General Opinion Formation Games with Social Group Membership (Discussion Paper)*

Vittorio Bilò[1,†], Diodato Ferraioli[2,†] and Cosimo Vinci[1,2,*,†]

[1] Università del Salento, Piazza Tancredi, 7, 73100 Lecce, Italy
[2] Università degli Studi di Salerno, Via Giovanni Paolo II, 132, 84084 Fisciano (SA), Italy

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 performance 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 $\lambda \in [1/5, 1/2]$. Moreover, in many settings, we provide even better upper bounds on the prices of anarchy and stability, which are tight under mild assumptions.

Keywords

Opinion Formation Games, Pure Nash Equilibria, Price of Anarchy, Price of Stability

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.

Discussion Papers - 22nd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022)

* This work was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets", by "GNCS–INdAM" and by the PON R&I 2014-2020 Project TEBAKA "Sistema per acquisizione conoscenze di base del territorio".

An extended version of this work appears in the proceedings of IJCAI 2022 [1].

* Corresponding author.
† These authors contributed equally.
vittorio.bilo@unisalento.it (V. Bilò); dferraioli@unisa.it (D. Ferraioli); cosimo.vinci@unisalento.it (C. Vinci)
0000-0001-7848-4904 (V. Bilò); 0000-0002-7962-5200 (D. Ferraioli); 0000-0001-7741-9342 (C. Vinci)

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 [2, 3], 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) [4], 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 [5, 6], or limited/local interactions [7, 8], or dynamic settings where social relationships and internal beliefs evolve over time [9, 10, 11, 12, 13].

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 [14, 15] 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). This degree of generality, but only restricted to attraction phenomenon, has been considered before only by [10].

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 [14, 15]), 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 influence 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 influence 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., [3, 10]), 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, e.g., facility location with heterogeneous preferences [16], content publishing [17] and isolation games [18, 19].

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., [10, 11, 3, 5, 2, 20, 6, 8]). 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 [3, 5, 10, 11]. 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 [21, 22, 23]. A third approach, finally, measures the distance from a *consensus* [9, 13].

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 performance, 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 [24], 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 classes of well-studied games, by proving tight numerical bounds.

Due to space limitation, for all claims in the next sections, proofs are omitted. We refer the interested reader to the full version for more details.

2. Model and Definitions

Generalized Opinion Formation Games. Let $N := [n]$ be a set of n *agents*, where $[n]$ denotes the set $\{1, \ldots, n\}$. 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).

We define $\boldsymbol{s} = (s_1, \ldots, s_n)$ to 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 [4, 3] and their generalization considered in [10].

A GOF game $\mathcal{O}$ is *well-ordered* if it is convex, and we can organize the agents in clusters $S_1, S_2, \ldots, S_k$ such that:

i each cluster is a non-empty set of consecutive agents;

ii 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 sub-game restricted to each cluster is an aggregation game;

iii 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 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 [4, 3] 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, [3, 10]).

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 $f_{i,j}(0) = \infty$.

Given $\lambda \in (0, 1)$, the *λ-social influence cost* $\mathsf{SUM}_\lambda(\boldsymbol{x})$ of the opinion profile $\boldsymbol{x}$ is defined as $\sum_i(\lambda \cdot c_{pu,i}(\boldsymbol{x}) + (1-\lambda) \cdot c_{pr,i}(\boldsymbol{x})) = 2\lambda \sum_{i>j} f_{i,j}(|x_i - x_j|) + (1-\lambda) \sum_i g_i(|x_i - s_i|)$, 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 $OPT_\lambda(\mathcal{O}) := \inf_{\boldsymbol{x} \in \mathsf{SP}(\mathcal{O})} \mathsf{SUM}_\lambda(\boldsymbol{x})$.

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 performance 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 performance of an equilibrium of $\mathcal{O}$ and the optimal λ-social influence cost of $\mathcal{O}$.

3. Equilibrium Existence

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 [25], 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.

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

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

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.*

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* $\lambda \in (0,1)$.

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* $\mathsf{PoS}_\lambda(\mathcal{O}) \leq \frac{\max\{2\lambda, 1-\lambda\}}{\min\{2\lambda, 1-\lambda\}}$. *This result is tight. Indeed, 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* $\mathsf{PoS}_\lambda(\mathcal{O}) \geq \frac{\max\{2\lambda, 1-\lambda\}}{\min\{2\lambda, 1-\lambda\}} - \epsilon$.

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

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 [24] 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 [26, 10].

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

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

We additionally show that the upper bound of Theorem 5 if often tight, even for the price of stability. Indeed, under mild assumptions, the proof arguments used to obtain the upper bound can be reversed via strong duality (by following a similar approach as in [27, 28, 29]) to derive tight lower bounds for the 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 $(x_1, x_2) = (1/2 - r, 1/2 + r)$ for some $r \in [0, 1/2]$ and the social optimum is $(y_1, y_2) = (1/2 - t, 1/2 + t)$ for some $t \in [0, 1/2]$.

Finally, as an application of Theorem 6, we show how our findings apply to some specific classes of games. The upper bound of Theorem 6 can be applied to derive tight bounds on the λ-price of anarchy of aggregation games with influence functions of type αx^p, where $p > 1$ is fixed and $\alpha \geq 0$ depends on the agent indexes (i.e., $\alpha := \alpha_i$ for private influence functions and $\alpha := \alpha_{i,j}$ for public ones); we stress that such influence functions have been already considered for standard opinion formation games (see, for instance, [3, 10]). As an additional application, we also study the λ-price of anarchy of isolation games with public influence functions of type α/x and generic private influence functions.

5. Future Directions

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 [30]) and for election control [31]. 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.

References

[1] V. Bilò, D. Ferraioli, C. Vinci, General opinion formation games with social group membership, in: IJCAI, 2022, pp. 88–94.

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

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

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

[5] F. Chierichetti, J. Kleinberg, S. Oren, On discrete preferences and coordination, Journal of Computer and System Sciences 93 (2018) 11 – 29.

[6] D. Ferraioli, P. Goldberg, C. Ventre, Decentralized dynamics for finite opinion games, Theoretical Computer Science 648 (2016) 96–115.

[7] D. Fotakis, V. Kandiros, V. Kontonis, S. Skoulakis, Opinion dynamics with limited information, in: WINE, 2018, pp. 282–296.

[8] D. Fotakis, D. Palyvos-Giannas, S. Skoulakis, Opinion dynamics with local interactions, in: IJCAI, 2016, pp. 279–285.

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

[10] K. Bhawalkar, S. Gollapudi, K. Munagala, Coevolutionary opinion formation games, in: STOC, 2013, pp. 41–50.

[11] V. Bilò, A. Fanelli, L. Moscardelli, Opinion formation games with dynamic social influences, Theoretical Computer Science 746 (2018) 444–458.

[12] D. Ferraioli, C. Ventre, Social pressure in opinion games, in: IJCAI, 2017, pp. 3661–3667.

[13] V. Auletta, A. Fanelli, D. Ferraioli, Consensus in opinion formation processes in fully evolving environments, in: AAAI, 2019, pp. 6022–6029.

[14] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Generalized discrete preference games, in: IJCAI, 2016, pp. 53–59.

[15] E. Acar, G. Greco, M. Manna, Group reasoning in social environments, in: AAMAS, 2017, pp. 1296–1304.

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

[17] V. Bilò, M. Flammini, C. Vinci, The quality of content publishing in the digital era, in: ECAI, 2020, pp. 35–42.

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

[19] Y. Zhao, W. Chen, S. Teng, The isolation game: A game of distances, in: ISAAC, 2008, pp. 148–159.

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

[21] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Minority becomes majority in social networks, in: WINE, 2015, pp. 74–88.

[22] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Information retention in

heterogeneous majority dynamics, in: WINE, 2017, pp. 30–43.

[23] V. Auletta, D. Ferraioli, G. Greco, Reasoning about consensus when opinions diffuse through majority dynamics, in: IJCAI, 2018, pp. 49–55.

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

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

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

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

[28] V. Bilò, L. Moscardelli, C. Vinci, Uniform mixed equilibria in network congestion games with link failures, in: ICALP, 2018, pp. 146:1–146:14.

[29] F. Benita, V. Bilò, B. Monnot, G. Piliouras, C. Vinci, Data-driven models of selfish routing: Why price of anarchy does depend on network topology, in: WINE, 2020, pp. 252–265.

[30] D. Kempe, J. M. Kleinberg, E. Tardos, Maximizing the spread of influence through a social network, Theory of Computing 11 (2015) 105–147.

[31] B. Wilder, Y. Vorobeychik, Controlling elections through social influence, in: AAMAS, 2018, pp. 265–273.