



On (approximate) pure Nash equilibria in preference games

Angelo Fanelli, Dimitris Fotakis

► To cite this version:

Angelo Fanelli, Dimitris Fotakis. On (approximate) pure Nash equilibria in preference games. 2022. hal-03865843

HAL Id: hal-03865843

<https://hal.science/hal-03865843>

Preprint submitted on 22 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On (approximate) pure Nash equilibria in preference games

Angelo Fanelli,¹ Dimitris Fotakis,²

¹ CNRS, France

² National Technical University of Athens, Greece

Abstract

In this work we introduce and study a general model of preference game. Our game belongs to the large class of games in which each agent aims at compromising between her innate preferences and the choices of the neighbours. Our game encompasses two models, i.e., the full information and the partial information settings. While in the former each agent is fully aware of the strategies of the neighbours, in the latter each agent has only access to a single representative strategy, which can be regarded as the output of an aggregation function. Both settings are quite general and allow for a systematic study of the fundamental questions related to the existence and efficiency of (approximate) equilibria in the strategic interaction of agents in social networks. Interestingly, our bounds on the price of anarchy are very general and are derived in terms of key features of the game (i.e., social impact, boundary and stretch). Moreover, our results show that the two settings behave very similarly, thus confirming our conjecture that in general the lack of information does not affect the existence and efficiency of the (approximate) equilibria.

1 Introduction

We refer to preference games as the class of games in which a set of finite agents share a common set of strategies, that are embedded in an (approximate) metric space, and in which each agent aims at choosing the strategy that represents a good compromise between her preferred strategy and the neighbours' strategies. An agent's cost is typically formulated as the sum of two quantities: the first quantity is the cost for choosing a strategy different from her preferred one, usually quantified as the (approximate) distance between her strategy and the preferred one, while the second quantity is the cost for choosing a strategy which is far from the neighbours' strategies.

Numerous games studied in the literature can be cast as preference games, including models of opinion formation (see e.g., (Bindel, Kleinberg, and Oren 2015; Bhawalkar, Gollapudi, and Munagala 2013; Ghaderi and Srikant 2014)), or their variants such as discrete preference games, (see e.g., (Chierichetti, Kleinberg, and Oren 2018; Auletta et al. 2016, 2017; Lolakapuri et al. 2019)).

In this paper we introduce and study a quite general model of preference game which encompasses the full and the partial information settings. In the full information setting, inspired by the discrete preference model introduced in (Chierichetti, Kleinberg, and Oren 2018), we assume that each agent is fully aware of the strategies of her neighbours and respond optimally to them. Instead, in the partial information setting, we assume that each agent has access to a single representative strategy, which can be regarded as the output of an aggregation function that condenses the strategies of her neighbors into a single one. In order to motivate this latter model, we may think of the traditional political voting scenario, where voters have fixed innate preferences over the candidates, but polls (which is a form of preference aggregation) might cause the voters to change their vote (see also (Epitropou et al. 2019), where the model involves an estimation of the average opinion of the agents).

Before presenting the related work (Section 1.3) and outline our conceptual and technical contributions (Section 1.4), we formally describe our model (Section 1.1) and introduce technical concepts (Section 1.2). We next present our results on the existence of (approximate) pure Nash equilibria (Section 2) and the price of anarchy (Section 3 and Section 4). We conclude the paper with an application of our findings on a special case of preference game which we call k -approval voting game (Section 5).

1.1 Preference games

We assume a finite population N of $n \geq 2$ agents and an underlying strategy space Z , which may be either *discrete* or *continuous*. Each agent i has a fixed *preferred strategy* s_i . Let $S = \{s_1, \dots, s_n\}$. One of the key ingredient of our model is to allow preferred strategies not to belong to Z . This ensures the agents to have more elaborate preferred strategies and to account for preferences that possibly cannot be fully disclosed in public. So, we assume a strategy “universe” U , and that $Z \subseteq U$ (e.g., Z may be discrete, while U may be the convex hull of Z) and $S \subseteq U$. Z and S could be disjoint or not. So, we let $N_S = \{i \in N : s_i \in Z\}$ denote the set of agents with preferred strategy in Z ; we say that the game is *restricted* if $N_S = \emptyset$, *unrestricted* if $N_S = N$, and *semi-restricted* otherwise.

We assume a distance function $d : U \times U \rightarrow \mathbb{R}_{\geq 0}$, which quantifies dissimilarity between the strategies; we let d be

a *approximate metric* over U (see Section 1.2 for a formal definition). This function will be used to quantify the agents' costs in a given strategy profile.

The influence exercised by j 's strategy on the strategy choice of agent i is quantified by the *influence weight* $w_{ij} \in [0, 1]$. We refer to the set of agents j such that $w_{ij} > 0$ as the set of i 's *neighbours*. Another key ingredient of our model is to assume that the influence weights w_{ij} are *not necessarily symmetric*, have $w_{ii} = 0$ and are *normalized* so that $\sum_j w_{ij} = 1$, i.e., the total influence exercised to any agent i sums up to 1. For every agent i , we denote by $\vec{w}_i$ the vector of influence weights $(w_{ij})_{j \in N}$.

We refer to any vector $\vec{z} \in Z^n$ as a *state (strategy profile)* of the game, in which the i -th component, denoted by $\vec{z}(i)$, represents the strategy of i . If $\vec{z} = (z, \dots, z)$, we say that $\vec{z}$ is a *consensus* on z . We denote by $\vec{z}_{-i}$ the partial strategy profile obtained from $\vec{z}$ after removing i 's strategy.

The overall cost $c_i(\vec{z})$ incurred by i in $\vec{z}$ can be summarized as the convex combination between two components, i.e.,

$$c_i(\vec{z}) = \alpha d(s_i, \vec{z}(i)) + (1 - \alpha) K(\vec{w}_i, \vec{z}),$$

where the first component $d(s_i, \vec{z}(i))$ is the *innate cost* of i for $\vec{z}(i)$, and the second component $K(\vec{w}_i, \vec{z})$ is the *disagreement cost*; the agents only influence each other through this latter term. In this paper we will propose two different definitions of K , corresponding to the *full information* and *partial information* settings. The coefficient $\alpha \in (0, 1)$ quantifies the confidence of each agent i on her preferred strategy s_i , and often referred to as the *self-confidence* of agent i . We say the agents are *compliant* when $\alpha \in (0, 1/2)$, *balanced* when $\alpha = 1/2$, and *stubborn* otherwise. As usual, each agent i selects her strategy $\vec{z}(i)$ so as to minimize her cost.

Full information. We assume that each agent has full access to the strategies of the neighbours. We define the disagreement cost of i as

$$K(\vec{w}_i, \vec{z}) = \sum_{j \neq i} w_{ij} d(\vec{z}(i), \vec{z}(j)). \quad (1)$$

When K is defined as in (1), we will denote the cost of each agent i in $\vec{z}$ as $c_i^F(\vec{z})$.

Partial information. We assume that each agent does not have full access to the strategies of the neighbours but to an aggregate value. The major ingredient is an *aggregation function* aggr which, for each agent i , maps $\vec{w}_i$ and $\vec{z}_{-i}$ to an aggregated strategy $\text{aggr}(\vec{w}_i, \vec{z}_{-i}) \in Z$ that “summarizes” the strategies in $\vec{z}_{-i}$. For brevity, we will usually write $\text{aggr}_i(\vec{z}_{-i})$, instead of $\text{aggr}(\vec{w}_i, \vec{z}_{-i})$. We define the disagreement cost of i as

$$K(\vec{w}_i, \vec{z}) = d(\vec{z}(i), \text{aggr}_i(\vec{z}_{-i})). \quad (2)$$

Our results hold for general aggregation functions that satisfy *unanimity* i.e., if $\vec{x}_{-i} = (x, \dots, x)$ then $\text{aggr}_i(\vec{x}_{-i}) = x$, and *consistency* i.e., if $\vec{x}$ and $\vec{y}$ satisfy $\sum_{j \neq i} w_{ij} d(\vec{x}(j), \vec{y}(j)) = 0$ (this quantity corresponds to the relative distance $\pi_i(\vec{x}, \vec{y})$ between $\vec{x}$ and $\vec{y}$, as defined in

Section 1.2) then $\text{aggr}_i(\vec{x}_{-i}) = \text{aggr}_i(\vec{y}_{-i})$. We refer to aggregation functions that satisfy unanimity and consistency as *feasible* (see Section 1.2 for some examples of feasible aggregation functions). We remark that choosing the aggregation value as an element of Z , instead of U , strongly affects the behaviour of the game. On the other hand, choosing aggr_i as an element of U requires more assumptions (e.g., on the strategy space, the aggregation function, or the distance function) in order to derive significant results; we leave this as an interesting direction for further research.

When K is defined as in (2), we will denote the cost of each agent i in $\vec{z}$ as $c_i^P(\vec{z})$.

A case study: k -approval voting game. As an application of our findings, we introduce an example of preference game in which the strategy space Z consists of all binary strings of length m with k ones, while the preferred strategies lie in the convex hull of Z . Formally, $Z = \{\vec{u} \in \{0, 1\}^m \mid \sum_{\ell=1}^m \vec{u}(\ell) = k\}$, $U = \{\vec{u} \in [0, 1]^m \mid \sum_{\ell=1}^m \vec{u}(\ell) = k\}$ and an agent i can choose any $s_i \in U$ as her preferred strategy. We refer to the case $m = 2$ and $k = 1$ as the *binary strategies* setting, since $|Z| = 2$. We assume that the distance function d is L_2^2 ; namely, for every $\vec{u}, \vec{v} \in U$,

$$d(\vec{u}, \vec{v}) = \sum_{\ell=1}^m (\vec{u}(\ell) - \vec{v}(\ell))^2.$$

As a special case of preference game, also here we can distinguish two versions of the game, namely the full and the partial information. This game is motivated by the k -approval voting and generalises the opinion formation model with binary strategies (Ferraioli, Goldberg, and Ventre 2016); we discuss in more details about this relation in Section 1.3.

1.2 Notation, Definitions and Preliminaries

In this section, we introduce some additional notation and discuss some important preliminaries.

Approximate metric. We assume that, for some $1 \leq \rho_Z \leq \rho$, $d : U \times U \mapsto \mathbb{R}_{\geq 0}$ satisfies (i) $d(x, x) = 0$, for all $x \in U$; (ii) symmetry, i.e., $d(x, y) = d(y, x)$, for all $x, y \in U$; (iii) (approximate) triangle inequality, i.e., $d(x, y) \leq \rho(d(x, z) + d(z, y))$, for all $x, y, z \in U$; and (iv) $d(x, y) \leq \rho_Z(d(x, z) + d(z, y))$, for all $x, y, z \in Z$. In this case, we say that d is a (ρ, ρ_Z) -*approximate metric*. We say that d is an *exact metric* when $\rho = \rho_Z = 1$, and that d is *uniform* when it is an exact metric such that $d(x, y) = 1$ for all $x \neq y$.

Examples of Feasible Aggregation Functions. These are aggregation functions that satisfy *unanimity* and *consistency*. Notable examples of feasible aggregation functions are the Fréchet median and the Fréchet mean. Given any state $\vec{z}$, the *Fréchet median* of agent i in $\vec{z}$, denoted by $\text{med}_i(\vec{z}_{-i})$, is any strategy that minimizes the weighted sum of its distances to the strategies in $\vec{z}_{-i}$, i.e.,

$$\text{med}_i(\vec{z}_{-i}) \in \arg \min_{y \in Z} \sum_{j \neq i} w_{ij} d(y, \vec{z}(j)). \quad (3)$$

The *Fréchet mean* of agent i in $\vec{z}$, denoted by $\text{mean}_i(\vec{z}_{-i})$, is any strategy in Z that minimizes the weighted sum of its squared distances to the strategies in $\vec{z}_{-i}$, i.e.,

$$\text{mean}_i(\vec{z}_{-i}) \in \arg \min_{y \in Z} \sum_{j \neq i} w_{ij} d^2(y, \vec{z}(j)). \quad (4)$$

Approximate Pure Nash Equilibria, Social Optima and Price of Anarchy. For any state $\vec{z}$ and any strategy z , we let $[\vec{z}_{-i}, z]$ be the new state obtained from $\vec{z}$ by replacing its i -component $\vec{z}(i)$ with z and keeping the remaining components unchanged, i.e., $[\vec{z}_{-i}, z](i) = z$ and $[\vec{z}_{-i}, z](j) = \vec{z}(j)$ for every $j \neq i$. For $\gamma \geq 1$, a strategy $z^* \in Z$ is a γ -approximate best response of agent i wrt. $\vec{z}_{-i}$ if $c_i([\vec{z}_{-i}, z^*]) \leq \gamma c_i([\vec{z}_{-i}, z])$, for every $z \in Z$ (when $\gamma = 1$, z^* is simply called *best response*). We say that a state $\vec{z}$ is a γ -approximate pure Nash equilibrium (or simply γ -equilibrium, for brevity) if every agent i is playing a γ -approximate stable strategy wrt. $\vec{z}_{-i}$, i.e., for every agent i and every strategy $z \in Z$, we have $c_i(\vec{z}) \leq \gamma c_i([\vec{z}_{-i}, z])$; if $\gamma = 1$, we simply refer to $\vec{z}$ as an *exact pure Nash equilibrium* (or simply *equilibrium*). We use $E_\gamma \subseteq Z^n$ to denote the set of all γ -approximate pure Nash equilibria of a given preference game. We say that a strategy $x^* \in Z$ is a γ -approximate dominant strategy of agent i if $c_i([\vec{x}_{-i}, x^*]) \leq \gamma c_i([\vec{x}_{-i}, y])$, for all states $\vec{x}_{-i}$ and $y \in Z$ (when $\gamma = 1$, x^* is simply called *dominant strategy*).

We measure the efficiency of each state $\vec{z}$ according to a *social objective*. We consider two social objectives, the *social cost* $\text{SUM}(\vec{z}) = \sum_{i \in N} c_i(\vec{z})$ and the *maximum cost* $\text{MAX}(\vec{z}) = \max_{i \in N} c_i(\vec{z})$. A state $\vec{o}$ is *optimal* wrt. SUM , if $\text{SUM}(\vec{o}) \leq \text{SUM}(\vec{z})$, for all states $\vec{z}$. We let $O_{\text{SUM}} \subseteq Z^n$ be the set of optimal states wrt. SUM , i.e., $O_{\text{SUM}} = \arg \min_{\vec{z} \in Z^n} \text{SUM}(\vec{z})$. Optimal states wrt. MAX and O_{MAX} are defined similarly.

For any $\vec{o} \in O_{\text{SUM}}$, the *price of anarchy* of a game wrt. SUM is $\text{POA}^{\text{SUM}} = \max_{\vec{e} \in E_1} \frac{\text{SUM}(\vec{e})}{\text{SUM}(\vec{o})}$, if $\text{SUM}(\vec{o}) > 0$; if $\text{SUM}(\vec{o}) = 0$ then $\text{POA}^{\text{SUM}} = +\infty$, when $E_1 \neq O_{\text{SUM}}$, and $\text{POA}^{\text{SUM}} = 1$, when $E_1 = O_{\text{SUM}}$. The definition of POA^{MAX} is similar.

Equivalence and Relative Distance between States. For every pair of states $\vec{x}, \vec{y}$, $D(\vec{x}, \vec{y}) = \{j \in N : \vec{x}(j) \neq \vec{y}(j)\}$ denotes the set of agents with different strategies in $\vec{x}$ and $\vec{y}$. If $D(\vec{x}, \vec{y}) = \emptyset$, we say that $\vec{x}$ and $\vec{y}$ are *globally equivalent*. For every agent i and all pairs of states $\vec{x}, \vec{y}$, we define the *relative distance* between $\vec{x}$ and $\vec{y}$ for i as $\pi_i(\vec{x}, \vec{y}) = \sum_{j \neq i} w_{ij} d(\vec{x}(j), \vec{y}(j))$. Notice that if $\vec{y} = [\vec{x}_{-i}, x]$, for any x , then $\pi_i(\vec{x}, \vec{y}) = 0$. When $\pi_i(\vec{x}, \vec{y}) = 0$, we say that $\vec{x}$ and $\vec{y}$ are *equivalent* for i . Notice that $D(\vec{x}, \vec{y}) = \emptyset$ implies $\pi_i(\vec{x}, \vec{y}) = 0$ (while the converse may not be true).

Social Impact, Boundary and Stretch. These are the three features of the game which fully characterize the inefficiency of the equilibria, as we will see in Section 3.

The *social impact* of agent i is $\delta_i = \sum_{j \in N} w_{ji}$ and quantifies the intensity by which i influences the environment. The *global social impact* is $\delta = \max_{j \in N} \delta_j$; it quantifies the asymmetry between the influence received and exercised

by any agent in the social network. Since $\sum_{j \neq i} w_{ij} = 1$ for every i , we observe that $\sum_{i \in N} \delta_i = n$ and $\delta \in [1, n-1]$.

The *boundary* of agent i , denoted by β_i quantifies how much closer a strategy x can be to s_i compared against an equilibrium strategy $\vec{e}(i)$ of i . Formally,

$$\beta_i = \min_{\vec{e} \in E, x \neq \vec{e}(i)} \frac{d(x, s_i)}{d(x, \vec{e}(i))}. \quad (5)$$

For nontrivial instances, there always exists a strategy x with $d(x, \vec{e}(i)) > 0$. Thus, β_i is well defined. The (*global*) *boundary* is $\beta = \min_{j \in N} \beta_j$.

The *stretch* τ_i of agent i quantifies how sensitive the aggregation function is wrt. the changes of the state of the game. Formally, we let

$$\hat{\tau}_i = \inf_{\substack{\vec{e} \in E \\ \vec{y} \in Z^n \\ \pi_i(\vec{e}, \vec{y}) > 0}} \left\{ r \geq 0 \mid \frac{d(\text{aggr}_i(\vec{e}_{-i}), \text{aggr}_i(\vec{y}_{-i}))}{\pi_i(\vec{e}, \vec{y})} \leq r \right\}.$$

We define $\tau_i = \max\{\hat{\tau}_i, 1\}$, so as to account for the case where $\pi_i(\vec{e}, \vec{y}) = 0$ and $d(\text{aggr}_i(\vec{e}_{-i}), \text{aggr}_i(\vec{y}_{-i})) = 0$. Since τ_i is used in the price of anarchy bounds, we can restrict its definition to optimal states $\vec{o}$, instead of arbitrary states $\vec{y}$. We use this more restricted definition in the proofs of Proposition 7. We define the (*global*) *stretch* as $\tau = \max_{j \in N} \tau_j$. At the conceptual level, the global stretch quantifies how much we lose, in terms of equilibrium efficiency, because agents only have access to an aggregate of the strategies in $\vec{z}$, instead of $\vec{z}$ itself; in short, the stretch is the parameter that measures the lack of information. We will give the formal evidence of this last statement in Lemma 3 and Remark 1.

1.3 Related Work

Our preference game with full information is closely related to both the discrete preference model, introduced in (Chierichetti, Kleinberg, and Oren 2018), and the class of opinion formation games (see e.g., (Bindel, Kleinberg, and Oren 2015; Bhawalkar, Gollapudi, and Munagala 2013; Ghaderi and Srikant 2014)), whose principles go back to the classical works of (Degroot 1974) and (Friedkin and Johnsen 1990) on opinion formation.

In the discrete preference model (Chierichetti, Kleinberg, and Oren 2018), differently from our model, the set of strategies is discrete, d is an exact metric and the influence weights are symmetric and not normalized. The authors show that their model is a potential game, that the price of anarchy can be unbounded and that the price of stability is at most 2. Moreover, the properties of the discrete preference game for richer metrics, such as tree metrics, were studied. Recently, (Lolakapuri et al. 2019) proved that computing a pure Nash equilibrium for such game is PLS-complete, even in a very restricted setting. The model proposed in (Chierichetti, Kleinberg, and Oren 2018) has been generalized in (Auletta et al. 2016) and consistency between preferred strategies and equilibrium strategies were considered in (Auletta et al. 2017).

On the other hand, opinion formation games are characterized by the fact that the set of strategies is the set of points

on the real line, d is L_2^2 (i.e., the square of the L_2 norm, motivated by repeated averaging in (Degroot 1974; Friedkin and Johnsen 1990)) and the influence weights are usually symmetric and not normalized. In the last few years, there has been considerable interest in equilibrium properties (e.g., existence, computational complexity, convergence, price of anarchy and stability) of opinion formation (see e.g., (Bindel, Kleinberg, and Oren 2015; Bhawalkar, Gollapudi, and Munagala 2013; Ghaderi and Srikant 2014)). The main message is that these games are well-behaved due to their continuous strategy space. Specifically, they admit a unique equilibrium which can be computed efficiently and has small price of anarchy for symmetric weights.

Our k -approval voting model bears some resemblance to opinion formation with binary strategies (Ferraioli, Goldberg, and Ventre 2016). Differently from our model, their model is restricted only to binary strategies, the influence weights are symmetric and not normalized. They proved that the price of anarchy is unbounded for $\alpha \leq 1/2$.

The partial information setting has strong connections with the class of aggregative games, e.g., (Kukushkin 2004; Dindos and Mezzetti 2006; Jensen 2010; Acemoglu and Jensen 2013; Babichenko 2017). In this class of games the cost of each agent depends on the his own strategy and an aggregate of the strategies of all the agents.

Aggregating preferences under some metric function has been received attention in algorithms (see e.g., (Ailon, Charikar, and Newman 2008)) and in social choice (see e.g., (Anshelevich et al. 2018)). Ours is the first work where aggregation under metric dissimilarity functions is used in modelling preference games.

1.4 Our Contribution

Our conceptual contribution is the *new model of preference game which encompasses the full and the partial information settings*. While the full information model is mostly inspired by the discrete preference game (Chierichetti, Kleinberg, and Oren 2018), the partial information one introduces a novel feature in preference games, that is the aggregation function. Both settings are quite general and allow for a systematic study of the fundamental questions related to the existence and efficiency of (approximate) pure Nash equilibria in the strategic interaction of agents in social networks.

On the technical side, we provide a comprehensive set of very general results on the existence and the structure of equilibria, and on the price of anarchy of the game. *The first general message of these results is that the two settings, despite being very different from each other, behave very similarly, thus confirming our conjecture that in general the lack of information does not affect the existence and efficiency of the (approximate) equilibria. The second message is that our bounds on the price of anarchy are very general and are derived in terms of key features of the game (i.e., social impact, boundary and stretch). The third message is that low self-confidence levels (i.e., $\alpha \leq 1/2$) help with the existence of equilibria and simplify their structure, while high self-confidence levels (i.e., $\alpha > 1/2$) help with the price of anarchy.*

Specifically, in Section 2 we show that, unless $\alpha > 1/2$ and the game is restricted, there always exists a ρ -approximate equilibrium both for the full and the partial information case. In fact we observe that the two settings behave almost similarly. Specifically, if the agents are compliant ($\alpha < 1/2$), any *consensus* (i.e., a state where all agents adopt the same strategy) is a ρ -equilibrium both for the full information case (Theorem 1) and the partial information one (Theorem 2); if the agents are stubborn ($\alpha > 1/2$), we show that (when the game is unrestricted), in both settings, the state $\vec{s}$ in which each agent i adopts her preferred strategy s_i is a ρ -equilibrium, and moreover such equilibrium is unique when $\rho = 1$ (Theorem 3); lastly, if the agents are balanced ($\alpha = 1/2$), the set of ρ -equilibria in both settings is seemingly larger, as both the set of consensuses and $\vec{s}$ are ρ -equilibria (theorems 1, 2 and 3). Existence of approximate equilibria for the restricted case with $\alpha > 1/2$ requires more assumptions (e.g., on the strategy space, the aggregation function, or the distance function) and is an interesting direction for further research.

In Section 3 we consider the price of anarchy wrt. SUM and MAX. We express the bounds on the price of anarchy wrt. social impact, the boundary and stretch of the game. Interestingly, also the bounds on the price of anarchy hold for both the full and the partial information settings, thus proving that the lack of information does not have any impact also on the inefficiency of the equilibria. Specifically, we observe that for compliant and balanced agents ($\alpha \in (0, 1/2]$), the price of anarchy can be either unbounded for the restricted case (Proposition 1) or $+\infty$ for the unrestricted one (Proposition 2); these hold for both objectives. So, if self-confidence level is low, the price of anarchy of our preference game behaves similarly to that of the discrete preference game (Chierichetti, Kleinberg, and Oren 2018) and of the opinion formation game with binary strategies (Ferraioli, Goldberg, and Ventre 2016). On the other hand, if the agents are stubborn ($\alpha > 1/2$), we show that the price of anarchy is bounded from above by $\rho_z \left(1 + \delta \frac{(1-\alpha)}{\alpha} \frac{\tau}{\beta}\right)$, wrt. SUM (Theorem 4), and by $\rho_z \left(1 + \frac{(1-\alpha)}{\alpha} \frac{\tau}{\beta}\right)$, wrt. MAX (Theorem 5). We also show that these bounds hold with $\tau = 1$ for the full information setting (Lemma 3 and Remark 1), thus proving that τ is the parameter which quantifies the lack of information. Notice that, in order to provide bounds on the price of anarchy only as functions of α , δ and ρ_z (notice that δ is always upper bounded by $n - 1$), we need additional assumptions on the structure of the game which allow us to estimate the values of τ (for the partial information case) and β ; we illustrate several examples in Section 4.

In Section 4, we provide upper bounds on τ and β for several significant cases. Interestingly, our bounds on β depend only on α , while our bounds on τ may depend on the aggregation function, the metric space and the influence weights.

In Section 5, we focus on the k -approval voting game. As we said in the introduction, this game is a spacial case of preference game, it is motivated by the k -approval voting and generalizes the opinion formation model with binary strategies (Ferraioli, Goldberg, and Ventre 2016). Since d

(defined as L_2^2) is a $(2, 1)$ -approximate metric (i.e., $\rho = 2$ and $\rho_Z = 1$, given that it is equivalent to the Hamming distance when restricted to Z), the results from Section 3 carry over to this special case. The main technical result of Section 5 is an upper bound on β (Theorem 6) which, combined with the results in Section 3, implies interesting upper bounds of $1 + \delta(2\alpha - 1)^{-2}\tau$ and $1 + (2\alpha - 1)^{-2}\tau$ on the price of anarchy wrt. SUM and MAX respectively, when $\alpha > 1/2$. We recall that these bounds, as for the general case, hold with $\tau = 1$ for the full information setting. Our analysis complements the results of (Ferraioli, Goldberg, and Ventre 2016) for the opinion formation model with binary strategies. For the partial information case we can derive specific bounds by applying the bounds on τ illustrated in Section 4. One of the most intriguing directions for further research is to determine under which assumptions the k -approval voting game admits equilibria for $\alpha > 1/2$.

2 Existence of Approximate Equilibria

In this section we characterize the existence of equilibria.

Compliant and Balanced Agents.

Theorem 1. *In the full information setting, if $\alpha \leq 1/2$ then any consensus $\vec{e} \in Z^n$ is a ρ -equilibrium.*

Lemma 1. *In the partial information setting, if $\alpha \leq 1/2$ then $\text{aggr}_i(\vec{x}_{-i})$ is a ρ -approximate best response for agent i with respect to $\vec{x}_{-i}$, for every partial strategy profile $\vec{x}_{-i}$. Moreover, if $\alpha < 1/2$ and $\rho = 1$, $\text{aggr}_i(\vec{x}_{-i})$ is the unique best response.*

The following theorem is an immediate consequences of Lemma 1.

Theorem 2. *In the partial information setting, if $\alpha \leq 1/2$, then any consensus $\vec{e} \in Z^n$ is a ρ -equilibrium.*

Stubborn and Balanced Agents.

Lemma 2. *If $\alpha \geq 1/2$ then s_i is a ρ -approximate dominant strategy for every agent $i \in N_S$. Moreover, if $\alpha > 1/2$ and $\rho = 1$, s_i is the unique dominant strategy.*

Theorems 3 is an immediate consequence of Lemma 2.

Theorem 3. *If $\alpha \geq 1/2$ and the game is unrestricted ($N_S = N$) then $\vec{s} = (s_1, s_2, \dots, s_n)$ is a ρ -approximate equilibrium. Moreover, if $\rho = 1$ then $\vec{s}$ is the unique equilibrium.*

3 Price of Anarchy

Compliant and Balanced Agents. This is the case in which the social cost of an equilibrium can be unboundedly larger than the social cost of any optimal configuration.

Proposition 1. *For $\alpha \in (0, 1/2]$, if the game is restricted ($N_S = \emptyset$), then there exist instances for which both POA^{SUM} and POA^{MAX} are unbounded.*

The proof of Proposition 1 considers a strategy universe $U = \{a, b, s\}$ with $Z = \{a, b\}$ and $s_i = s$ for all agents $i \in N$. If $d(b, s) = \epsilon$, for some arbitrarily small $\epsilon > 0$, $d(a, b) = 1$ and $d(a, s) = 1 + \epsilon$, d is a metric. Then, the consensus

on b is optimal, while the consensus on a is an equilibrium. If, instead, we consider $Z = U = \{a, s\}$, with $d(a, s) = 1$ and $s_i = s$ for all agents i , the state $\vec{s}$, in which every agent chooses s , is optimal with $\text{SUM}(\vec{s}) = \text{MAX}(\vec{s}) = 0$, while the consensus on a is an equilibrium. Hence, we show that:

Proposition 2. *For $\alpha \in (0, 1/2]$, if the game is unrestricted ($N_S = N$), then there exist instances with for which both POA^{SUM} and POA^{MAX} are unbounded.*

Stubborn Agents. In theorems 4 and 5, we show general bounds on the price of anarchy that depend on δ , τ , β and α . The proof of Lemma 3 follows from the equilibrium condition, the triangle inequality and the definition of stretch.

Lemma 3. *For every agent i , equilibrium $\vec{e}$ and state $\vec{z}$,*

$$c_i^F(\vec{e}) \leq \rho_Z c_i^F(\vec{z}) + \rho_Z(1 - \alpha)\pi_i(\vec{z}, \vec{e}),$$

and

$$c_i^P(\vec{e}) \leq \rho_Z c_i^P(\vec{z}) + \tau_i \rho_Z(1 - \alpha)\pi_i(\vec{z}, \vec{e}).$$

Proof. For the full information setting, using that $\vec{e}$ is an equilibrium, we have

$$\begin{aligned} c_i^F(\vec{e}) &\leq \alpha d(\vec{z}(i), s_i) + (1 - \alpha) \sum_{j \neq i} w_{ij} d(\vec{z}(i), \vec{e}(j)) \\ &\leq \alpha d(\vec{z}(i), s_i) + \rho_Z(1 - \alpha) \sum_{j \neq i} w_{ij} d(\vec{z}(i), \vec{z}(j)) \\ &\quad + \rho_Z(1 - \alpha) \sum_{j \neq i} w_{ij} d(\vec{z}(j), \vec{e}(j)) \\ &\leq \rho_Z c_i^F(\vec{z}) + \rho_Z(1 - \alpha)\pi_i(\vec{z}, \vec{e}), \end{aligned}$$

where the first inequality follows from the equilibrium condition, the second from the approximate triangle inequality and the third from the definition of the stretch, $c_i^F(\vec{z})$ and ρ_Z .

Analogously, in the partial information setting we have

$$\begin{aligned} c_i^P(\vec{e}) &\leq \alpha d(\vec{z}(i), s_i) + (1 - \alpha) d(\vec{z}(i), \text{aggr}_i(\vec{e}_{-i})) \\ &\leq \alpha d(\vec{z}(i), s_i) + \rho_Z(1 - \alpha) d(\vec{z}(i), \text{aggr}_i(\vec{z}_{-i})) \\ &\quad + \rho_Z(1 - \alpha) d(\text{aggr}_i(\vec{z}_{-i}), \text{aggr}_i(\vec{e}_{-i})) \\ &\leq \rho_Z c_i^P(\vec{z}) + \rho_Z(1 - \alpha) d(\text{aggr}_i(\vec{z}_{-i}), \text{aggr}_i(\vec{e}_{-i})) \\ &\leq \rho_Z c_i^P(\vec{z}) + \tau_i \rho_Z(1 - \alpha)\pi_i(\vec{z}, \vec{e}). \end{aligned}$$

Also in this case the first inequality follows from the equilibrium condition, the second from the approximate triangle inequality, the third from the definition of $c_i^P(\vec{z})$ and ρ_Z and the last from the definition of the stretch. $\square$

Lemma 4. *If $\alpha > 1/2$, for every equilibrium $\vec{e}$ and optimal state $\vec{o} \in O_{\text{SUM}}$ with $D(\vec{o}, \vec{e}) \neq \emptyset$, we have $D(\vec{o}, \vec{e}) \subseteq D(\vec{o}, \vec{s})$, where $\vec{s} = (s_1, s_2, \dots, s_n)$.*

Remark 1. *Since by definition $\tau_i \geq 1$, for simplicity of exposition, in theorems 4 and 5 we relax the claim of Lemma 3 and assume $c_i^F(\vec{e}) \leq \rho_Z c_i^F(\vec{z}) + \tau_i \rho_Z(1 - \alpha)\pi_i(\vec{z}, \vec{e})$, for every agent i , equilibrium $\vec{e}$ and state $\vec{z}$. This allows us to apply the same arguments without making a distinction between the full and partial information. This implies that the bounds on the price of anarchy of theorems 4 and 5 hold with $\tau = 1$ for the full information setting. This is the evidence of the fact that τ is the parameter that measure the lack of information.*

Theorem 4. If $\alpha > 1/2$, $\text{PoA}^{\text{SUM}} \leq \rho_Z \left(1 + \delta \frac{(1-\alpha)}{\alpha} \frac{\tau}{\beta}\right)$.

Proof. If $E \subseteq O_{\text{SUM}}$, then $\text{PoA}^{\text{SUM}} = 1$. Otherwise, let $\vec{e}$ be any equilibrium and $\vec{o} \in O_{\text{SUM}}$ be any optimal state with $D(\vec{o}, \vec{e}) \neq \emptyset$. We have

$$\begin{aligned} \text{SUM}(\vec{e}) &\leq \rho_Z \text{SUM}(\vec{o}) + \rho_Z(1-\alpha) \sum_{i \in N} \tau_i \pi_i(\vec{o}, \vec{e}) \\ &\leq \rho_Z \text{SUM}(\vec{o}) + \tau \rho_Z(1-\alpha) \sum_{i \in N} \sum_{j \in N \setminus \{i\}} w_{ij} d(\vec{o}(j), \vec{e}(j)) \\ &= \rho_Z \text{SUM}(\vec{o}) + \tau \rho_Z(1-\alpha) \sum_{j \in D(\vec{o}, \vec{e})} \delta_j d(\vec{o}(j), \vec{e}(j)), \end{aligned} \quad (6)$$

where the first inequality follows from Lemma 3 and the definition of SUM, the second from the definition of τ and π_i , and the last step from the definitions of $D(\vec{o}, \vec{e})$ and the social impact of j .

On the other hand,

$$\begin{aligned} \text{SUM}(\vec{o}) &\geq \alpha \sum_{i \in N} d(\vec{o}(i), s_i) = \alpha \sum_{i \in D(\vec{o}, \vec{s})} d(\vec{o}(i), s_i) \\ &\geq \alpha \sum_{i \in D(\vec{o}, \vec{e})} d(\vec{o}(i), s_i) \geq \alpha \sum_{i \in D(\vec{o}, \vec{e})} \beta_i d(\vec{o}(i), \vec{e}(i)), \end{aligned} \quad (7)$$

where the first inequality follows from Lemma 4 and the second from the definition of boundary (5). Notice that the last expression is strictly larger than 0 because $\alpha > 1/2$, $\beta_i > 0$ and $d(\vec{o}(i), \vec{e}(i)) > 0$ for every $i \in D(\vec{o}, \vec{e}) \neq \emptyset$.

Therefore, we can conclude that

$$\begin{aligned} \text{PoA}^{\text{SUM}} &\leq \frac{\rho_Z \text{SUM}(\vec{o}) + \tau \rho_Z(1-\alpha) \sum_{j \in D(\vec{o}, \vec{e})} \delta_j d(\vec{o}(j), \vec{e}(j))}{\text{SUM}(\vec{o})} \\ &\leq \rho_Z + \frac{\tau \rho_Z(1-\alpha) \sum_{j \in D(\vec{o}, \vec{e})} \delta_j d(\vec{o}(j), \vec{e}(j))}{\alpha \sum_{i \in D(\vec{o}, \vec{e})} \beta_i d(\vec{o}(i), \vec{e}(i))} \\ &\leq \rho_Z \left(1 + \delta \frac{(1-\alpha)}{\alpha} \frac{\tau}{\beta}\right), \end{aligned}$$

where the first inequality follows from (6), the second from (7), and the third from the definitions of δ , β . $\square$

Theorem 5. If $\alpha > 1/2$, $\text{PoA}^{\text{MAX}} \leq \rho_Z \left(1 + \frac{(1-\alpha)}{\alpha} \frac{\tau}{\beta}\right)$.

The proof of Theorem 5 is similar to that of Theorem 4.

4 Some Bounds on the Boundary and Stretch

Next, we show bounds on the boundary and the stretch under quite general assumptions on the structure of the game.

We start with a technical claim which will be used to bound β .

Proposition 3. If $\rho_Z = 1$ then for any equilibrium $\vec{e} \in E$, any agent i , and any strategy $x \neq \vec{e}(i)$, we have

$$d(x, s_i) \geq d(\vec{e}(i), s_i) - \frac{(1-\alpha)}{\alpha} d(x, \vec{e}(i)).$$

Proposition 4. If $\alpha > 1/2$ and $\rho = 1$ then $\beta \geq \frac{2\alpha-1}{2\alpha}$.

Proof. For any equilibrium $\vec{e} \in E$, any agent i , and any strategy $x \neq \vec{e}(i)$, we have

$$\begin{aligned} d(x, s_i) &\geq d(\vec{e}(i), s_i) - \frac{(1-\alpha)}{\alpha} d(x, \vec{e}(i)) \\ &\geq \left[d(x, \vec{e}(i)) - d(x, s_i) \right] - \frac{(1-\alpha)}{\alpha} d(x, \vec{e}(i)) \\ &\geq \left(\frac{2\alpha-1}{\alpha} \right) d(x, \vec{e}(i)) - d(x, s_i), \end{aligned}$$

where the first inequality follows from Proposition 3 and the remaining inequalities from the triangle inequality. The theorem follows by adding $d(x, s_i)$ to each side of the previous inequality and dividing by 2. $\square$

Proposition 5. For $\alpha > 1/2$, if either the game is unrestricted ($N_S = N$) or d is uniform then $\beta = 1$.

We proceed to show some upper bounds on the stretch. Since $\tau = 1$ in the full information case, we next restrict our attention on the partial information case.

Proposition 6. It always holds that $\tau \leq \frac{d^{\max}(Z)}{w^{\min} d^{\min}(Z)}$, where $d^{\max}(Z) = \max_{x,y \in Z} d(x,y)$ is the diameter of Z , $d^{\min}(Z) = \min_{x \neq y: d(x,y) > 0} d(x,y)$ is the minimum positive distance in Z , and $w^{\min} = \min_{i \neq j: w_{ij} > 0} w_{ij}$ is the minimum positive influence weight.

Proof. In the definition of τ_i , we have $\tau_i = 1$, if $\pi_i(\vec{e}, \vec{y}) = d(\text{aggr}_i(\vec{e}_{-i}), \text{aggr}_i(\vec{y}_{-i})) = 0$. Hence, we assume that $\pi_i(\vec{e}, \vec{y}) > 0$.

We let $w^{\min}(i) = \min_{j \neq i: w_{ij} > 0} w_{ij}$. Then, $\pi_i(\vec{e}, \vec{y}) = \sum_{j \neq i} w_{ij} d(\vec{e}(j), \vec{y}(j)) \geq w^{\min}(i) d^{\min}(Z)$, because $\vec{e}(j), \vec{y}(j) \in Z$, there is at least one positive term in the sum, and if either $w_{ij} = 0$ or $d(\vec{e}(j), \vec{y}(j)) = 0$, the corresponding term is 0. Moreover, $d(\text{aggr}_i(\vec{e}_{-i}), \text{aggr}_i(\vec{y}_{-i})) \leq d^{\max}(Z)$, because $\text{aggr}_i(\vec{e}_{-i}), \text{aggr}_i(\vec{y}_{-i}) \in Z$. Therefore,

$$\tau_i \leq \frac{d^{\max}(Z)}{w^{\min}(i) d^{\min}(Z)}.$$

Using that the global stretch $\tau = \max_{i \in N} \tau_i$ and that $w^{\min} = \min_{i \in N} w^{\min}(i)$, we conclude the proof of the proposition. $\square$

Unless we impose additional structures, the upper bound of Proposition 6 is essentially the best possible, because there are examples where a small change in a single coordinate of a state moves the aggregate to a diametrically different strategy. E.g., consider $Z = \{0, 1\}$, $\vec{e}_{-i} = (0, 1, 1)$, $\vec{y}_{-i} = (0, 0, 1)$, $w_{i1} = w_{i3} = \frac{1-\varepsilon}{2}$ and $w_{i2} = \varepsilon$, and the Fréchet median as aggregation function. Clearly, $\text{aggr}_i(\vec{e}_{-i}) = 1$, while $\text{aggr}_i(\vec{y}_{-i}) = 0$.

For the following propositions, we restrict the definition of τ to optimal states $\vec{o}$ (either $\vec{o} \in O_{\text{SUM}}$ or $\vec{o} \in O_{\text{MAX}}$), instead of arbitrary states $\vec{y}$.

Proposition 7. If either every optimum (wrt. any objective) is a consensus or every equilibrium is a consensus then

- a) $\tau \leq 2$, if $\rho_Z = 1$ and aggr is the Fréchet median;
- b) $\tau = 1$, if d is L_2^2 and aggr is the Fréchet median.

5 k -Approval Voting Game

In this section we discuss how the results developed in the previous sections can be used, together with other technical considerations, to characterize the equilibria and their efficiency in the k -approval voting game.

Existence of approximate equilibria. The approval voting model does not admit an exact potential function (due to the asymmetry of w_{ij} , even for binary strategies). Instead, since d is a $(2, 1)$ -approximate metric, Theorem 1, Theorem 2 and Theorem 3 apply to this setting, implying the existence of 2-approximate pure Nash equilibria in both the full and partial information settings. Moreover, by using (Ferraioli, Goldberg, and Ventre 2016, Observation 2.2), we can show that our approval voting game admits an equilibrium for the game with binary strategies when $\alpha \in [0, 1]$.

Price of anarchy. In order to bound the price of anarchy for the k -approval voting game, we first derive a lower bound on β .

Theorem 6. *For every integers k, m , with $0 < k < m$, if $\alpha > 1/2$ then $\beta \geq \left(\frac{2\alpha-1}{2\alpha}\right)^2$ for the k -approval voting game.*

Proof. To obtain a lower bound on β_i , we fix an equilibrium $\vec{e} \in E$ and a strategy $x \in Z$ with $\vec{e}(i) \neq x$, and find $s_i^* = \arg \min_{s_i \in U} d(x, s_i)/d(x, \vec{e}(i))$. Then, we conclude that $\beta_i \geq d(x, s_i^*)/d(x, \vec{e}(i))$. Throughout the proof, we let $[m] = \{1, \dots, m\}$. We also let $x(j)$, $s_i(j)$ and $\vec{e}(i)(j)$ denote the j -th coordinate (bit) of strategies x , s_i and $\vec{e}(i)$.

To this end, for any fixed strategy $x \in Z$, let $C_0 = \{j \in [m] : x(j) = 0\}$ and let $C_1 = \{j \in [m] : x(j) = 1\}$ be indices j for which $x(j)$ is 0, for C_0 , and 1, for C_1 . We recall that $|C_1| = k$ and $|C_0| = m - k$. Moreover, for the equilibrium strategy $\vec{e}(i) \in Z$, we let $C_{01} = \{j \in C_0 : \vec{e}(i)(j) = 1\}$ and $C_{10} = \{j \in C_1 : \vec{e}(i)(j) = 0\}$ be indices j where $x(j) \neq \vec{e}(i)(j)$. Since both x and $\vec{e}(i)$ include exactly k ones and $m - k$ zeros, $|C_{01}| = |C_{10}|$. In the following, we denote $|C_{01}| = |C_{10}| = \ell$, for simplicity, for some $1 \leq \ell \leq k$. Using this notation, we can write that $d(x, \vec{e}(i)) = 2\ell$.

Our goal is to find $\min_{s_i \in U} d(x, s_i)/(2\ell)$, for any fixed $\ell \in \{1, \dots, k\}$, under the constraints that $s_i \in U$, $x \in Z$, and $\vec{e}(i) \in Z$ is an equilibrium strategy. We first observe that

$$\frac{d(x, s_i)}{2\ell} = \frac{1}{2\ell} \left(\sum_{j \in C_0} s_i(j)^2 + \sum_{j \in C_1} (1 - s_i(j))^2 \right), \quad (8)$$

$$\text{where } \sum_{j \in C_0 \cup C_1} s_i(j) = k \text{ and } s_i \in [0, 1]^m. \quad (9)$$

Since $\vec{e}$ is an equilibrium and d is an exact metric on Z (d is identical to the Hamming distance on Z), we can apply Proposition 3 and obtain

$$d(\vec{e}(i), s_i) - d(x, s_i) \leq \frac{(1 - \alpha)}{\alpha} d(x, \vec{e}(i)) = \frac{(1 - \alpha)}{\alpha} 2\ell. \quad (10)$$

Moreover, for any $j \notin C_{10} \cup C_{01}$, $x(j) = \vec{e}(i)(j)$, we have

$$\begin{aligned} d(\vec{e}(i), s_i) - d(x, s_i) &= \sum_{j \in C_{10}} (s_i(j)^2 - (1 - s_i(j))^2) \\ &\quad + \sum_{j \in C_{01}} ((1 - s_i(j))^2 - s_i(j)^2) \\ &= 2 \sum_{j \in C_{10}} s_i(j) - 2 \sum_{j \in C_{01}} s_i(j), \end{aligned}$$

where, for the last equality, we use the fact that $|C_{10}| = |C_{01}|$. Using (10), we conclude that for any equilibrium strategy $\vec{e}(i) \in Z$ and any strategy $x \in Z$ with $d(x, \vec{e}(i)) = 2\ell$, s_i satisfies

$$\sum_{j \in C_{10}} s_i(j) - \sum_{j \in C_{01}} s_i(j) \leq \ell(1 - \alpha)/\alpha. \quad (11)$$

Therefore, we can find $\min_{s_i \in S} d(x, s_i)/(2\ell)$ by computing an s_i that minimizes (8), subject to (9) and (11). We can do so by solving a strictly convex minimization problem with linear constraints.

In the supplementary material, we show that the minimizer s_i^* of (8), subject to (9) and (11), achieves an objective value of $d(x, s_i^*)/(2\ell) = \left(\frac{2\alpha-1}{2\alpha}\right)^2$. Moreover, s_i^* is defined for any strategy $x \in Z$ and any equilibrium strategy $\vec{e}(i) \in Z$ with $d(x, \vec{e}(i)) = 2\ell > 0$, and always attains the same objective value $d(x, s_i^*)/(2\ell) = \left(\frac{2\alpha-1}{2\alpha}\right)^2$. $\square$

The proof of Theorem 6 actually shows something stronger, namely that for any equilibrium strategy $\vec{e}(i) \in Z$ and any $x \in Z$, with $x \neq \vec{e}(i)$, there is a preferred strategy $s_i^* \in S$, defined as in the proof of Theorem 6, such that $d(x, s_i^*)/d(x, \vec{e}(i)) = \left(\frac{2\alpha-1}{2\alpha}\right)^2$. Therefore, the proof of Theorem 6 shows that for any agent i , $\beta_i = \left(\frac{2\alpha-1}{2\alpha}\right)^2$.

Combined with Theorem 4, Theorem 5 and the fact that $\rho_Z = 1$, Theorem 6 implies that, when $\alpha \geq 1/2$, for the k -approval voting game

$$\text{POA}^{\text{SUM}} \leq 1 + \delta \frac{4\alpha(1 - \alpha)}{(2\alpha - 1)^2} \tau \leq 1 + \delta(2\alpha - 1)^{-2} \tau,$$

and

$$\text{POA}^{\text{MAX}} \leq 1 + \delta \frac{4\alpha(1 - \alpha)}{(2\alpha - 1)^2} \tau \leq 1 + (2\alpha - 1)^{-2} \tau.$$

We recall that these bounds hold with $\tau = 1$ for the full information case (see Remark 1). These results complements the results in (Ferraioli, Goldberg, and Ventre 2016) for the binary opinion formation model. By applying Proposition 7, we obtain that these bounds hold with $\tau = 1$ also for the partial information setting when the aggregation function is the Fréchet median and either every optimum (wrt. any objective) or every equilibrium is a consensus.

References

- Acemoglu, D.; and Jensen, M. K. 2013. Aggregate comparative statics. *Games and Economic Behavior*, 81: 27–49.
- Ailon, N.; Charikar, M.; and Newman, A. 2008. Aggregating inconsistent information: Ranking and clustering. *Journal of the ACM*, 55(5).
- Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27–51.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2016. Generalized Discrete Preference Games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 53–59.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2017. Robustness in discrete preference games. In *In Proc. of the 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1314–1321.
- Babichenko, Y. 2017. Fast Convergence of Best-Reply Dynamics in Aggregative Games. *Mathematics of Operations Research*, 43(1): 333–346.
- Bhawalkar, K.; Gollapudi, S.; and Munagala, K. 2013. Co-evolutionary Opinion Formation Games. In *In Proc. of the 45th ACM Symposium on Theory of Computing (STOC)*, 41–50.
- Bindel, D.; Kleinberg, J. M.; and Oren, S. 2015. How bad is forming your own opinion? *Games and Economic Behavior*, 92: 248–265.
- Chierichetti, F.; Kleinberg, J. M.; and Oren, S. 2018. On discrete preferences and coordination. *Journal of Computer and System Sciences*, 93: 11–29.
- Degroot, M. H. 1974. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345): 118–121.
- Dindos, M.; and Mezzetti, C. 2006. Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. *Games and Economic Behavior*, 54(2): 261–292.
- Epitropou, M.; Fotakis, D.; Hoefer, M.; and Skoulakis, S. 2019. Opinion Formation Games with Aggregation and Negative Influence. *Theory Computing Systems*, 63(7): 1531–1553.
- Fanelli, A. 2021. On coevolutionary opinion formation. *Unpublished (personal communication)*.
- Ferraioli, D.; Goldberg, P. W.; and Ventre, C. 2016. Decentralized dynamics for finite opinion games. *Theoretical Computer Science*, 648: 96–115.
- Friedkin, N. E.; and Johnsen, E. C. 1990. Social influence and opinions. *The Journal of Mathematical Sociology*, 15(3-4): 193–206.
- Ghaderi, J.; and Srikant, R. 2014. Opinion Dynamics in Social Networks with Stubborn Agents: Equilibrium and Convergence Rate. *Automatica*, 50: 3209–3215.
- Jensen, M. 2010. Aggregative games and best-reply potentials. *Economic Theory*, 43: 45–66.
- Kukushkin, N. S. 2004. Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48(1): 94–110.
- Lolakapuri, P. R.; Bhaskar, U.; Narayanam, R.; Parija, G. R.; and Dayama, P. S. 2019. Computational Aspects of Equilibria in Discrete Preference Games. In *In Proc. of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 471–477.

Appendix

A Feasibility of Fréchet Mean and Fréchet Median

The following proposition shows that both aggregation functions (4) and (3) are feasible.

Proposition 8. *The Fréchet mean and the Fréchet median are feasible aggregation rules.*

Proof. We prove the statement for the Fréchet mean. A virtually identical argument applies to the Fréchet median.

Let i be any agent, and $\vec{x}$ and $\vec{y}$ be any pair of states. It is straightforward to verify unanimity, namely, that if $\vec{x}_{-i}$ is a consensus on $x \in Z$, then $\text{mean}(\vec{x}_{-i}) = x$. In fact, every term of the summation would be 0 in x .

We proceed to prove consistency. Let us assume that $\sum_{j \neq i} w_{ij} d(\vec{x}(j), \vec{y}(j)) = 0$. If $\vec{x}(j) = \vec{y}(j)$ for all coordinates $j \neq i$, then $\vec{x}_{-i}$ and $\vec{y}_{-i}$ are identical and $\text{mean}_i(\vec{x}_{-i}) = \text{mean}_i(\vec{y}_{-i})$. So, let us assume that for some coordinates j , $\vec{x}(j) \neq \vec{y}(j)$. Since $\sum_{j \neq i} w_{ij} d(\vec{x}(j), \vec{y}(j)) = 0$, it must be $w_{ij} = 0$ for all coordinates j with $\vec{x}(j) \neq \vec{y}(j)$. Equivalently, for every j with $w_{ij} > 0$, we have that $\vec{x}(j) = \vec{y}(j)$. Therefore, for every $y \in Z$, it holds that

$$\begin{aligned} \sum_{j \in N \setminus \{i\}} w_{ij} d^2(y, \vec{x}(j)) &= \sum_{\substack{j \in N \setminus \{i\} \\ w_{ij} > 0}} w_{ij} d^2(y, \vec{x}(j)) \\ &= \sum_{\substack{j \in N \setminus \{i\} \\ w_{ij} > 0}} w_{ij} d^2(y, \vec{y}(j)) \\ &= \sum_{j \in N \setminus \{i\}} w_{ij} d^2(y, \vec{y}(j)), \end{aligned}$$

which implies that $\text{mean}_i(\vec{x}_{-i}) = \text{mean}_i(\vec{y}_{-i})$. $\square$

B Missing Proofs

Proof of Theorem 1

Proof. Let i be any agent, $\vec{e}$ any consensus, and $x \neq \vec{e}(i)$ any strategy. Using the facts that $d(\vec{e}(i), \vec{e}(j)) = 0$, we obtain that:

$$\begin{aligned} c_i^F(\vec{e}) &= \alpha d(s_i, \vec{e}(i)) \\ &\leq \rho \left(\alpha d(s_i, x) + \alpha d(x, \vec{e}(i)) \right) \end{aligned} \quad (12)$$

$$\leq \rho \alpha d(s_i, x) + \rho(1 - \alpha) d(x, \vec{e}(i)) \quad (13)$$

$$= \rho \alpha d(s_i, x) + \rho(1 - \alpha) \sum_{j \neq i} w_{ij} d(x, \vec{e}(j)) \quad (14)$$

$$= \rho \alpha d(s_i, x) + \rho(1 - \alpha) \sum_{j \neq i} w_{ij} d(x, \vec{e}(j)) \quad (15)$$

$$= \rho c_i^F([\vec{e}_{-i}, x]),$$

where (12) follows from the approximate triangle inequality, (13) from $\alpha \leq 1/2$, (14) from $\sum_{j \neq i} w_{ij} = 1$, and (15) from $\vec{e}(i) = \vec{e}(j)$. $\square$

Proof of Lemma 1

Proof. Let i be any agent, $\vec{x}$ any state, and $x \neq \text{aggr}_i(\vec{x}_{-i})$ any strategy. Using the facts that $\text{aggr}_i(\vec{x}_{-i}) \in Z$ and $d(\text{aggr}_i(\vec{x}_{-i}), \text{aggr}_i(\vec{x}_{-i})) = 0$, we obtain that:

$$\begin{aligned} c_i^P([\vec{x}_{-i}, \text{aggr}_i(\vec{x}_{-i})]) &= \alpha d(\text{aggr}_i(\vec{x}_{-i}), s_i) \\ &\leq \rho \left(\alpha d(x, s_i) + \alpha d(x, \text{aggr}_i(\vec{x}_{-i})) \right) \end{aligned} \quad (16)$$

$$\begin{aligned} &\leq \rho \alpha d(x, s_i) + \rho(1 - \alpha) d(x, \text{aggr}_i(\vec{x}_{-i})) \\ &= \rho c_i^P([\vec{x}_{-i}, x]), \end{aligned} \quad (17)$$

where (16) follows from the approximate triangle inequality and (17) from $\alpha \leq 1/2$. Notice that if $\alpha < 1/2$ then (17) is strict and hence, when $\rho = 1$, $\text{aggr}_i(\vec{x}_{-i})$ is the unique best response. $\square$

Proof of Lemma 2

Proof. Let us assume that $\alpha \geq 1/2$. Let i be any agent in N_S and let $\vec{x}$ be any state with $\vec{x}(i) \neq s_i$ (recall that $|Z| \geq 2$).

Since $s_i \in Z$, in the full information setting we have

$$\begin{aligned} c_i^F([\vec{x}_{-i}, s_i]) &= \alpha d(s_i, s_i) + (1 - \alpha) \sum_{j \neq i} w_{ij} d(s_i, \vec{x}(j)) \\ &= (1 - \alpha) \sum_{j \neq i} w_{ij} d(s_i, \vec{x}(j)) \\ &\leq \rho(1 - \alpha) \sum_{j \neq i} w_{ij} \left(d(s_i, \vec{x}(i)) + d(\vec{x}(i), \vec{x}(j)) \right) \\ &\leq \rho \alpha d(s_i, \vec{x}(i)) + \rho(1 - \alpha) \sum_{j \neq i} w_{ij} d(\vec{x}(i), \vec{x}(j)) \\ &= \rho c_i^F(\vec{x}), \end{aligned}$$

where the first inequality follows from the approximate triangle inequality and the second from $\sum_{j \neq i} w_{ij} = 1$ and $\alpha \geq 1/2$. Notice that if $\alpha < 1/2$, then the second inequality is strict and hence, when $\rho = 1$, s_i is the unique dominant strategy.

Analogously, in the partial information setting we have

$$\begin{aligned} c_i^P([\vec{x}_{-i}, s_i]) &= \alpha d(s_i, s_i) + (1 - \alpha) d(s_i, \text{aggr}_i(\vec{x}_{-i})) \\ &= (1 - \alpha) d(s_i, \text{aggr}_i(\vec{x}_{-i})) \\ &\leq \rho(1 - \alpha) \left(d(x_i, s_i) + d(x_i, \text{aggr}_i(\vec{x}_{-i})) \right) \\ &\leq \rho \alpha d(x_i, s_i) + \rho(1 - \alpha) d(x_i, \text{aggr}_i(\vec{x}_{-i})) \\ &= \rho c_i^P(\vec{x}). \end{aligned}$$

Also here, the first inequality follows from the approximate triangle inequality and the second from $\alpha \geq 1/2$. If $\alpha < 1/2$, then the second inequality is strict, hence, s_i is a strictly dominant strategy of agent i . $\square$

Proof of Proposition 1

Proof. We consider a set of instances with $U = \{a, b, s\}$, $Z = \{a, b\}$ and $S = \{s\}$, (i.e., $s_i = s$, for all $i \in N$), where a, b, s are three distinct elements. Notice that $N_S = \emptyset$. We

recall that $0 < \alpha \leq 1/2$. Since a, b, s are three distinct elements, the distance between every two elements of U is non-negative. Let the distance between b and s be an arbitrarily small positive number, i.e., $d(b, s) = \epsilon > 0$. By Theorem 2, the state $\vec{e}$ in which every agent choses a is an equilibrium; hence, $\text{SUM}(\vec{e}) = n\alpha d(a, s)$ and $\text{MAX}(\vec{e}) = \alpha d(a, s)$. On the other hand, in the state $\vec{o}$, in which every agent choses b , $\text{SUM}(\vec{o}) = n\alpha d(b, s) = n\alpha\epsilon > 0$ and $\text{MAX}(\vec{o}) = \alpha d(b, s) = \alpha\epsilon > 0$. We can conclude that $\text{POA}^{\text{SUM}} = \text{POA}^{\text{MAX}} \geq \frac{d(a, s)}{\epsilon}$. Since ϵ is an arbitrarily small positive number, the claim follows. $\square$

Proof of Lemma 4

Proof. Let i be any agent in $D(\vec{o}, \vec{e})$. We prove the claim by showing that i belongs also to $D(\vec{o}, \vec{s})$. If $i \in N_S$ then, by Theorem 3, $\vec{e}(i) = s_i$. Since $\vec{o}(i) \neq \vec{e}(i)$, we also have that $\vec{o}(i) \neq s_i$, i.e., $i \in D(\vec{o}, \vec{s})$. If $i \in N \setminus N_S$ then s_i does not belongs to Z . Since $\vec{o}(i)$ belongs to Z , this trivially implies that $\vec{o}(i) \neq s_i$, i.e., $i \in D(\vec{o}, \vec{s})$. $\square$

Proof of Theorem 5

Proof. Let us consider the full information setting first.

If $E \subseteq O_{\text{MAX}}$, then trivially $\text{POA}^{\text{MAX}} = 1$.

Otherwise, let $\vec{e} \in E$ and $\vec{o} \in O_{\text{MAX}}$ be two states such that $D(\vec{o}, \vec{e}) \neq \emptyset$. We recall that $\alpha > 1/2$. Let i be one of the agents with maximum cost at equilibrium, i.e., $c_i^F(\vec{e}) = \text{MAX}(\vec{e})$. We have

$$\begin{aligned} \text{MAX}(\vec{e}) &\leq \rho_Z \text{MAX}(\vec{o}) + \tau_i \rho_Z (1 - \alpha) \pi_i(\vec{o}, \vec{e}) \\ &\leq \rho_Z \text{MAX}(\vec{o}) + \tau \rho_Z (1 - \alpha) \sum_{j \in N \setminus \{i\}} w_{ij} d(\vec{o}(j), \vec{e}(j)), \end{aligned} \quad (18)$$

where the first inequality follows from Lemma 3 and the second from the definition of τ and π_i . Notice that, if none of the agents in $N \setminus \{i\}$ is in $D(\vec{o}, \vec{e})$, then every term in the summation in (18) is 0 and the inequality becomes $\text{MAX}(\vec{e}) \leq \text{MAX}(\vec{o})$, which implies $\text{POA}^{\text{MAX}} = 1$.

Therefore, let us assume that $\{N \setminus \{i\}\} \cap D(\vec{o}, \vec{e}) \neq \emptyset$ and let j^* be any agent in such set maximizing the distance between the strategy at the optimum and the one at the equilibrium, i.e., $j^* \in \arg \max_{j \in D(\vec{o}, \vec{e})} d(\vec{o}(j), \vec{e}(j))$. Then we can write

$$\begin{aligned} \text{MAX}(\vec{e}) &\leq \rho_Z \text{MAX}(\vec{o}) + \tau \rho_Z (1 - \alpha) \sum_{j \in N \setminus \{i\}} w_{ij} d(\vec{o}(j^*), \vec{e}(j^*)) \\ &= \rho_Z \text{MAX}(\vec{o}) + \tau \rho_Z (1 - \alpha) d(\vec{o}(j^*), \vec{e}(j^*)), \end{aligned} \quad (19)$$

where first step follows from (18) and the definition of j^* , and the second from $\sum_{j \in N} w_{ij} = 1$.

On the other side,

$$\begin{aligned} \text{MAX}(\vec{o}) &\geq c_{j^*}^F(\vec{o}) \geq \alpha d(\vec{o}(j^*), s_{j^*}) \\ &\geq \alpha \beta_i \cdot d(\vec{o}(j^*), \vec{e}(j^*)), \end{aligned} \quad (20)$$

where the last inequality follows from the definition of boundary (5). Notice that the last expression is strictly larger

than 0, because $\alpha > 1/2$, $\beta_i > 0$ and $d(\vec{o}(j^*), \vec{e}(j^*)) > 0$. Therefore, we can conclude that

$$\begin{aligned} \text{POA}^{\text{MAX}} &\leq \frac{\rho_Z \text{MAX}(\vec{o}) + \tau \rho_Z (1 - \alpha) d(\vec{o}(j^*), \vec{e}(j^*))}{\text{MAX}(\vec{o})} \\ &\leq \rho_Z + \frac{\tau \rho_Z (1 - \alpha) d(\vec{o}(j^*), \vec{e}(j^*))}{\alpha \beta_i \cdot d(\vec{o}(j^*), \vec{e}(j^*))} \\ &\leq \rho_Z \left(1 + \frac{(1 - \alpha) \tau}{\alpha \beta_i}\right), \end{aligned}$$

where the first inequality follows from (19), and the last from (20) and the definition of β .

For the partial information setting, the same argument applies by replacing c_i^F with c_i^P . $\square$

Proof of Proposition 3

Proof. Let us first focus on the full information setting. We apply the equilibrium condition and obtain that

$$\begin{aligned} \alpha d(\vec{e}(i), s_i) + (1 - \alpha) \sum_{j \neq i} w_{ij} d(\vec{e}(i), \vec{e}(j)) \\ \leq \alpha d(x, s_i) + (1 - \alpha) \sum_{j \neq i} w_{ij} d(x, \vec{e}(j)), \end{aligned}$$

which, by triangle inequality on Z , implies

$$\begin{aligned} d(x, s_i) &\geq d(\vec{e}(i), s_i) - \frac{(1 - \alpha)}{\alpha} \left[\sum_{j \neq i} w_{ij} d(x, \vec{e}(j)) \right. \\ &\quad \left. - \sum_{j \neq i} w_{ij} d(\vec{e}(i), \vec{e}(j)) \right] \\ &\geq d(\vec{e}(i), s_i) - \frac{(1 - \alpha)}{\alpha} d(x, \vec{e}(i)). \end{aligned}$$

This inequality holds also for the partial information setting. In fact, in this case, if we apply the equilibrium condition, we obtain

$$\begin{aligned} \alpha d(\vec{e}(i), s_i) + (1 - \alpha) d(\vec{e}(i), \text{aggr}_i(\vec{e}_{-i})) \\ \leq \alpha d(x, s_i) + (1 - \alpha) d(x, \text{aggr}_i(\vec{e}_{-i})), \end{aligned}$$

from which, by applying the triangle inequality on Z , we have

$$\begin{aligned} d(x, s_i) &\geq d(\vec{e}(i), s_i) - \frac{(1 - \alpha)}{\alpha} \left[d(x, \text{aggr}_i(\vec{e}_{-i})) \right. \\ &\quad \left. - d(\vec{e}(i), \text{aggr}_i(\vec{e}_{-i})) \right] \\ &\geq d(\vec{e}(i), s_i) - \frac{(1 - \alpha)}{\alpha} d(x, \vec{e}(i)). \end{aligned}$$

$\square$

Proof of Proposition 5

Proof. Let us first assume that the game is unrestricted, i.e., $N_S = N$. Let us consider any equilibrium $\vec{e} \in E$, any agent i and any strategy $x \neq \vec{e}(i)$ (recall that $|Z| \geq 2$). In order to prove the claim, we need to show that $d(x, s_i) = d(x, \vec{e}(i))$. Assuming $\alpha > 1/2$, by Theorem 3, we get $\vec{e}(i) = s_i$, which implies that $d(x, s_i) = d(x, \vec{e}(i))$.

Next, we assume that (U, d) is a uniform metric. As before, we consider any equilibrium $\vec{e} \in E$, any agent i and any strategy $x \neq \vec{e}(i)$. In order to prove the claim, we need to show that $d(x, s_i) = d(x, \vec{e}(i))$. Notice that, since $x \neq \vec{e}(i)$, we have $d(x, \vec{e}(i)) = 1$. Therefore, in order for the claim to hold, it must be that $d(x, s_i) = 1$. By contradiction, let us assume that $d(x, s_i) = 0$, or equivalently $x = s_i$. This implies that i belongs to N_S . If we assume $\alpha > 1/2$, by Theorem 3, we get $s_i = \vec{e}(i)$, and by transitivity $x = \vec{e}(i)$, i.e., $d(x, \vec{e}(i)) = 0$, contradicting the hypothesis. $\square$

Proof of Proposition 7

In the proof of Proposition 7, we use the following result:

Theorem 7 ((Fanelli 2021)). *Let i be any agent, $\vec{x}$ any state, and y any strategy in Z . If d is L_2^2 and aggr_i is the Fréchet median then*

$$\begin{aligned} & \sum_{j \in N \setminus \{i\}} w_{ij} d(y, \vec{x}(j)) \\ &= d(y, \text{aggr}_i(\vec{x}_{-i})) + \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{x}_{-i}), \vec{x}(j)). \end{aligned}$$

We are now ready to prove the proposition.

Proof. We proof the clam for the case in which every optimum (wrt. any objective) is a consensus. The proof for the case in which every equilibrium is a consensus is symmetric and follows by simply exchanging the roles of $\vec{e}$ and $\vec{o}$ in the argument below.

We start with the proof of (a). Let us consider any equilibrium $\vec{e} \in E$, any social optimum $\vec{o}$ (either $\vec{o} \in O_{\text{SUM}}$ or $\vec{o} \in O_{\text{MAX}}$) and any agent i . By the definition of the Fréchet median,

$$\begin{aligned} & \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{e}_{-i}), \vec{e}(j)) \\ & \leq \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{o}_{-i}), \vec{e}(j)). \end{aligned} \quad (21)$$

Hence, we have that

$$\begin{aligned} & d(\text{aggr}_i(\vec{o}_{-i}), \text{aggr}_i(\vec{e}_{-i})) \\ &= \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{o}_{-i}), \text{aggr}_i(\vec{e}_{-i})) \end{aligned} \quad (22)$$

$$\begin{aligned} & \leq \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{o}_{-i}), \vec{e}(j)) \\ & \quad + \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{e}_{-i}), \vec{e}(j)) \end{aligned} \quad (23)$$

$$\leq 2 \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{o}_{-i}), \vec{e}(j)) \quad (24)$$

$$= 2 \sum_{j \in N \setminus \{i\}} w_{ij} d(\vec{o}(j), \vec{e}(j)) = 2 \cdot \pi_i(\vec{o}, \vec{e}), \quad (25)$$

where (22) follows from $\sum_{j \in N \setminus \{i\}} w_{ij} = 1$, (23) from the triangle inequality, (24) from (21) and (25) from the fact that

$\vec{o}$ is a consensus. Hence $\hat{\tau}_i \leq 2$. Using that $\tau_i = \max\{\hat{\tau}_i, 1\}$ and $\tau = \max_{i \in N} \tau_i$, we conclude the proof of (a).

Let us prove (b) now. Also in this case, let us consider any equilibrium $\vec{e} \in E$, any social optimum $\vec{o}$ (either $\vec{o} \in O_{\text{SUM}}$ or $\vec{o} \in O_{\text{MAX}}$) and any agent i . We have that

$$\begin{aligned} & d(\text{aggr}_i(\vec{o}_{-i}), \text{aggr}_i(\vec{e}_{-i})) \\ & \leq \sum_{j \in N \setminus \{i\}} w_{ij} d(\text{aggr}_i(\vec{o}_{-i}), \vec{e}(j)) \end{aligned} \quad (26)$$

$$= \sum_{j \in N \setminus \{i\}} w_{ij} d(\vec{o}(j), \vec{e}(j)) = \pi_i(\vec{o}, \vec{e}), \quad (27)$$

where (26) follows from Theorem 7 by setting $y = \text{aggr}_i(\vec{o}_{-i})$ and $\vec{x} = \vec{e}$, and (27) from the fact that $\vec{o}$ is a consensus. Hence $\hat{\tau}_i = 1$. Using that $\tau_i = \max\{\hat{\tau}_i, 1\}$ and $\tau = \max_{i \in N} \tau_i$, we conclude the proof of (b). $\square$

Missing Details from the Proof of Theorem 6

Proof. We complete the proof of Theorem 6, by observing that we can find $\min_{s_i \in S} d(x, s_i)/(2\ell)$ by computing an s_i that minimizes (8), subject to (9) and (11). Since (8) is a strictly convex function of $s_i(1), \dots, s_i(m)$ and the constraints of (9) and (11) are linear on $s_i(1), \dots, s_i(m)$, we have to solve a strictly convex minimization problem with linear constraints.

We observe that

$$s_i^*(j) = \begin{cases} \frac{2\alpha - 1}{2\alpha} & \text{if } j \in C_{01} \\ \frac{1}{2\alpha} & \text{if } j \in C_{10} \\ 0 & \text{if } j \in C_0 \setminus C_{01} \\ 1 & \text{if } j \in C_1 \setminus C_{10} \end{cases}$$

is feasible (i.e., it lies in S and satisfies (11)) and achieves an objective value of $d(x, s_i^*)/(2\ell) = \left(\frac{2\alpha-1}{2\alpha}\right)^2$. Moreover, s_i^* is defined for any strategy $x \in Z$ and any equilibrium strategy $\vec{e}(i) \in Z$ with $d(x, \vec{e}(i)) = 2\ell > 0$ and always attains the same objective value $d(x, s_i^*)/(2\ell) = \left(\frac{2\alpha-1}{2\alpha}\right)^2$.

To confirm that s_i^* is indeed a minimizer of the above strictly convex minimization problem, we observe that the KKT optimality conditions are satisfied at s_i^* with KKT multipliers $\mu = \frac{2\alpha-1}{2\ell\alpha}$ for (11) and 0 for all the constraints of (9). To see this, let

$$\begin{aligned} f(y) &= \frac{1}{2\ell} \left(\sum_{j \in C_0} y(j)^2 + \sum_{j \in C_1} (1 - y(j))^2 \right) \quad \text{and} \\ g(y) &= \sum_{j \in C_{10}} y(j) - \sum_{j \in C_{01}} y(j) - \ell(1 - \alpha)/\alpha \end{aligned}$$

We first observe that complementary slackness constraints are satisfied, since $g(s_i^*) = 0$ and all other KKT multipliers are 0. Primal and dual feasibility constraints are also satisfied. Regarding stationarity constraints, we observe that $\nabla_j f(s_i^*) = 0$, for any $j \notin C_{01} \cup C_{10}$, $\nabla_j f(s_i^*) = \frac{2\alpha-1}{2\ell\alpha}$, for all $j \in C_{01}$, and $\nabla_j f(s_i^*) = \frac{1-2\alpha}{2\ell\alpha}$, for all $j \in C_{10}$. Similarly, $\nabla_j g(s_i^*) = 0$ for any $j \notin C_{01} \cup C_{10}$, $\nabla_j g(s_i^*) = -1$,

for all $j \in C_{01}$, and $\nabla_j g(s_i^*) = 1$, for all $j \in C_{10}$. Therefore, stationarity constraints are also satisfied with KKT multiplier $\mu = \frac{2\alpha-1}{2\ell\alpha}$. $\square$