General Opinion Formation Games with Social Group Membership

Vittorio Bilò[1] , **Diodato Ferraioli**[2] , **Cosimo Vinci**[2,3]
[1]University of Salento, Italy
[2]University of Salerno, Italy
[3]Gran Sasso Science Institute, Italy
vittorio.bilo@unisalento.it, {dferraioli, cvinci}@unisa.it

Abstract

Modeling how agents form their opinions is of paramount importance for designing marketing and electoral campaigns. In this work, we present a new framework for opinion formation which generalizes the well-known Friedkin-Johnsen model by incorporating three important features: (i) social group membership, that limits the amount of influence that people not belonging to the same group may lead on a given agent; (ii) both attraction among friends, and repulsion among enemies; (iii) different strengths of influence lead from different people on a given agent, even if the social relationships among them are the same.

We show that, despite its generality, our model always admits a pure Nash equilibrium which, under opportune mild conditions, is even unique. Next, we analyze the performances of these equilibria with respect to a social objective function defined as a convex combination, parametrized by a value $\lambda \in [0,1]$, of the costs yielded by the untruthfulness of the declared opinions and the total cost of social pressure. We prove bounds on both the price of anarchy and the price of stability which show that, for not-too-extreme values of λ, performance at equilibrium are very close to optimal ones. For instance, in several interesting scenarios, the prices of anarchy and stability are both equal to $\max\{2\lambda, 1-\lambda\}/\min\{2\lambda, 1-\lambda\}$ which never exceeds 2 for $\lambda \in [1/5, 1/2]$.

1 Introduction

In recent years, a lot of attention has been devoted to studying how people form their opinions, and how the social media affect the opinion formation process. Understanding these aspects is of fundamental importance for analysing and forecasting electoral flows and implement suitable electoral campaigns, or for marketing purposes.

Most of the approaches proposed in the literature usually assume that people try to "imitate" their "friends". This is, for example, the case of the celebrated DeGroot (DG) model [Degroot, 1974; Bindel *et al.*, 2015], where opinions are continuous and repeatedly updated to the average of the opinions expressed by one's friends. Among the most relevant generalizations of the DG model is the one of Friedkin-Johnsen (FJ) [Friedkin and Johnsen, 1990], in which people have an internal belief about the matter in object that limits in some way the influence of friends. Other approaches consider discrete opinion spaces [Chierichetti *et al.*, 2018; Ferraioli *et al.*, 2016], or limited/local interactions [Fotakis *et al.*, 2018; Fotakis *et al.*, 2016], or dynamic settings where social relationships and internal beliefs evolve over time [Hegselmann and Krause, 2002; Bhawalkar *et al.*, 2013; Bilò *et al.*, 2018a; Ferraioli and Ventre, 2017; Auletta *et al.*, 2019].

All these models, however, focus on imitative behaviour only. Indeed, there are many examples in which our opinion is not only influenced by imitation of our friends, but also by rejection of our "enemies". One example arises from youth subcultures, where peoples belonging to two different subcultures, even if a strict relation exists among them (e.g., they are relatives or they are in the same school), try to make opposite choices about style and interests, with the goal to distinguish each from the other. Another example comes from politics, where the position of a party about a topic sometimes arises more in opposition to adversaries rather than from principles and values. To the best of our knowledge, very few works considered this mixture of attraction and repulsion in opinion formation [Auletta *et al.*, 2016; Acar *et al.*, 2017] and, in any case, they limit the modelling of attraction/repulsion to a logic setting, which can only be applied to discrete opinions.

Both examples described above also highlight a fundamental feature of opinion formation that most of the discussed works neglect: membership in social groups. Indeed, followers of a subculture (e.g., hipsters) are used to limit their musical interests to the genre of reference of this subculture (e.g. indie), even if they are influenced by people listening to different music styles. Similarly, people belonging to a party usually support only opinions "allowed" by that party, despite the amount of social pressure they may face.

Yet another limitation of most of the considered models is that they assume a strength of attraction (or dis-attraction) that is the same for each pair of friends (enemies), possibly diversified only by a scaling factor measuring the weight of the social relationship. However, it may not be the case that hipster guys are attracted in the same way by emo peers and by geek peers, even if they all share the same social relationship. Similarly, the position of a right party on a given topic may be influenced in different ways by a center party or by an extreme-right party, even if the right party shares the same contacts with the other two (e.g., they are always allied at elections). Such degree of generality, although only restricted to attraction phenomenon, has been considered before only by [Bhawalkar *et al.*, 2013].

In this work, we tackle all the above limitations by proposing a new, general, model in which people choose their opinion by trying to simultaneously imitate their "friends" and distinguish themselves from their "enemies". We allow opinions to be chosen from a continuous set (differently from [Auletta *et al.*, 2016; Acar *et al.*, 2017]), and model social group membership by limiting the set of choices of each agent within the boundaries imposed by her social group. Finally, we also allow the strength of attraction and repulsion to be completely arbitrary and pair-specific, and not only influenced by the weights of the social relationships.

Specifically, we model this opinion formation framework as a cost minimization game with n agents, in which each agent belonging to a social group chooses an opinion whose distance from her private belief cannot exceed a certain threshold yielded by the boundaries of the group. In other words, while an agent is allowed to change her/his opinion, this opinion cannot lead this agent too far away from the cluster (social group) she/he feels to belong to.

As a consequence of her choice and of the choices of all the others, each agent i experiences a cost which depends on n functions: an increasing function g_i (*private cost function*), which measures the cost of agent i for disagreeing with her own belief, and $n-1$ functions $f_{i,j}$ for each $j \neq i$ (*public cost functions*), which measure the cost of the social pressure. In particular, $f_{i,j}$ is increasing (resp., decreasing) when agent j is a friend (resp., an enemy) of agent i. We stress that, despite of the huge mathematical challenges met in dealing with non-binary enemy relationships (one of the novelty of our model), most of our results only require all these functions to be continuous. Hence, our work provides a significant advancement along the direction of designing new models for opinion formation which may yield a good compromise between simplicity (needed for an analytical study) and expressive power.

Nevertheless, we also focus on special classes of games, that we name *well-ordered*, which turn out to enjoy interesting theoretical properties, while still spanning many realistic settings. Specifically, we consider opinion formation games that include the following additional properties: (i) the social groups do not intersect (and thus the opinions of the members of a group are always different from the opinions of the members of other groups), and (ii) all cost functions are strictly convex (i.e., the marginal increment of the cost strictly increases (resp., decreases) as the distance between opinions increases). The first property is realized when the social group membership is sufficiently strong to avoid any overlap of the opinions of agents belonging to different groups, despite they may influence each other. The second property is highly motivated in opinion dynamics, too. Indeed, convex cost functions model scenarios in which (a) the urgency of fixing the disagreement with close friends quickly grows as the disagreement becomes larger and larger, and similarly, (b) putting distance among enemies becomes more and more urgent when their opinions are close to each other. Furthermore, we point out that convexity is a common assumption in opinion formation games (see, e.g., [Bindel *et al.*, 2015; Bhawalkar *et al.*, 2013]), in which the influence functions are convex by hypothesis or coincide with some specific convex functions (e.g., quadratic or higher degree polynomials).

In light of the above considerations, our opinion formation framework and the special case of well-ordered games are able to include and generalize most of the previously defined models. Moreover, they can have multiple applications even in settings departing from opinion formation, such as facility location with heterogeneous preferences [Serafino and Ventre, 2016], content publishing [Bilò *et al.*, 2020] and isolation games [Bilò *et al.*, 2011; Zhao *et al.*, 2008].

Our contribution. We show that any game induced by our model admits at least a *pure Nash equilibrium* (i.e., a stable configuration in which each agent cannot reduce her cost via a unilateral change of opinion). We stress that this result does not require convexity or any other restrictive assumption to hold. In general, a game may admit different equilibria; however, we show that it is unique in well-ordered opinion formation games (that, differently from general games, must satisfy some convexity assumptions).

Next, we focus on the evaluation of the quality of equilibria through the concepts of *Price of Anarchy (PoA)* and *Price of Stability (PoS)*, by following the literature on the topic (see, e.g., [Bhawalkar *et al.*, 2013; Bilò *et al.*, 2018a; Bindel *et al.*, 2015; Chierichetti *et al.*, 2018; Degroot, 1974; Epitropou *et al.*, 2019; Ferraioli *et al.*, 2016; Fotakis *et al.*, 2016]. Indeed, PoA and PoS are used to better understand the social degradation caused by opinion formation phenomena that often appear in several real-life scenarios (e.g., political polls, trends formation, etc...). Moreover, PoA and PoS results play a practical role in establishing when the intervention of social planner is necessary, and when there is no need of altering the evolution of the system: whenever PoA/PoS are high, intervention of social planner may be welcome.

In this work, we focus on different ways to evaluating the quality of an equilibrium. A first approach uses the *utilitarian social cost*, defined as the sum of the agents' costs. This direction has been taken, e.g., in [Bindel *et al.*, 2015; Chierichetti *et al.*, 2018; Bhawalkar *et al.*, 2013; Bilò *et al.*, 2018a]. A second approach emphasizes the *truthfulness* of the declared opinions, by bounding how much the social pressure deviates the agents' opinions from their private beliefs. This metric has been considered in [Auletta *et al.*, 2015; Auletta *et al.*, 2017; Auletta *et al.*, 2018]. A third approach, finally, measures the distance from a *consensus* [Hegselmann and Krause, 2002; Auletta *et al.*, 2019].

We believe that all these approaches are useful and meaningful. Not only, but it is often useful and meaningful to have, for example, equilibria that are close to be truthful (or close to be a consensus) and, at the same time, represent a good compromise for the society as a whole. For this reason, we propose to measure the performance of an equilibrium by means of the λ-*social influence cost*, obtained by summing the cost of untruthfulness scaled by λ and the cost of social pressure (i.e., distance from a consensus) scaled by $1 - \lambda$, for any $\lambda \in [0, 1]$. Observe that, by setting $\lambda = 1/2$, $\lambda = 1$ and $\lambda = 0$, respectively, we re-obtain the above three metrics.

Our results highlight how PoA and PoS with respect to λ-social influence cost vary as the parameters of the system change: this will provide practically useful suggestions about the direction in which possible interventions of a social planner should occur. For example, our results suggest that, in order to guarantee that opinion formation converges to states with good social performances, one should try to avoid enemy relation or one should try to assure that social groups are "closed" as described in the definition of well-ordered games. Hence, the social planner may be interested in designing campaigns to enforce these properties.

Specifically we prove that for extreme values of λ (i.e., $\lambda = 0$ or $\lambda = 1$), the PoA and the PoS can grow arbitrarily large, as it may be impossible to reach an equilibrium that is a consensus or a truthful profile when considering agents with general cost functions and possessing both attraction and dis-attraction attitudes. Nevertheless, we surprisingly show that the PoA and the PoS are usually not very large when λ is sufficiently far from the extremes. Specifically, we prove that the PoS is always (i.e., we do not require convexity or other assumptions) bounded by $\frac{\max\{2\lambda, 1-\lambda\}}{\min\{2\lambda, 1-\lambda\}} = O\left(\max\left\{\frac{1}{\lambda}, \frac{1}{1-\lambda}\right\}\right)$. The same bound holds even for the PoA in well-ordered opinion formation games, while in general the PoA can be unbounded. Moreover, when the cost functions obey some additional mild assumptions, better bounds on the PoA are possible. The technique used to prove this result may be of independent interest: a generalization of the primal-dual technique introduced in [Bilò, 2018], and applied for the first time in this setting. We additionally show that these bounds are often tight. In the full version, we also provide applications of our general results to specific class of well-studied games, by proving tight numerical bounds.

2 Model and Definitions

Mathematical Notation. Given an integer n, let $[n]$ denote the set $\{1, \ldots, n\}$. Given n convex functions $F_i : A \subseteq \mathbb{R}^n \to \mathbb{R}$, a vector $\boldsymbol{v} = (v_1, \ldots, v_n) \in \mathbb{R}^n$, and $\boldsymbol{x} := (x_1, \ldots, x_n) \in A$, let $\langle (F_i(\boldsymbol{x}))_i, \boldsymbol{v} \rangle := \sum_{i \in [n]} \frac{\partial F_i(\boldsymbol{x})}{\partial x_i^{sgn(v_i)}} \cdot v_i$, where $\frac{\partial F_i(\boldsymbol{x})}{\partial x_i^{sgn(v_i)}}$ denotes the right-derivative (resp. left-derivative) of F_i in $\boldsymbol{x}$ w.r.t. to variable x_i, if $v_i \geq 0$ (resp. $v_i < 0$)[1]. If the apex $sgn(v_i)$ is missing, we implicitly refer to the right-derivative. We conventionally assume that an indeterminate form of type $c/0$ is equal to ∞ if $c > 0$. Other indeterminate forms will be specified when they occur.

Generalized Opinion Formation Games. Let $N := [n]$ be a set of n *agents*. Each agent $i \in [n]$ has a *private belief* $s_i \in [0, 1]$. In order to model the membership of agents to a social group, and the influence that this has on her opinion, we assume that each agent i has a *maximum left (resp. right) deviation value* $d_i^- \in [0, s_i]$ (resp. $d_i^+ \in [0, 1 - s_i]$). These values, determined with respect to the social group at which one belongs, limit the extent at which the opinion of an agent may change (and thus the extent at which this opinion may differ from the opinion of other agents in the same group).

Let $\boldsymbol{s} = (s_1, \ldots, s_n)$ be the *private belief vector* and let $\boldsymbol{d} = ((d_1^-, d_1^+), \ldots, (d_n^-, d_n^+))$ be the *maximum deviation vector*. We assume w.l.o.g. that $0 \leq s_1 \leq \ldots \leq s_n \leq 1$. Each agent $i \in [n]$ declares a *public opinion* $x_i \in [0, 1]$ (equivalently denoted as the *strategy* of agent i) such that $-d_i^- \leq x_i - s_i \leq d_i^+$. We let $\boldsymbol{x} = (x_1, \ldots, x_n)$ denote the resulting *opinion profile*. Ideally, $s_i - d_i^-$ and $s_i + d_i^+$ correspond to the boundaries of the social group to which agent i belongs (thus our constraint on the public opinion essentially states that i always remains within her own social group).

Each agent $i \in [n]$ in an opinion profile $\boldsymbol{x}$ incurs in a *public influence cost* $c_{pu,i}(\boldsymbol{x})$ defined as $c_{pu,i}(\boldsymbol{x}) := \sum_{j \in [n]: j \neq i} f_{i,j}(|x_i - x_j|)$, where, for any pair (i, j) such that $i \neq j$, $f_{i,j} : (0, 1] \to \mathbb{R}_{\geq 0}$ is called the (i, j)-*public influence function* and satisfies the following properties: (i) $f_{i,j}(x) = f_{j,i}(x)$ for any $x \in (0, 1]$; (ii) $f_{i,j}$ is continuous in $(0, 1]$; (iii) $\exists \lim_{x \to 0^+} f_{i,j}(x) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. Observe that we do not have any constraints on the slope of $f_{i,j}$. This allows us to model both attraction among friends (when $f_{i,j}$ is increasing), and repulsion among enemies (when $f_{i,j}$ is decreasing). Also, by choosing different functions for each pair of agents, we can represent different strengths of attraction and repulsion.

Moreover, each agent $i \in [n]$ in an opinion profile $\boldsymbol{x}$ incurs also in a *private influence cost* $c_{pr,i}(\boldsymbol{x})$ defined as $c_{pr,i}(\boldsymbol{x}) := g_i(|x_i - s_i|)$, where $g_i : [0, 1] \to \mathbb{R}_{\geq 0}$ is a continuous and non-decreasing function called the *i-private influence function*. The *influence cost* of agent $i \in [n]$ is defined as $c_i(\boldsymbol{x}) = c_{pu,i}(\boldsymbol{x}) + c_{pr,i}(\boldsymbol{x})$, i.e., the influence cost of each agent is given by the sum of her public and private influence costs. We assume that, for any agent $i \in [n]$, there exists at least a non-null public influence function $f_{i,j}$ for some $j \neq i$. The tuple $\mathcal{O} = (N, \boldsymbol{s}, \boldsymbol{d}, (f_{i,j})_{i \neq j}, (g_i)_i)$ is called *generalized opinion formation game* (GOF game).

Classes of GOF games. Given a GOF game $\mathcal{O}$, let $\mathcal{F}(\mathcal{O})$ and $\mathcal{G}(\mathcal{O})$ denote the set of non-null public and private influence functions of $\mathcal{O}$, respectively. A GOF game is *convex* if all the functions in $\mathcal{F}(\mathcal{O})$ and $\mathcal{G}(\mathcal{O})$ are convex, thus implying that the marginal increment of the cost increases (resp., decreases) as the distance between opinions increases.

A GOF game $\mathcal{O}$ is *unconstrained* if $d_i^- = d_i^+ = 1$ for any $i \in [n]$, i.e., if social group membership is not considered. $\mathcal{O}$ is an *isolation* game if all the functions in $\mathcal{F}(\mathcal{O})$ are non-increasing, i.e., every agent wants to be as far as possible from other agents. $\mathcal{O}$ is an *aggregation* game if all the functions in $\mathcal{G}(\mathcal{O})$ are non-decreasing, i.e., every agent wants to imitate her friends. The class of unconstrained aggregation games includes the standard opinion formation games introduced in [Friedkin and Johnsen, 1990; Bindel *et al.*, 2015] and their generalization considered in [Bhawalkar *et al.*, 2013].

A GOF game $\mathcal{O}$ is *well-ordered* if it is convex, and we can

[1]The right-derivative (resp. left-derivative) of a function $F : \mathbb{R}^n \to \mathbb{R}$ w.r.t. variable x_i calculated in $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$ is defined as the limit for $t \to 0^+$ (resp. $t \to 0^-$) of $(f(x_1, \ldots, x_i + t, \ldots, x_n) - f(x_1, \ldots, x_i, \ldots, x_n))/t$. Observe that left and right derivatives of a convex function always exist.

organize the agents in clusters $S_1, S_2, \ldots, S_k$ such that:

1. each cluster is a non-empty set of consecutive agents;
2. for any cluster S_r and $i, j \in S_r$, we have that each public influence function $f_{i,j}$ is non-decreasing, i.e., the subgame restricted to each cluster is an aggregation game;
3. for any $r \in [k-1]$, and any $i \in S_r$ and $j \in S_{r+1}$, we have that $d_i^+ + d_j^- \leq s_j - s_i$, i.e., for any opinion profile $\boldsymbol{x}$, and for any $i, j \in [n]$ and $r \in [k-1]$ with $i \in S_r$ and $j \in S_{r+1}$, we have that $x_r \leq x_{r+1}$.

Roughly speaking, in well-ordered games, all groups of agents are organized in disjoint intervals on the line, in such a way that one agent belonging to a group cannot express an opinion outside the corresponding interval. Despite this geometric structure of the agents' opinions is undoubtedly an extreme choice, it provides a realistic model for many settings. Indeed, it is often the case that changes in the structure of social groups (or clusters) do not occur among existing groups, but only as a side-effect of the birth of new groups, that may be endogenously provoked by alliances between extremists of existing groups (e.g., as for parties), or exogenously by the creation of a new product (e.g., in youth subcultures).

Observe that classical opinion formation games [Friedkin and Johnsen, 1990; Bindel *et al.*, 2015] are well-ordered, since they can be represented with a unique cluster containing all agents. Observe also that for isolation games to be well-ordered, we need each cluster to be a singleton.

Finally, a well-ordered GOF game is *regular* if all functions in $\mathcal{F}(\mathcal{O})$ and $\mathcal{G}(\mathcal{O})$ are continuously differentiable (i.e., the left and right derivatives are equal), and the derivative in $x = 0$ is null for each function $g \in \mathcal{G}(\mathcal{O})$ and function $f_{i,j} \in \mathcal{F}(\mathcal{O})$ with i, j belonging to the same cluster. Observe that the differentiability of the cost functions models the absence of "jumps" in the individual costs while the agents continuously change their public opinions, and it is a standard assumption in several opinion formation games (see, for instance, [Bindel *et al.*, 2015; Bhawalkar *et al.*, 2013]).

Pure Nash Equilibria and λ-Social Influence Cost. Given an opinion profile $\boldsymbol{x}$ and $y_i \in [0,1]$, let $(\boldsymbol{x}_{-i}, y)$ denote the opinion profile in which strategy x_i is replaced with y_i. An opinion profile $\boldsymbol{x}$ is a *(pure Nash) equilibrium* if, for any $i \in [n]$, we have that $c_i(\boldsymbol{x}) \leq c_i(\boldsymbol{x}_{-i}, y_i)$ for any (feasible) strategy y_i of agent i, i.e., no agent can reduce her influence cost via a unilateral change of strategy. Let $\mathsf{E}(\mathcal{O})$ denote the set of equilibria of game $\mathcal{O}$ and let $\mathsf{SP}(\mathcal{O})$ denote the set of opinion profiles of $\mathcal{O}$. We exclude from $\mathsf{SP}(\mathcal{O})$ and $\mathsf{E}(\mathcal{O})$ all the opinion profiles $\boldsymbol{x}$ such that $x_i = x_j$ and $f_{i,j}(0) = \infty$ (i.e., the influence cost of some agent in $\boldsymbol{x}$ is ∞).

Given $\lambda \in (0,1)$, the λ-*social influence cost* of a given opinion profile $\boldsymbol{x}$ is defined as

$$\mathsf{SUM}_\lambda(\boldsymbol{x}) := \sum_i (\lambda \cdot c_{pu,i}(\boldsymbol{x}) + (1-\lambda) \cdot c_{pr,i}(\boldsymbol{x}))$$
$$= 2\lambda \sum_{i>j} f_{i,j}(|x_i - x_j|) + (1-\lambda) \sum_i g_i(|x_i - s_i|),$$

i.e., it is a convex combination under parameter λ of the sum of all the public influence costs and the sum of all the private influence costs. Let $OPT_\lambda(\mathcal{O}) := \inf_{\boldsymbol{x} \in \mathsf{SP}(\mathcal{O})} \mathsf{SUM}_\lambda(\boldsymbol{x})$.[2]

Remark 1. The utilitarian social cost considered in many works about opinion formation games [Bindel *et al.*, 2015; Chierichetti *et al.*, 2018; Bhawalkar *et al.*, 2013; Bilò *et al.*, 2018a; Epitropou *et al.*, 2019] is equivalent to the λ-social influence cost with $\lambda = 1/2$. Opinion thruthfulness, i.e. how far the opinions at equilibrium are from the private beliefs [Auletta *et al.*, 2015; Auletta *et al.*, 2017; Auletta *et al.*, 2018], can be measured simply by the λ-social influence cost with λ close to 1. Similarly, understanding whether an equilibrium gets close to a consensus or not [Hegselmann and Krause, 2002; Auletta *et al.*, 2019], can be done by the λ-social influence cost with λ close to 0. Thus, the definition of the λ-social influence cost encompasses all these three commonly adopted measures of performance of equilibria, and allows also combination of them.

λ-Price of Anarchy and λ-Price of Stability. To evaluate the performance of equilibria with respect to the λ-social influence, we define the following concepts: the λ-*price of anarchy* of game $\mathcal{O}$, defined as $\mathsf{PoA}_\lambda(\mathcal{O}) := \sup_{\boldsymbol{x} \in \mathsf{E}(\mathcal{O})} \frac{\mathsf{SUM}_\lambda(\boldsymbol{x})}{OPT_\lambda(\mathcal{O})}$, which is the worst-case ratio between the performances of an equilibrium of $\mathcal{O}$ and the optimal λ-social influence cost of $\mathcal{O}$, and the λ-*price of stability* of game $\mathcal{O}$, defined as $\mathsf{PoS}_\lambda(\mathcal{O}) := \inf_{\boldsymbol{x} \in \mathsf{E}(\mathcal{O})} \frac{\mathsf{SUM}_\lambda(\boldsymbol{x})}{OPT_\lambda(\mathcal{O})}$, which is the best-case ratio between the performances of an equilibrium of $\mathcal{O}$ and the optimal λ-social influence cost of $\mathcal{O}$.

3 Equilibrium Existence

We first prove that an equilibrium always exists. Moreover, for strictly convex well-ordered games, it turns out that the equilibrium is unique. Due to space limitation, for most of the claims in this and in the next sections, proofs are omitted or only sketched. We refer the interested reader to the full version for more details.

3.1 General Case

In order to prove that any GOF game possesses an equilibrium, we use a potential function argument. Given a GOF game $\mathcal{O}$, a function $\Phi : \mathsf{SP}(\mathcal{O}) \to \mathbb{R}_{\geq 0}$ is a *potential function* of $\mathcal{O}$ if $\Phi(\boldsymbol{x}) - \Phi(\boldsymbol{x}_{-i}, y_i) = c_i(\boldsymbol{x}) - c_i(\boldsymbol{x}_{-i}, y_i)$ for any opinion profile $\boldsymbol{x}$, any agent $i \in [n]$, and any strategy y_i of agent i. Let Φ be the function such that $\Phi(\boldsymbol{x}) := \sum_{i>j} f_{i,j}(|x_i - x_j|) + \sum_i g_i(|x_i - s_i|)$, for any opinion profile $\boldsymbol{x}$. It is not hard to check that following lemmas hold.

Lemma 1. *Given a GOF game $\mathcal{O}$, Φ is a potential function of $\mathcal{O}$.*

Lemma 2. *Φ admits a global minimum point.*

As shown in [Monderer and Shapley, 1996], any global minimum of a potential function is a pure Nash equilibrium. Thus, with the help of Lemma 1 and Lemma 2, we can prove the following theorem.

[2]As the considered game is not discrete, showing that there exists an opinion profile minimizing a certain function is not immediate. Anyway, by using the same arguments as in Lemma 2 below, one can easily show that there exists an opinion profile $\boldsymbol{y}$ that minimizes the λ-social influence cost (i.e., $\mathsf{SUM}_\lambda(\boldsymbol{y}) = OPT_\lambda(\mathcal{O})$).

Theorem 1. *Any GOF game $\mathcal{O}$ admits at least a pure Nash equilibrium. In particular, all global minimum points of Φ are equilibria*[3].

3.2 Well-ordered Games

In the case of well-ordered GOF games, we have a better characterization of the set of equilibria.

We first give some preliminary lemmas.

Lemma 3. *Let $\mathcal{O}$ be a well-ordered game. Then the potential function Φ is a convex function defined on a convex set.*

Lemma 4. *Let $\mathcal{O}$ be a convex GOF game, and let $\boldsymbol{x}$ be an equilibrium of $\mathcal{O}$. Then $\langle (c_i(\boldsymbol{x}))_i, \boldsymbol{y}-\boldsymbol{x}\rangle \geq 0$ for any opinion profile $\boldsymbol{y}$ of $\mathcal{O}$.*

We are now ready to prove the equilibria characterization for well-ordered games.

Theorem 2. *Let $\mathcal{O}$ be a well-ordered GOF game. Then: (i) the set of equilibria of $\mathcal{O}$ coincides with the set of global minimum points of Φ; (ii) the set of equilibria is a convex set; (iii) if the non-null public influence functions are strictly convex, then there exists a unique equilibrium.*

Sketch of the proof. We show claim (i) only. By Theorem 1, we trivially have that any global minimum point of Φ is an equilibrium, thus it suffices to show that any equilibrium is a global minimum point of Φ. Let $\boldsymbol{x}$ be an equilibrium. By Lemma 4, and by using the definition of potential function, we have that $0 \leq \langle (c_i(\boldsymbol{x}))_i, \boldsymbol{y}-\boldsymbol{x}\rangle = \sum_{i\in[n]} \frac{\partial \Phi(\boldsymbol{x})}{\partial x_i^{sgn(y_i-x_i)}} \cdot (y_i - x_i)$, for any opinion profile $\boldsymbol{y}$ of $\mathcal{O}$. Thus, $\boldsymbol{x}$ is a local minimum point of Φ. As Φ is a convex function on a convex set by Lemma 3, then $\boldsymbol{x}$ is a global minimum point of Φ. □

4 The Efficiency of GOF Games

We have that the λ-price of anarchy can be unbounded, even for unconstrained and convex isolation games with two agents and linear functions.

Theorem 3. *There is an unconstrained convex isolation GOF game with two agents s.t. $\mathsf{PoA}_\lambda(\mathcal{O}) = \infty$ for any $\lambda \in (0,1)$.*

Proof. Consider an unconstrained GOF game $\mathcal{O}$ having two agents $1, 2$ with $s_1 = 0$, $s_2 = 1$, private and public influence functions such that $g_1(x) = g_2(x) = x/2$, and $f_{1,2}(x) = 1 - x$ for any $x \in [0,1]$. We observe that $\mathcal{O}$ is a convex isolation game. Let $\boldsymbol{x} = (1,0)$. One can easily show that $\boldsymbol{x}$ is an equilibrium of $\mathcal{O}$. Let $\boldsymbol{y} = (0,1)$. We have that $\mathsf{SUM}_\lambda(\boldsymbol{x}) = 1-\lambda$ and $\mathsf{SUM}_\lambda(\boldsymbol{y}) = 0$. Thus $\mathsf{PoA}_\lambda(\mathcal{O}) \geq \frac{\mathsf{SUM}_\lambda(\boldsymbol{x})}{\mathsf{SUM}_\lambda(\boldsymbol{y})} = \frac{1-\lambda}{0} = \infty$[4]. □

[3] By using similar arguments as in [Monderer and Shapley, 1996], one can also show that any asynchronous best-response dynamics (in which at most one agent at a time unilaterally deviates in favor of her best strategy) converges eventually to a pure Nash equilibrium.

[4] We observe that $\boldsymbol{x}$ and $\boldsymbol{y}$ are the unique equilibria of $\mathcal{O}$, i.e., there are multiple equilibria that constitute a non-convex set of the euclidean space. Thus, Theorem 2 does not hold if the considered game is not well-ordered, even for convex isolation games.

Differently from the λ-price of anarchy, for the λ-price of stability we get a tight bound which is parametrized by $\lambda \in (0,1)$ and is always finite.

Theorem 4. *Given a GOF game $\mathcal{O}$, we have that $\mathsf{PoS}_\lambda(\mathcal{O}) \leq \frac{\max\{2\lambda,1-\lambda\}}{\min\{2\lambda,1-\lambda\}}$.*

Sketch of the proof. Let $\boldsymbol{x}$ be a global minimum point of potential function Φ and let $\boldsymbol{y}$ be an opinion profile of $\mathcal{O}$ minimizing the λ-social influence cost. By Theorem 1, $\boldsymbol{x}$ is an equilibrium, and one can easily show that $\mathsf{SUM}_\lambda(\boldsymbol{x}) \leq \max\{2\lambda, 1-\lambda\} \cdot \Phi(\boldsymbol{x}) \leq \max\{2\lambda, 1-\lambda\} \cdot \Phi(\boldsymbol{y}) \leq \frac{\max\{2\lambda,1-\lambda\}}{\min\{2\lambda,1-\lambda\}} \cdot \mathsf{SUM}_\lambda(\boldsymbol{y})$. □

We provide matching lower bounds holding even for unconstrained aggregation (or isolation) games with two agents and linear functions.

Theorem 5. *For any $\epsilon > 0$, there exists an unconstrained and convex aggregation (isolation, resp.) game $\mathcal{O}$ with two agents and linear public and private influence functions such that $\mathsf{PoS}_\lambda(\mathcal{O}) \geq \frac{\max\{2\lambda,1-\lambda\}}{\min\{2\lambda,1-\lambda\}} - \epsilon$.*

5 Better Bounds for Well-ordered Games

In the following theorem, we show that the upper bound on the λ-price of stability established in Theorem 4 extends to the λ-price of anarchy if the considered game is well-ordered.

Theorem 6. *Given a well-ordered GOF game $\mathcal{O}$, we have that $\mathsf{PoA}_\lambda(\mathcal{O}) \leq \frac{\max\{2\lambda,1-\lambda\}}{\min\{2\lambda,1-\lambda\}}$.*

Sketch of the proof. Let $\boldsymbol{x}$ be an equilibrium maximizing the λ-social influence cost, and let $\boldsymbol{y}$ be an opinion profile minimizing the λ-social influence cost. By Theorem 2, we have that $\boldsymbol{x}$ minimizes the potential function Φ. Thus, the same inequalities as in Theorem 4 hold, and the claim follows. □

Remark 2. As suggested above, convex unconstrained aggregation games are a subclass of well-ordered games, and thus above upper bound holds even for this subclass. This immediately implies that the bound is tight by Theorem 5.

Remark 3. The upper bound provided in Theorem 6 implies that the λ-price of anarchy of well-ordered GOF games is equal to 1 (i.e., optimal) if $\lambda = 1/3$, and is at most 2 for any $\lambda \in [1/5, 1/2]$. Thus, by carefully choosing how to balance the measures of social cost, untruthfulness, and lack of consensus, we achieve that the performances of the game are always good, despite agents behave in a completely uncoordinated and selfish way. Observe that, by our bound also emerges that the performances of equilibria can be very bad whenever λ is close to 1 (resp. close to 0): this is exactly what we expect, since instances converging at the equilibrium to a consensus or to the truthful profile are extremely rare.

If the considered well-ordered game is also regular, we can obtain a generally better upper bound on the λ-price of anarchy, that depends also on the specific public and private influence functions, and not only on λ. To prove these bounds, we generalize the primal-dual method introduced in [Bilò, 2018] by incorporating some topological properties of regular well-ordered games. This approach, which is of independent interest and exports for the first time the primal-dual method outside the realm of congestion games, shares some similarities with the notion of local smoothness [Roughgarden and Schoppmann, 2015; Bhawalkar *et al.*, 2013].

Given $\theta, q, r, t \geq 0$ and a real function h, let $\eta_q(\theta, h, r, t) = \frac{q\cdot h(r)+\theta(t-r)\cdot\frac{\partial h(r)}{\partial r}}{q\cdot h(t)}$, with the convention that $c/0 := \infty$ if $c > 0$ and $c/0 := 1$ if $c \leq 0$.

Theorem 7. *Let $\mathcal{O}$ be a regular well-ordered GOF game. Then, for any fixed $\theta \geq 0$, we have that*

$$\mathsf{PoA}_\lambda(\mathcal{O}) \leq \sup_{\substack{f\in\mathcal{F}(\mathcal{O}),\\ g\in\mathcal{G}(\mathcal{O}),\\ x,y,\hat{x},\hat{y}\in[0,1]}} \max\left\{\eta_{2\lambda}(\theta, f, x, y), \eta_{1-\lambda}(\theta, g, \hat{x}, \hat{y})\right\}. \tag{1}$$

Proof. Given a GOF game $\mathcal{O}$ and two opinion profiles $\boldsymbol{x}$ and $\boldsymbol{y}$, we will shorten the notation by setting $x_{i,j} := |x_i - x_j|$, $\hat{x}_i := |x_i - s_i|$, $y_{i,j} := |y_i - y_j|$, and $\hat{y}_i := |y_i - s_i|$ for any i, j. We observe that the above quantities belong to interval $[0, 1]$. We start by giving a preliminary lemma.

Lemma 5. *Let $\mathcal{O}$ be a regular well-ordered GOF game. We have that $\langle (c_i(\boldsymbol{x}))_i, \boldsymbol{y} - \boldsymbol{x}\rangle \leq \sum_{i>j}(y_{i,j} - x_{i,j})\frac{\partial f_{i,j}(x_{i,j})}{\partial x_{i,j}} + \sum_i(\hat{y}_i - \hat{x}_i)\frac{\partial g_i(\hat{x}_i)}{\partial \hat{x}_i}$ for any opinion profiles $\boldsymbol{x}, \boldsymbol{y}$ of $\mathcal{O}$.*

To show the theorem we resort to the primal-dual method [Bilò, 2018]. Let $\boldsymbol{x}$ be an equilibrium of $\mathcal{O}$ maximizing the λ-social influence cost and let $\boldsymbol{y}$ be an opinion profile of $\mathcal{O}$ minimizing the λ-social influence cost. Since $\mathsf{PoA}_\lambda(\mathcal{O})$ is upper bounded by a constant for any $\lambda \in (0, 1)$ (by Theorem 6), we assume w.l.o.g. that $OPT_\lambda(\mathcal{O}) > 0$.

Let us consider the following linear program LP in variables $\alpha_{i,j} \geq 0$ and $\beta_i \geq 0$:

$$\max\ 2\lambda\sum_{i>j}\alpha_{i,j}f_{i,j}(x_{i,j}) + (1-\lambda)\sum_i \beta_i g_i(\hat{x}_i), \tag{2}$$

$$s.t.\ \sum_{i>j}\alpha_{i,j}(y_{i,j} - x_{i,j}))\frac{\partial f_{i,j}(x_{i,j})}{\partial x_{i,j}} + \sum_i \beta_i(\hat{y}_i - \hat{x}_i)\frac{\partial g_i(\hat{x}_i)}{\partial \hat{x}_i} \geq 0, \tag{3}$$

$$2\lambda\sum_{i>j}\alpha_{i,j}f_{i,j}(y_{i,j}) + (1-\lambda)\sum_i \beta_i g_i(\hat{y}_i) = 1. \tag{4}$$

We have that the optimal value of LP is an upper bound on $\mathsf{PoA}_\lambda(\mathcal{O})$. Indeed, by considering the assignment of variables $(\alpha_{i,j})_{i,j}$ and $(\beta_i)_i$ such that $\alpha_{i,j} = 1/OPT_\lambda(\mathcal{O})$ for any i, j and $\beta_i = 1/OPT_\lambda(\mathcal{O})$ for any i, we have the following facts: (i) (2) is equal to $\mathsf{SUM}(\boldsymbol{x})/\mathsf{SUM}(\boldsymbol{y}) = \mathsf{PoA}_\lambda(\mathcal{O})$; (ii) the left-hand part of (3) is at most equal to $\langle (c_i(\boldsymbol{x}))_i, \boldsymbol{y} - \boldsymbol{x}\rangle$ (by Lemma 5), and $\langle (c_i(\boldsymbol{x}))_i, \boldsymbol{y} - \boldsymbol{x}\rangle$ is non-negative as $\boldsymbol{x}$ is an equilibrium (by Lemma 4), thus (4) is satisfied by the considered assignment; (iii) the left-hand part of (4) is $\mathsf{SUM}(\boldsymbol{y})/\mathsf{SUM}(\boldsymbol{y}) = 1$, thus (4) is satisfied by the considered assignment.

We conclude that there exists a feasible assignment of LP whose optimal value is $\mathsf{PoA}_\lambda(\mathcal{O})$. Now, we compute the dual DLP of LP, which has two dual variables $\theta \geq 0$ and η associated to the first and the second constraint of LP, respectively:

$$\min\ \eta,\quad s.t.\ 2\lambda\cdot f_{i,j}(y_{i,j})\cdot\eta \geq 2\lambda\cdot f_{i,j}(x_{i,j}) + \theta\cdot(y_{i,j} - x_{i,j})\cdot\frac{\partial f_{i,j}(x_{i,j})}{\partial x_{i,j}},\ \forall i \neq j,$$

$$(1-\lambda)\cdot g_i(\hat{y}_i)\cdot\eta \geq (1-\lambda)\cdot g_i(\hat{x}_i) + \theta\cdot(\hat{y}_i - \hat{x}_i)\cdot\frac{\partial g_i(\hat{x}_i)}{\partial \hat{x}_i}, \forall i,$$

Fix $\theta \geq 0$ and let η be equal to the right-hand part of inequality (1). If $\eta = \infty$ then the claim easily follows, as ∞ is a trivial upper bound on the λ-price of anarchy of $\mathcal{O}$. If $\eta < \infty$, by exploiting the structure of DLP, we have that (θ, η) is a feasible solution of DLP. Thus, by weak-duality, η is an upper bound on the optimal value of LP, i.e., it is an upper bound on the λ-price of anarchy of $\mathcal{O}$, and the claim follows. □

Remark 4. By setting $\theta := \max\{2\lambda, 1-\lambda\}$, the right-hand part of (1) is at most equal to $\max\{2\lambda, 1-\lambda\}/\min\{2\lambda, 1-\lambda\}$, thus showing that the upper bound provided in Theorem 7 cannot be worse than the one given in Theorem 6. Furthermore, if the set of non-null private influence functions is empty, by setting $\theta = 2\lambda$, the right-hand part of (1) gets equal to 1, i.e., the λ-price of anarchy equals 1.

Tight Lower Bounds. Under mild assumptions, the proof arguments of Theorem 7 can be reversed via strong duality (by following a similar approach as in [Bilò and Vinci, 2017; Bilò *et al.*, 2018b; Benita *et al.*, 2020]) to derive tight lower bounds for the λ-price of stability, holding even for games with two agents. The general structure of the lower bound is the following: (i) we have two agents with private beliefs equal to $1/2 - s$ and $1/2 + s$ for some $s \in [0, 1/2]$ (ii) there is a unique equilibrium $\boldsymbol{r} = (1/2 - r, 1/2 + r)$ for some $r \in [0, 1/2]$ and the social optimum is $\boldsymbol{t} = (1/2 - t, 1/2 + t)$ for some $t \in [0, 1/2]$.

Applications. The upper bound of Theorem 7 can be applied to derive tight (or almost tight) bounds on the λ-price of anarchy of GOF games with cost functions $h_{i,j}(x)$ of type $\alpha_{i,j}x^p$ or $\alpha_{i,j}/x$, for fixed $\alpha_{i,j} \geq 0$ and $p > 1$. We stress that cost functions of type $\alpha_{i,j}x^p$ have been already considered for standard opinion formation games (see, for instance, [Bindel *et al.*, 2015; Bhawalkar *et al.*, 2013]).

6 Conclusions

In this work, we provided a new model for opinion formation that encompasses social group membership, and both attraction and repulsion among agents. In this way, we try to model many aspects of opinion formation that occur in real-world examples, such as youth subcultures or political parties. We proved that equilibria always exist and provided tight bounds on their quality.

We believe that our model can be useful for analyzing and forecasting the diffusion of opinion in social networks, and suggesting specific strategies for marketing (e.g., for target advertising [Kempe *et al.*, 2015]) and for election control [Wilder and Vorobeychik, 2018]. Another interesting direction would be to embed our opinion formation process in an evolving setting: this would give useful hints on the processes that lead to radical changes in cultures and styles.

Acknowledgments

This work was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets" and by "GNCS–INdAM".

References

[Acar *et al.*, 2017] E. Acar, G. Greco, and M. Manna. Group reasoning in social environments. In *AAMAS*, pages 1296–1304, 2017.

[Auletta *et al.*, 2015] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. Minority becomes majority in social networks. In *WINE*, pages 74–88, 2015.

[Auletta *et al.*, 2016] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. Generalized discrete preference games. In *IJCAI*, pages 53–59, 2016.

[Auletta *et al.*, 2017] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. Information retention in heterogeneous majority dynamics. In *WINE*, pages 30–43, 2017.

[Auletta *et al.*, 2018] V. Auletta, D. Ferraioli, and G. Greco. Reasoning about consensus when opinions diffuse through majority dynamics. In *IJCAI*, pages 49–55, 2018.

[Auletta *et al.*, 2019] V. Auletta, A. Fanelli, and D. Ferraioli. Consensus in opinion formation processes in fully evolving environments. In *AAAI*, pages 6022–6029, 2019.

[Benita *et al.*, 2020] F. Benita, V. Bilò, B. Monnot, G. Piliouras, and C. Vinci. Data-driven models of selfish routing: Why price of anarchy does depend on network topology. In *WINE*, pages 252–265, 2020.

[Bhawalkar *et al.*, 2013] K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *STOC*, pages 41–50, 2013.

[Bilò and Vinci, 2017] V. Bilò and C. Vinci. On the impact of singleton strategies in congestion games. In *ESA*, pages 17:1–17:14, 2017.

[Bilò *et al.*, 2011] V. Bilò, M. Flammini, G. Monaco, and L. Moscardelli. On the performances of nash equilibria in isolation games. *Journal of Combinatorial Optimization*, 22(3):378 – 391, 2011.

[Bilò *et al.*, 2018a] V. Bilò, A. Fanelli, and L. Moscardelli. Opinion formation games with dynamic social influences. *Theoretical Computer Science*, 746:444–458, 2018.

[Bilò *et al.*, 2018b] V. Bilò, L. Moscardelli, and C. Vinci. Uniform mixed equilibria in network congestion games with link failures. In *ICALP*, pages 146:1–146:14, 2018.

[Bilò *et al.*, 2020] V. Bilò, M. Flammini, and C. Vinci. The quality of content publishing in the digital era. In *ECAI*, pages 35–42, 2020.

[Bilò, 2018] V. Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory of Computing Systems*, 62(5):1288 – 1317, 2018.

[Bindel *et al.*, 2015] D. Bindel, J. M. Kleinberg, and S. Oren. How bad is forming your own opinion? *Games and Economic Behavior*, 92:248–265, 2015.

[Chierichetti *et al.*, 2018] F. Chierichetti, J. Kleinberg, and S. Oren. On discrete preferences and coordination. *Journal of Computer and System Sciences*, 93:11 – 29, 2018.

[Degroot, 1974] M. Degroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[Epitropou *et al.*, 2019] M. Epitropou, D. Fotakis, M. Hoefer, and S. Skoulakis. Opinion formation games with aggregation and negative influence. *Theory Comput. Syst.*, 63(7):1531–1553, 2019.

[Ferraioli and Ventre, 2017] D. Ferraioli and C. Ventre. Social pressure in opinion games. In *IJCAI*, pages 3661–3667, 2017.

[Ferraioli *et al.*, 2016] D. Ferraioli, P. Goldberg, and C. Ventre. Decentralized dynamics for finite opinion games. *Theoretical Computer Science*, 648:96–115, 2016.

[Fotakis *et al.*, 2016] D. Fotakis, D. Palyvos-Giannas, and S. Skoulakis. Opinion dynamics with local interactions. In *IJCAI*, pages 279–285, 2016.

[Fotakis *et al.*, 2018] D. Fotakis, V. Kandiros, V. Kontonis, and S. Skoulakis. Opinion dynamics with limited information. In *WINE*, pages 282–296, 2018.

[Friedkin and Johnsen, 1990] N. Friedkin and E. Johnsen. Social influence and opinions. *The Journal of Mathematical Sociology*, 15(3-4):193–206, 1990.

[Hegselmann and Krause, 2002] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5:1–24, 2002.

[Kempe *et al.*, 2015] D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147, 2015.

[Monderer and Shapley, 1996] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[Roughgarden and Schoppmann, 2015] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory*, 156:317–342, 2015.

[Serafino and Ventre, 2016] P. Serafino and C. Ventre. Heterogeneous facility location without money. *Theoretical Computer Science*, 636(C):27–46, 2016.

[Wilder and Vorobeychik, 2018] B. Wilder and Y. Vorobeychik. Controlling elections through social influence. In *AAMAS*, pages 265–273, 2018.

[Zhao *et al.*, 2008] Y. Zhao, W. Chen, and S.H. Teng. The isolation game: A game of distances. In *ISAAC*, pages 148–159, 2008.