s_i| \geq \frac{1}{2}|s_{\sigma(i)} - s_i| \geq \frac{1}{3}|s_i - s_{\eta(i)}|$ contradicting the fact that $i \in \mathcal{S}$. So, $z_{\sigma(i)} < s_{\sigma(i)}$. Also, $z_{\sigma(i)} \geq s_i$ (since neither i is covered nor $\sigma(i)$ is kangaroo). Since $\sigma(i)$ is small, $s_{\sigma(i)} - z_{\sigma(i)} < \frac{1}{3}|s_{\sigma(i)} - s_{\eta(\sigma(i))}| \leq \frac{1}{3}(s_{\sigma(i)} - s_i)$, i.e., $z_{\sigma(i)} > \frac{2}{3}s_{\sigma(i)} + \frac{1}{3}s_i$. Hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{3}(s_{\sigma(i)} - s_i)$, which contradicts $i \in \mathcal{S}$. \square

Claim 12. For every two players $i, i' \in S$, $\sigma(i) \neq \sigma(i')$.

Proof. Assume otherwise and let $\sigma(i) = \sigma(i') = j$ with $i < i'$. If $z_j \notin [s_i, s_{i'}]$, then the cost of either i or i' is at least $\frac{1}{2}(s_{i'} - s_i)$, contradicting the fact that both players are small. Hence, $z_j \in [s_i, s_{i'}]$. Notice that $s_j \in [s_i, s_{i'}]$ as well, otherwise either i or i' would be covered by j . Now the fact that i and i' are small implies that $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{i'}(\mathbf{z}, \mathbf{s}) <$

$\frac{1}{3}|s_i - s_{\eta(i)}| + \frac{1}{3}|s_{i'} - s_{\eta(i')}| \leq \frac{1}{3}(s_j - s_i) + \frac{1}{3}(s_{i'} - s_j) = \frac{1}{3}(s_{i'} - s_i)$. On the other hand, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{i'}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}(z_j - s_i) + \frac{1}{2}(s_{i'} - z_j) = \frac{1}{2}(s_{i'} - s_i)$, a contradiction. \square

Using T as an abbreviation of $L \cup (\mathcal{C} \setminus \mathcal{K})$, we have

$$\begin{aligned} \text{SC}(\mathbf{z}, \mathbf{s}) &= \sum_{i=1}^n \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &\geq \sum_{i \in \mathcal{S}: \sigma(i) \in \mathcal{K}} (\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)}) \\ &\quad + \sum_{i \in \mathcal{S}: \sigma(i) \in T} (\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s})) \\ &\quad + \sum_{i \in \mathcal{K}} (\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i) + \sum_{i \in L \setminus N(S)} \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &\geq \frac{1}{3} \sum_{i \in \mathcal{S}: \sigma(i) \in \mathcal{K}} |s_i - s_{\eta(i)}| \\ &\quad + \frac{1}{3} \sum_{i \in \mathcal{S}: \sigma(i) \in T} (|s_i - s_{\eta(i)}| + |s_{\sigma(i)} - s_{\eta(\sigma(i))}|) \\ &\quad + \frac{1}{3} \sum_{i \in \mathcal{K}} \left(|s_i - s_{\eta(i)}| + \sum_{j \in \mathcal{C}_i} |s_j - s_{\eta(j)}| \right) \\ &\quad + \frac{1}{3} \sum_{i \in L \setminus N(S)} |s_i - s_{\eta(i)}| \\ &\geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|, \end{aligned}$$

as desired. The first inequality follows by the player classification and due to Claims 11 and 12. The second one follows by Claims 8, 9, and 10, and by the definition of large players. The third one follows since the players enumerated in the first two sums at its left cover the whole set S (by Claim 11). \square

Theorem 13. The price of anarchy of 1-COF games over pure Nash equilibria is at most 3.

Proof. Let us consider a 1-COF game with n players and belief vector \mathbf{s} . Let \mathbf{z}^* be an optimal opinion vector and recall that $\eta(i)$ is the player that minimizes the distance $|s_i - s_{\eta(i)}|$. By Lemma 7, we have

$$\text{SC}(\mathbf{z}^*, \mathbf{s}) \geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|. \quad (1)$$

Now, consider any pure Nash equilibrium \mathbf{z} of the game. We will show that

$$\text{SC}(\mathbf{z}, \mathbf{s}) \leq \sum_{i=1}^n |s_i - s_{\eta(i)}|. \quad (2)$$

The theorem will then follow by (1) and (2).

In particular, we will show that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq |s_i - s_{\eta(i)}|$ for each player i . Let us assume that $\eta(i) = i - 1$; the case $\eta(i) = i + 1$ is symmetric. We distinguish between four cases.

- **Case I:** $\sigma(i) = i - 1$. By Lemma 1d, we have $s_{i-1} \leq z_i \leq s_i$. Then, clearly, $\text{cost}_i(\mathbf{z}, \mathbf{s}) = |s_i - z_i| \leq |s_i - s_{i-1}|$ as desired.
- **Case II:** $\sigma(i) = i + 1$ and $\sigma(i - 1) = i$. By Lemmas 1b and 1d, we have $s_{i-1} \leq z_{i-1} \leq s_i \leq z_i$. Since player i has player $i + 1$ as neighbor, we have $|z_{i+1} - s_i| \leq |s_i - z_{i-1}|$. Hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) = |z_i - s_i| \leq |z_{i+1} - s_i| \leq |s_i - z_{i-1}| \leq |s_i - s_{i-1}|$.
- **Case III:** $\sigma(i) = i + 1$, $\sigma(i - 1) = i - 2$, and $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq \text{cost}_{i-1}(\mathbf{z}, \mathbf{s})$. By the definition of the function σ and Lemma 1b, we have $z_{i-2} \leq z_{i-1} \leq s_{i-1} \leq s_i \leq z_i \leq z_{i+1}$. We have

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &\leq 2 \text{cost}_{i-1}(\mathbf{z}, \mathbf{s}) - \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &= |s_{i-1} - z_{i-2}| - |z_i - s_i| \\ &\leq |z_i - s_{i-1}| - |z_i - s_i| \\ &= |s_i - s_{i-1}|. \end{aligned}$$

The second inequality follows since player $i - 2$ (instead of i) is the neighbor of player $i - 1$.

- **Case IV:** $\sigma(i) = i + 1$, $\sigma(i - 1) = i - 2$, and $\text{cost}_i(\mathbf{z}, \mathbf{s}) > \text{cost}_{i-1}(\mathbf{z}, \mathbf{s})$.

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &< 2 \text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{cost}_{i-1}(\mathbf{z}, \mathbf{s}) \\ &= |z_{i+1} - s_i| - |s_{i-1} - z_{i-1}| \\ &\leq |s_i - z_{i-1}| - |s_{i-1} - z_{i-1}| \\ &= |s_i - s_{i-1}|. \end{aligned}$$

The second inequality follows since player $i + 1$ (instead of $i - 1$) is the neighbor of player i .

This completes the proof. \square

5 Lower Bounds on the Price of Anarchy

We now present lower bounds on the price of anarchy over pure and mixed Nash equilibria. We provide full proofs for the case of pure Nash equilibria in 1-COF games only.

Theorem 14. *The price of anarchy of 1-COF games over pure Nash equilibria is at least 3.*

Proof. Let $\lambda \in (0, 1)$ and consider a 1-COF game with 6 players and belief vector $\mathbf{s} = (-10 - \lambda, -10 - \lambda, -2 - \lambda, 2 + \lambda, 10 + \lambda, 10 + \lambda)$. This game is depicted in Figure 5. We can show that the opinion vector $\mathbf{z} = (-10 - \lambda, -10 - \lambda, -6 - \lambda, 6 + \lambda, 10 + \lambda, 10 + \lambda)$ is a pure Nash equilibrium with social cost $\text{SC}(\mathbf{z}, \mathbf{s}) = 8$. The first two players suffer zero cost as they follow each other and their opinions coincide with their beliefs; the same holds also for the last two players. For the third player, it is $\sigma(3) \in \{1, 2\}$ since $|z_1 - s_3| = |z_2 - s_3| = 8 < |z_4 - s_3| = 8 + 2\lambda$ and z_3 is in the middle of the interval $[-10 - \lambda, -2 - \lambda]$; hence, $\text{cost}_3(\mathbf{z}, \mathbf{s}) = 4$. Similarly, we have $\sigma(4) \in \{5, 6\}$, z_4 lies in the middle of the interval $[2 + \lambda, 10 + \lambda]$ and $\text{cost}_4(\mathbf{z}, \mathbf{s}) = 4$. Hence, \mathbf{z} is indeed a pure Nash equilibrium.

Now, consider the opinion vector

$$\bar{\mathbf{z}} = \left(-10 - \lambda, -10 - \lambda, \frac{-2 - \lambda}{3}, \frac{2 + \lambda}{3}, 10 + \lambda, 10 + \lambda \right)$$




Figure 5: The 1-COF game that is used to derive the lower bounds on the price of anarchy over pure and mixed Nash equilibria. Each point corresponds to a belief and $[x]$ denotes that x players have a particular belief.

which yields a social cost of $\text{SC}(\bar{\mathbf{z}}, \mathbf{s}) = \frac{8+4\lambda}{3}$; here, again, the first and last two players have zero cost, but players 3 and 4 now each have cost $\frac{4+2\lambda}{3}$ since they follow each other. The optimal social cost is at most $\text{SC}(\bar{\mathbf{z}})$ and, hence, the price of anarchy is at least $\frac{\text{SC}(\mathbf{z}, \mathbf{s})}{\text{SC}(\bar{\mathbf{z}}, \mathbf{s})} = \frac{3}{1+\lambda/2}$. The theorem follows by setting λ arbitrarily close to 0. \square

We now consider the case of mixed Nash equilibria and we remark that the next lower bound is greater than the upper bound of Theorem 13 on the price of anarchy over pure Nash equilibria. Hence, allowing randomized strategies may lead to more inefficient outcomes and, therefore, there is a clear separation between the classes of pure and mixed Nash equilibria. The proof relies again on the 1-COF game presented in Figure 5.

Theorem 15. *The price of anarchy of 1-COF games over mixed Nash equilibria is at least 6.*

We conclude this section with lower bounds on the price of anarchy of k -COF games, with $k \geq 2$.

Theorem 16. *The price of anarchy of k -COF games over pure Nash equilibria is at least $k + 1$ for $k \geq 3$, and at least $18/5$ for $k = 2$.*

Theorem 17. *The price of anarchy of k -COF games over mixed Nash equilibria is at least $k + 2$ for $k \geq 3$, and at least $24/5$ for $k = 2$.*

6 Conclusion and Open Problems

We have introduced the class of k -COF games. Our findings indicate that the quality of their equilibria grows linearly with the neighborhood size k . Still, there exists a gap between our lower and upper bounds for $k \geq 2$ and closing this gap is a challenging technical task. Furthermore, we know that mixed equilibria are strictly worse for 1-COF games but we have been unable to prove upper bounds on their price of anarchy. Is their price of anarchy still linear? Also, whether the price of stability depends on k or is at most a (small) constant is an interesting and wide open question. Another natural question is about the complexity of pure Nash equilibria in k -COF games for $k \geq 2$. We conjecture that there exists a polynomial time algorithm for computing them, but finding such an algorithm remains elusive.

Acknowledgements

This work was partially supported by Caratheodory research grant E.114 from the University of Patras and by a PhD scholarship from the Onassis Foundation. The first author would like to thank Pino Persiano for fruitful discussions at early stages of this work.

References

- [Anshelevich *et al.*, 2008] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [Auletta *et al.*, 2015] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. Minority becomes majority in social networks. In *Proceedings of the 11th International Conference on Web and Internet Economics (WINE)*, pages 74–88, 2015.
- [Auletta *et al.*, 2016] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. Generalized discrete preference games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2016.
- [Bhawalkar *et al.*, 2013] Kshipra Bhawalkar, Sreenivas Golapudi, and Kamesh Munagala. Coevolutionary opinion formation games. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 41–50, 2013.
- [Bilò *et al.*, 2016] Vittorio Bilò, Angelo Fanelli, and Luca Moscardelli. Opinion formation games with dynamic social influences. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, pages 444–458, 2016.
- [Bindel *et al.*, 2015] David Bindel, Jon M. Kleinberg, and Sigal Oren. How bad is forming your own opinion? *Games and Economic Behavior*, 92:248–265, 2015.
- [Chierichetti *et al.*, 2013] Flavio Chierichetti, Jon M. Kleinberg, and Sigal Oren. On discrete preferences and coordination. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 233–250, 2013.
- [DeGroot, 1974] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [Fotakis *et al.*, 2016] Dimitris Fotakis, Dimitris Palyvos-Giannas, and Stratis Skoulakis. Opinion dynamics with local interactions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 279–285, 2016.
- [Friedkin and Johnsen, 1990] Noah E. Friedkin and Eugene C. Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–205, 1990.
- [Hegselmann and Krause, 2002] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- [Holme and Newman, 2006] Petter Holme and M. E. J. Newman. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E*, 74, 2006.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
- [Mossel and Tamuz, 2014] Elchanan Mossel and Omer Tamuz. Opinion exchange dynamics. *arXiv: 1401.4770*, 2014.
- [Olshevsky and Tsitsiklis, 2009] Alex Olshevsky and John N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.
- [Schwind *et al.*, 2015] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1590–1596, 2015.
- [Tsang and Larson, 2014] Alan Tsang and Kate Larson. Opinion dynamics of skeptical agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 277–284, 2014.