Generalized Discrete Preference Games

Vincenzo Auletta
University of Salerno
auletta@unisa.it

Ioannis Caragiannis
University of Patras
caragian@upatras.gr

Diodato Ferraioli
University of Salerno
dferraioli@unisa.it

Clemente Galdi
University of Naples "Federico II"
clemente.galdi@unina.it

Giuseppe Persiano
University of Salerno
giuper@unisa.it

Abstract

Recently, much attention has been devoted to discrete preference games to model the formation of opinions in social networks. More specifically, these games model the agents' strategic decision of expressing publicly an opinion, which is a result of an interplay between the agent's private belief and the social pressure. However, these games have very limited expressive power; they can model only very simple social relations and they assume that all the agents respond to social pressure in the same way. In this paper, we define and study the novel class of *generalized discrete preference games*. These games have additional characteristics that can model social relations to allies or competitors and complex relations among more than two agents. Moreover, they introduce different levels of strength for each relation, and they personalize the dependence of each agent to her neighborhood.

We show that these novel games admit generalized ordinal potential functions and, more importantly, that every two-strategy game that admits a generalized ordinal potential function is structurally equivalent to a generalized discrete preference game. This implies that the games in this novel class capture the full generality of two-strategy games in which the existence of (pure) equilibria is guaranteed by topological arguments.

1 Introduction

Understanding the emergent behavior of groups of interacting agents is an important research challenge in the social sciences and in AI. One of the problems considered along this direction is to understand how opinions are formed and expressed in a social context, e.g., how an opinion diffuses in a social network when each single agent adapts her own belief in response to the opinions expressed by her "friends." The recent research activity on this subject is extensive, both in AI and in multiagent systems [Pryymak *et al.*, 2012; Tsang and Larson, 2014] and in CS at large [Acemoglu and Ozdaglar, 2011; Bindel *et al.*, 2011; Mossel and Tamuz, 2014], as well as in sociology, economics, physics, and epidemiology.

A classical model for studying the diffusion and the adoption of opinions in a social network has been proposed by Friedkin and Johnsen [1990] as a refinement of a previous model introduced by DeGroot [1974]. The model assumes that each agent has a *private belief*, but the *opinion* she eventually expresses can be different from her belief, and it is given by the outcome of a repeated averaging between her belief and the opinions of individuals with whom she has social relations. Recently, Bindel *et al.* [2011] considered this model and proved that, under mild assumptions, whenever beliefs and opinions belong to $[0, 1]$, the repeated averaging process leads to a *unique* equilibrium, describing the opinion that each agent eventually expresses.

Two recent papers — by Ferraioli *et al.* [2012] and by Chierichetti *et al.* [2013] — focus on a refinement of this model, where beliefs and opinions are discrete and can take two possible values: 0 and 1. This simple restriction has important effects. Clearly, the discrete nature of the preferences does not allow for averaging anymore and several nice properties of the opinion formation models mentioned above — such as the uniqueness of the outcome — are lost. In contrast, it now seems natural to assume that each agent behaves strategically and aims to pick the most beneficial (or less costly) opinion for her, where the benefit (or cost) depends both on her internal belief and on the opinions of individuals with whom she has social relations.

This immediately defines a *discrete preference game*. With the notation that we will use later, let us define the cost of agent i when the binary strategies of the n agents are given by the vector $\mathbf{s} = (\mathbf{s}(1), \ldots, \mathbf{s}(n))$ as

$$c_i(\mathbf{s}) = \alpha \cdot |\mathbf{s}(i) - \mathbf{b}(i)| + (1-\alpha) \cdot \sum_{j \in N(i)} |\mathbf{s}(i) - \mathbf{s}(j)|, \quad (1)$$

where $\mathbf{b}(i) \in \{0, 1\}$ denotes the belief of agent i and $N(i)$ is the set of individuals with whom she has social relations (e.g., friends in the social network). Note that the cost has two components that depend on the distance of the agent's strategy from her internal belief and from the strategies of her neighbors, respectively. The parameter $\alpha \in [0, 1]$ adjusts the relative importance of the two terms. Intuitively, it measures the degree of coordination the agents seek.

The class of discrete preference games was recently considered by Chierichetti *et al.* [2013]. They assess the quality of states using the total (or social) cost of the agents.

Their findings include tight bounds on the price of stability for discrete preference games as well as conditions that imply that states of minimum social cost are always equilibria. These results come in contrast to the price of anarchy that can be unbounded; the same result has also been observed independently by Ferraioli *et al.* [2012]. Recently, Auletta *et al.* [2015] also proved that a best response dynamics can lead from the truthful profile (in which agents' opinion coincides with their belief) to an equilibrium in which the majority has been subverted.

Still discrete preference games have evident limitations. In order to overcome them, we significantly generalize the model; this is our main conceptual contribution. Chierichetti *et al.* [2013] generalized these games to more than two strategies per agent. However, this kind of generalization faces the problem that it is not clear whether there exists an ideal metric for representing distances among opinions, and different results are obtained by adopting different metrics. In addition, games with two strategies per agent have been proved to be as general (and challenging) as possible from the computational complexity point of view; see [Fabrikant *et al.*, 2004]. For these reasons, we keep the two-strategy restriction and instead consider richer relations among agents. For example, the facts that each agent treats her neighbors equally, she seeks agreement to all of them, and all social relations include only two agents are obvious limitations of discrete preference games. In contrast, we would like to (1) model social relations to both allies or competitors, (2) include complex relations between more than two agents, (3) introduce different levels of strength for each relation, and (4) personalize the dependence of each agent to its neighborhood.

So, motivated by games that are inspired by constraint satisfaction optimization problems — such as cut games and party affiliation games; see, e.g., [Balcan *et al.*, 2013; Caragiannis *et al.*, 2014; Wooldridge *et al.*, 2013] — we define the broader class of *generalized discrete preference games*. In these games, the strategy of each agent is again a binary opinion. Social relations (or social constraints) are weighted Boolean formulas over the strategies of subsets of agents. Then, the cost of each agent depends on the total weight of the unsatisfied formulas whose outcome depend on the opinion of the agent; the particular dependence can be different from agent to agent. We give the detailed definition of generalized discrete preference games in Section 3, where we also present and discuss several examples.

Our main technical contributions are presented in Section 4. First, we kill many birds with one stone: in Theorem 1, we show that the negative sum of the weights of the satisfied formulas is a generalized ordinal potential function for such a game. Note that the same potential function may correspond to very different games and essentially describes the general structure of the Nash dynamics graph of each of them. In addition, depending on some details in the definition of the agents' costs, the potential function can be proved to be exact, weighted, or ordinal potential. As a consequence, we have that generalized discrete preference games always admit pure Nash equilibria, and that a finite sequence of best-response strategy updates always converges to an equilibrium. Potential functions usually reveal a structure in the underlying game and have turned out to be useful for a more deep understanding of it. For example, they have been useful in proving bounds on the price of stability [Anshelevich *et al.*, 2008] or in bounding the time required to reach equilibria or efficient states after best-response play [Awerbuch *et al.*, 2008]. We hope that the potential function we define can play such important roles for generalized ordinal potential games.

Second, and probably more importantly, in Theorem 4 we show that every game with two strategies per agent that admits a generalized ordinal potential is structurally equivalent (in particular, *better-response equivalent*) to a generalized discrete preference game. This implies that generalized discrete preference games capture the full generality of two-strategy games in which the existence of pure equilibria is guaranteed by topological arguments. This result is similar in spirit to the equivalence proved by Monderer and Shapley [1996] between exact potential and congestion games. This equivalence has been extremely useful for a deep understanding of the former class of games. On the other side, generalized ordinal potential games are still not well understood, even if it has been proved that they have many important applications such as in load balancing [Even-Dar *et al.*, 2007] and in wireless spectrum management [Wong *et al.*, 2014]. We hope that generalized discrete preference games can play the same role as congestion games for improving our knowledge about generalized ordinal potential games. Ideally, we would like to be able to prove properties of generalized ordinal potential games by focusing on equivalent generalized discrete preference games.

Other Related Work. The paper that is most closely related to ours is the one by Chierichetti *et al.* [2013] which we have discussed above. Independently, Ferraioli *et al.* [2012] have considered the same class of games, mostly focusing on the study of (noisy) best-response dynamics with respect to their convergence to equilibria or stable states. An interesting generalization has been proposed by Bhawalkar *et al.* [2013] who consider beliefs and opinions in $[0, 1]$; according to their model, in addition to the opinion expressed, the social relations of an agent are also part of her strategy and are selected so that connections to agents with similar opinions are more preferable.

In addition to the cited papers that study games involving Boolean variables and constraints, Bonzon *et al.* [2006] defined boolean games where, like generalized discrete preference games, the interdependence among agents is represented through boolean formulas involving their strategies. However, classical boolean games assume that formulas are not weighted, while we will see that weights are crucial for our equivalence result to hold. Recently, extensions to boolean games that include weights (e.g., see [Harrenstein *et al.*, 2014]) have been proposed, but still they are too restrictive for achieving our results.

Studies in social networks consider several phenomena related to the spread of social influence such as information cascading, network effects, epidemics, and more. The book of Easley and Kleinberg [2010] provides an excellent introduction to the theoretical treatment of such phenomena. From a different perspective, problems of this type have also been considered in the distributed computing literature, motivated by the need to control and restrict the influence of failures in distributed systems; e.g., see the survey by Peleg [2002] and the references therein.

2 Preliminaries

Discrete preference games are formally defined as follows. There are n agents; we use $[n] = \{1, 2, \ldots, n\}$ to denote their set. Each agent corresponds to a distinct node of an undirected graph $G = (V, E)$ that represents the *social network*; i.e., the network of social relations between the agents. Agent i has an (internal) *belief* $\mathbf{b}(i) \in \{0, 1\}$ and her strategy set consists of the two alternatives that she can declare as her *opinion* (or *preference*), i.e., $\mathbf{s}(i) \in \{0, 1\}$. A *strategy profile* (or, simply, a *profile* or *state*) is a vector of strategies, with one strategy per agent. We use bold symbols for profiles; i.e., $\mathbf{s} = (\mathbf{s}(1), \ldots, \mathbf{s}(n))$. For $y \in \{0, 1\}$, we denote as $\overline{y}$ the negation of y; i.e., $\overline{y} = 1 - y$.

At a profile $\mathbf{s}$, the cost of agent i is denoted by $c_i(\mathbf{s})$ and is defined as in (1). A profile $\mathbf{s}$ is a *pure Nash equilibrium* (or, simply, an equilibrium) if $c_i(\mathbf{s}) \leq c_i(\overline{\mathbf{s}(i)}, \mathbf{s}_{-i})$ for every agent i. The notation $(\overline{\mathbf{s}(i)}, \mathbf{s}_{-i})$ is used to refer to the state in which all agents besides agent i follow their strategies in state $\mathbf{s}$ and agent i follows strategy $\overline{\mathbf{s}(i)}$. As observed in [Chierichetti *et al.*, 2013] and [Ferraioli *et al.*, 2012], discrete preference games always have equilibria. This follows using potential function arguments; we discuss potential functions of different types below.

Following the standard terminology in the game theory literature, we say that a game is an *exact potential game* if there exists a function Φ defined over the states of the game such that the following condition holds: For every two states $(s, \mathbf{x}_{-i})$ and $(s', \mathbf{x}_{-i})$ that differ in the strategy of agent i, $c_i(s, \mathbf{x}_{-i}) - c_i(s', \mathbf{x}_{-i}) = \Phi(s, \mathbf{x}_{-i}) - \Phi(s', \mathbf{x}_{-i})$. It is a *weighted potential game* if there exists a function Φ and positive weights $(v_i)_{i \in [n]}$ so that the above condition becomes $c_i(s, \mathbf{x}_{-i}) - c_i(s', \mathbf{x}_{-i}) = v_i \cdot (\Phi(s, \mathbf{x}_{-i}) - \Phi(s', \mathbf{x}_{-i}))$. It is an *ordinal potential game* if there exists a function Φ such that the condition becomes $c_i(s, \mathbf{x}_{-i}) - c_i(s', \mathbf{x}_{-i}) > 0 \Leftrightarrow \Phi(s, \mathbf{x}_{-i}) - \Phi(s', \mathbf{x}_{-i}) > 0$; Finally, it is a *generalized ordinal potential game* if there exists a function Φ such that $c_i(s, \mathbf{x}_{-i}) - c_i(s', \mathbf{x}_{-i}) > 0 \Rightarrow \Phi(s, \mathbf{x}_{-i}) - \Phi(s', \mathbf{x}_{-i}) > 0$.

Note that each class generalizes the previous one. It is a folklore result — e.g., see Monderer and Shapley [1996] — that every (finite) generalized ordinal potential game has at least one Nash equilibrium and it can be reached through cost-decreasing strategy updates of the agents. Specifically, any local minimum of the potential function corresponds to an equilibrium of the game. The opposite — i.e., the fact that equilibria always correspond to local minima of the potential function — is true for ordinal potential games but it is not necessarily true for generalized ordinal potential games.

3 Generalized Discrete Preference Games

We are now ready to define the broader class of *generalized discrete preference* (GDP) games. Like in discrete preference games, again we assume that every agent i has a private belief $\mathbf{b}(i) \in \{0, 1\}$ and that her strategy $\mathbf{s}(i)$ is a preference from the strategy set $\{0, 1\}$. However, in contrast to discrete preference games, agents are not only interested in agreeing to their neighbors and more complex *constraints* can be used to represent their preferences. Specifically, we define a *social constraint* as a Boolean formula involving the opinions (but not the beliefs) of a subset of agents. For example, the social constraint $C(\mathbf{s}) = (\mathbf{s}(i) \wedge \overline{\mathbf{s}(j)}) \vee (\overline{\mathbf{s}(i)} \wedge \mathbf{s}(j))$ involves agents i and j and is satisfied when the opinions of i and j are different. In addition to social constraints, GDP games have *belief constraints*; these constraints are true when the expressed opinion of an agent coincides with her belief. Specifically, the belief constraint $B_i(\mathbf{s})$ of agent i is $B_i(\mathbf{s}) = (\mathbf{s}(i) \wedge \mathbf{b}(i)) \vee (\overline{\mathbf{s}(i)} \wedge \overline{\mathbf{b}(i)})$ and, unlike social constraints, involves only the opinion and the belief of agent i. Sometimes, we will use the equivalent constraint $B_i(\mathbf{s}) = \overline{\mathbf{s}(i)}$ (respectively, $B_i(\mathbf{s}) = \mathbf{s}(i)$) if $\mathbf{b}(i) = 0$ (respectively, $\mathbf{b}(i) = 1$). A set $\mathcal{C}$ of constraints is *feasible* if, for every agent i, at most one of the two belief constraints $\mathbf{s}(i)$ and $\overline{\mathbf{s}(i)}$ belongs to $\mathcal{C}$.

In a generalized discrete preference game, each constraint C has a non-negative *weight* $W(C)$. For any agent i and for any profile $\mathbf{s}$, we also denote by $w_i(\mathbf{s})$ the sum of the weights of the constraints involving agent i that are satisfied at $\mathbf{s}$. Furthermore, we denote by w_i the sum of the weights of all constraints that involve agent i and we denote by $w(\mathbf{s})$ the sum of all constraints satisfied at $\mathbf{s}$. Also, we let w denote the sum of the weights of all constraints.

The *cost* $c_i(\mathbf{s})$ of agent i in profile $\mathbf{s}$ is defined through a monotone non-decreasing function $F_i \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ on the weight of the constraints that involve i and are unsatisfied at $\mathbf{s}$. More precisely, we have $c_i(\mathbf{s}) = F_i(w_i - w_i(\mathbf{s}))$. Note that using different functions F_i we can model agents that evaluate their social relationships in different ways. We emphasize that, in the definition of the cost of agent i, the function F_i gets as input the *total* weight of *all* unsatisfied constraints that involve i. In a sense, an agent cares about all (belief and social) constraints she is involved in and cannot disregard any constraint. Observe that, through these agent-specific functions is possible to model agents that give different importance to the constraints in which they are involved, even if constraint weights are common to each agent. An example of such a setting is described later.

To sum up, a *generalized discrete preference game* is a tuple $([n], \mathcal{C}, W, (F_i)_{i \in [n]})$, where $\mathcal{C}$ is a feasible set of constraints, W is the weight function and F_i is the cost function of i. Next we give examples showcasing the expressiveness of GDP games.

Examples. We give examples of games that fall within our class, mostly focusing on notions that will be used again later. A constraint C is an *equality-constraint* if it is satisfied only when all agents involved in C have the same opinion. Formally, C is an equality-constraint on $N \subseteq [n]$ if $C(\mathbf{s}) = \left(\bigwedge_{i \in N} s_i\right) \vee \left(\bigwedge_{i \in N} \overline{s}_i\right)$. This allows to easily model a group of agents that want to agree with each other. A constraint C is an *or-constraint* if it is satisfied in all profiles except the ones in which the agents involved in C have a specific configuration. Formally, C is an or-constraint on N if $C(\mathbf{s}) = \bigvee_{i \in N} \ell_i$, where $\ell_i \in \{s_i, \overline{s}_i\}$. This allows to model the natural requirement of an agent that wants to avoid a specific configuration.

We also introduce another class of constraints, that, even though it is less natural than previous ones, it will turn out to be sometimes useful for characterizing the behavior of GDP games. Specifically, a constraint C is a *switch-constraint* if every time an agent that is involved in the constraint switches her preference then the constraint switches between satisfied and unsatisfied. Formally, C is a switch-constraint on set N of agents if $C(\mathbf{s}) = \bigoplus_{i \in N} \ell_i$, where $\ell_i \in \{s_i, \bar{s}_i\}$. Informally, a switch constraint represents a parity requirement.

According to the definition of GDP games, the weight of a constraint is the same for all agents involved in the constraint. Nevertheless, as we will show next, this restriction still allows to implement many natural cost functions[1].

In particular, very natural cost definitions can be obtained using specific choices for the F_i functions. For example, the cost function, according to which an agent tries to minimize the sum of the weights of the unsatisfied constraints in which she is interested, can be modeled by the *identity* function; that is, $F_i(r) = r$ for each $r \geq 0$.

In another natural example, the cost function assigns to the belief constraint an agent-specific weight that is different from the one assigned to social constraints. For example, consider the following cost function $c_i(\mathbf{s}) = \alpha(1 - B_i(\mathbf{s})) + (1-\alpha)\frac{d_i - n_i(\mathbf{s})}{d_i}$ where d_i is the number of social constraints in which i is involved and $n_i(\mathbf{s})$ is the number of social constraints in which i is involved among those that are satisfied at $\mathbf{s}$. This cost function then balances among the closeness to agent's own belief and the fraction of satisfied constraints in which she is involved. Even though in our definition the weight of a constraint is the same for all agents, we can implement this cost function through a *weighted identity* function F_i; i.e., $F_i(r) = v_i r$ for any $r \geq 0$, where $v_i > 0$. Indeed, we can set $W(B_i) = \alpha d_i$ for each $i \in [n]$, $W(C) = 1 - \alpha$ for each social constraint $C \in \mathcal{C}$ and $v_i = \frac{1}{d_i}$ for each $i \in [n]$.

Another natural class of cost functions is the one according to which agents are happy if and only if the sum of weights of the unsatisfied constraints is less than the sum of weights of satisfied constraints. This can be achieved by using the *majority* function, i.e., $F_i(r) = 1$ if $r \geq \frac{w_i}{2}$ and $F_i(r) = 0$ otherwise. We can also consider *threshold* functions, according to which the agents are happy if and only if the sum of weights of the unsatisfied constraints is less than some threshold T, i.e., $F_i(r) = 1$ if $r \geq T$ and $F_i(r) = 0$ otherwise.

Note that discrete preference games are the special case of GDP games in which all non-belief constraints are equality constraints for pairs of agents and the cost of each agent is defined using the identity function.

4 Existence of Equilibria

We present our structural results in this section. We begin by proving that GDP games are generalized ordinal potential games. As such, they always admit a pure Nash equilibrium.

[1]Clearly, it is possible to introduce in the game player-specific unrelated weights. However, this will break our equivalence result.

Theorem 1. *Let $\mathcal{G} = ([n], \mathcal{C}, W, (F_i)_{i\in[n]})$ be a generalized discrete preference game. The function $-w(\mathbf{s})$ is a generalized ordinal potential function for $\mathcal{G}$.*

Proof. Consider an agent i and profile $\mathbf{s}$. Without loss of generality, let us assume that $c_i(0, \mathbf{s}_{-i}) - c_i(1, \mathbf{s}_{-i}) > 0$; we will show that $w(1, \mathbf{s}_{-i}) - w(0, \mathbf{s}_{-i}) > 0$.

By the definition of the cost of agent i, we have $F_i(w_i - w_i(0, \mathbf{s}_{-i})) - F_i(w_i - w_i(1, \mathbf{s}_{-i})) > 0$ and, since the function F_i is monotone non-decreasing, $w_i(0, \mathbf{s}_{-i}) - w_i(1, \mathbf{s}_{-i}) < 0$. Now, for the total weight of constraints that do not depend on the opinion or the belief of agent i, we obviously have

$$w(0, \mathbf{s}_{-i}) - w_i(0, \mathbf{s}_{-i}) = w(1, \mathbf{s}_{-i}) - w_i(1, \mathbf{s}_{-i}). \tag{2}$$

Hence, $w(1, \mathbf{s}_{-i}) - w(0, \mathbf{s}_{-i}) > 0$, as desired. □

The function $-w(\mathbf{s})$ gives a stronger characterization for special subclasses of GDP games.

Proposition 2. *Let $\mathcal{G} = ([n], \mathcal{C}, W, (F_i)_{i\in[n]})$ be a generalized discrete preference game. Then, the function $-w(\mathbf{s})$ is: An exact potential for $\mathcal{G}$ if the F_i's are identity functions; a weighted potential for $\mathcal{G}$ if the F_i's are weighted identity functions; an ordinal potential for $\mathcal{G}$ if every constraint $C \in \mathcal{C}$ is a switch-constraint and the F_i's are majority functions.*

Proof. Consider an agent i and profile $\mathbf{s}$. Let us first assume that F_i is a weighted identity function with $F_i(r) = v_i \cdot r$. Hence, $c_i(0, \mathbf{s}_{-i}) - c_i(1, \mathbf{s}_{-i}) = v_i \cdot (w_i - w_i(0, \mathbf{s}_{-i})) - v_i \cdot (w_i - w_i(1, \mathbf{s}_{-i})) = v_i \cdot (w_i(1, \mathbf{s}_{-i}) - w_i(0, \mathbf{s}_{-i})) = v_i \cdot (w(1, \mathbf{s}_{-i}) - w(0, \mathbf{s}_{-i}))$, where the last equality follows by the observation (2) above. The claim for identity functions follows by setting $v_i = 1$ for every agent i.

Now, let us focus on a generalized discrete preference game with switch constraints and majority functions for the definition of agents' cost. Since we have already proved that $-w(\mathbf{s})$ is a generalized ordinal potential function, all that is left to prove is that $w(0, \mathbf{s}_{-i}) - w(1, \mathbf{s}_{-i}) < 0$ implies $c_i(0, \mathbf{s}_{-i}) - c_i(1, \mathbf{s}_{-i}) > 0$.

From observation (2), $w(0, \mathbf{s}_{-i}) - w(1, \mathbf{s}_{-i}) < 0$ implies that $w_i(0, \mathbf{s}_{-i}) - w_i(1, \mathbf{s}_{-i}) < 0$ for every $\mathbf{s}$ and i. Furthermore, for switch constraints we have that $w_i(0, \mathbf{s}_{-i}) + w_i(1, \mathbf{s}_{-i}) = w_i$. Therefore we obtain that $w_i(1, \mathbf{s}_{-i}) > \frac{w_i}{2}$ and $w_i(0, \mathbf{s}_{-i}) < \frac{w_i}{2}$. Since F_i is the majority function, it follows that $c_i(0, \mathbf{s}_{-i}) - c_i(1, \mathbf{s}_{-i}) > 0$. □

Let us note that Proposition 2 is, in a sense, "tight". For the first two cases of Proposition 2, it is easy to construct GDP games with non-identity and non-weighted identity functions that do not admit exact and weighted potential functions, respectively. We give examples indicating that, by slightly deviating from either of the assumptions of switch constraints or majority functions, we obtain games that do not admit an ordinal potential function.

Indeed, we will first show that if F_i are threshold (but not majority) functions, we may have a game that does not admit an ordinal potential function. Consider a game with two agents, the switch constraints $C_1 = s_1 \oplus \bar{s}_2$ and $C_2 = s_1 \oplus s_2$ and the belief constraints $B_1 = s_1$ and $B_2 = s_2$ with weights $W(C_1) = 1$, $W(C_2) = 2$, $W(B_1) = 1$ and $W(B_2) = 2$. F_1 is the majority function and F_2 is a threshold function with a threshold at 5. Suppose that an ordinal potential Φ exists. It is easy to see that $c_1(0,0) > c_1(1,0)$ and, hence, it must also be $\Phi(0,0) > \Phi(1,0)$. On the other side, we can easily verify that $c_2(0,0) = c_2(0,1)$, $c_1(0,1) = c_1(1,1)$ and $c_2(1,1) = c_2(1,0)$. Hence, it should also be $\Phi(0,0) = \Phi(0,1) = \Phi(1,1) = \Phi(1,0)$, a contradiction.

The same may happen if the F_i's are majority functions but there are non-switch constraints. Consider a game with three agents, the switch constraint $C_1 = s_1 \oplus \overline{s}_2$, two or-constraints $C_2 = s_1 \vee s_3$ and $C_3 = s_2 \vee s_3$, and belief constraints $B_1 = s_1$, $B_2 = s_2$, and $B_3 = s_3$ with weights $W(C_1) = W(B_1) = W(B_2) = W(B_3) = 3$, $W(C_2) = 2$, $W(C_3) = 8$. We can easily verify that $c_1(0,1,1) = 1$ but $c_1(1,1,1) = c_2(1,1,1) = c_2(1,0,1) = c_1(1,0,1) = c_1(0,0,1) = c_2(0,0,1) = c_2(0,1,1) = 0$. An ordinal potential function Φ should then satisfy $\Phi(0,1,1) > \Phi(1,1,1) = \Phi(1,0,1) = \Phi(0,0,1) = \Phi(0,1,1)$, a contradiction.

4.1 Better-Response Equivalence

We now define the concept of better-response equivalence in order to state the main theorem of this section.

Definition 3. *Two games $\mathcal{G} = ([n], (S_i)_{i\in[n]}, (c_i)_{i\in[n]})$ and $\mathcal{G}' = ([n], (S_i)_{i\in[n]}, (c'_i)_{i\in[n]})$ with the same set of profiles S are* better-response equivalent *if, for every pair of profiles $\mathbf{s}, \mathbf{s}' \in S$ that differ in the strategy of only one agent (say i), we have that $c_i(\mathbf{s}) > c_i(\mathbf{s}')$ if and only if $c'_i(\mathbf{s}) > c'_i(\mathbf{s}')$.*

Clearly, one could construct artificial classes of games that are better-response equivalent to generalized ordinal potential games. The importance of our next statement is that it involves a very *natural* class of games, namely GDP games.

Theorem 4. *Any two-strategy generalized ordinal potential game $\mathcal{G}$ is better-response equivalent to a generalized discrete preference game $\mathcal{G}'$.*

The rest of this subsection is devoted to proving Theorem 4. Without loss of generality, we assume that the strategy set of each agent of $\mathcal{G}$ is $\{0,1\}$. The profiles can then be identified with the nodes of an n-dimensional undirected hypercube. An undirected edge $(\mathbf{s}, \mathbf{s}')$ of the hypercube connects two profiles that differ in the strategy of a single agent; e.g., $\mathbf{s} = (0, \mathbf{s}_{-i})$ and $\mathbf{s}' = (1, \mathbf{s}_{-i})$, for some $i \in [n]$. Sometimes we denote edge $(\mathbf{s}, \mathbf{s}')$ by $(\mathbf{s}, i)$ or, equivalently, by $(\mathbf{s}', i)$ and call it a *dimension-i* edge.

Directing and Tagging Edges. If $(\mathbf{s}, \mathbf{s}')$ is a dimension-i edge and $c_i(\mathbf{s}) > c_i(\mathbf{s}')$, then we direct it from $\mathbf{s}$ to $\mathbf{s}'$. If instead $c_i(\mathbf{s}) = c_i(\mathbf{s}')$, the edge is left undirected. Notice that, since $\mathcal{G}$ is a generalized ordinal potential game, the partially directed hypercube is essentially the Nash dynamics graph of $\mathcal{G}$ and contains no cycle consisting only of directed edges. Then we tag undirected edges as *good* as long as we can do so without creating a cycle consisting solely of directed and good edges. Remaining edges are called *bad*.

Claim 5. *For each pair of profiles $\mathbf{s}, \mathbf{s}'$, there is an undirected path on the partially directed hypercube between $\mathbf{s}$ and $\mathbf{s}'$ that does not use bad edges.*

Proof. Suppose that all undirected paths from $\mathbf{s}$ to $\mathbf{s}'$ contain some bad edge. If we tag at least one bad edge as good, no cycle is created, a contradiction. $\square$

Defining the Constraints. Let us now start to describe the generalized discrete preference game $\mathcal{G}'$ that will turn out to be equivalent to $\mathcal{G}$. It has two types of constraints.

Node constraints: For every node $\mathbf{s}$ of the hypercube, $\mathcal{G}'$ has constraint $V(\mathbf{s})$ that is satisfied at all profiles except at $\mathbf{s}$. It is easy to see that $V(\mathbf{s})$ can be expressed as an or-constraint and that every agent is involved in $V(\mathbf{s})$.

Edge constraints: For every edge $(\mathbf{s}, i)$, $\mathcal{G}'$ has constraint $E(\mathbf{s}, i)$ that is satisfied at all profiles except at each endpoints, $(0, \mathbf{s}_{-i})$ and at $(1, \mathbf{s}_{-i})$. It is easy to see that $E(\mathbf{s}, i)$ can be expressed as an or-constraint that involves all agents except i.

Defining the Constraint Weights. In order to define the weights of the constraints, let us fix some additional notation. We will identify a constraint with its weight. Thus, the weight of the constraint associated with edge $(\mathbf{s}, i)$ is $E(\mathbf{s}, i)$ and the weight of the constraint associated with node $\mathbf{s}$ is $V(\mathbf{s})$. We also denote by $E(\mathbf{s})$ the weight of all edge constraints incident to $\mathbf{s}$; that is, $E(\mathbf{s}) = \sum_i E(\mathbf{s}, i)$. Similarly, for an agent i we denote by $E(i)$ the sum of the weights of all dimension-i edge constraints; that is, $E(i) = \sum_{\mathbf{s}} E(\mathbf{s}, i)$. Finally, we denote by $V = \sum_{\mathbf{s}} V(\mathbf{s})$ the sum of the weights of all node constraints and by $E = \sum_{\mathbf{s},i} E(\mathbf{s}, i)$ the sum of the weights of all edge constraints.

We assign weight 0 to each edge constraint, except for the ones corresponding to bad edges that have weight 2^n.

The weights of the node constraints are determined as follows. We remove all bad edges from the (partially directed) hypercube and contract good edges by merging endpoints into super-nodes. By Claim 5 and by the tagging procedure, the resulting graph is a connected directed acyclic graph (DAG); that is, ignoring direction, for every two super-nodes X and Y, the DAG contains a directed path between X and Y. Weights of the node constraints are determined so that the two following properties hold: If $\mathbf{s}$ and $\mathbf{s}'$ belong to the same super-node X of the DAG, then $V(\mathbf{s}) + E(\mathbf{s}) = V(\mathbf{s}') + E(\mathbf{s}')$, and we set $W(X) = V(\mathbf{s}) + E(\mathbf{s})$; If $\mathbf{s}$ and $\mathbf{s}'$ belong to different super-nodes X and Y and there is a directed path from X to Y, then $W(X) = V(\mathbf{s}) + E(\mathbf{s}) > V(\mathbf{s}') + E(\mathbf{s}') = W(Y)$.

Note that an assignment satisfying both these properties always exists. Indeed, we can determine the weights of the node constraints according to a topological ordering of the super-nodes of the DAG. For every super-node X of in-degree 0, we determine the node $\mathbf{s}^\star$ with the maximum $E(\mathbf{s}^\star)$ among all of its nodes and set $V(\mathbf{s}^\star) = N = (n+1)2^n$. Then we set $V(\mathbf{s}) = N - E(\mathbf{s}) + E(\mathbf{s}^\star)$ for all the other nodes $\mathbf{s}$ of X. Observe that $W(X) = N + E(\mathbf{s}^\star)$. Suppose now that we have set the weights for all super-nodes that have an out-going edge to super-node Y and let X be the such super-node that minimizes $W(X)$. Then for all nodes $\mathbf{s}'$ of Y we set $V(\mathbf{s}') = W(X) - 1 - E(\mathbf{s}')$. Observe that $W(Y) = W(X) - 1$.

We next prove that these weights are non-negative. Let $\mathbf{s}$ be a node, $Y_{\mathbf{s}}$ be the corresponding super-node. Consider the super-nodes of in-degree 0 from which it is possible to reach $Y_{\mathbf{s}}$ in the DAG and let X be the such super-node that minimizes $W(X)$. We have $V(\mathbf{s}) \geq W(X) - \mathtt{dist}(Y_{\mathbf{s}}, X) - E(\mathbf{s}')$, where $\mathtt{dist}(Y_{\mathbf{s}}, X)$ denotes the distance between X

and $Y_{\mathbf{s}}$ in the DAG. Since $W(X) \geq N$, $\mathtt{dist}(Y_{\mathbf{s}}, X) \leq 2^n - 1$ and $E(\mathbf{s}') \leq n2^n$, we have that our choice of N is sufficient to make $V(\mathbf{s}) \geq 0$ for every node $\mathbf{s}$.

Determining the F_i's. Remember that $w_i(\mathbf{s})$ denotes the sum of the weights of the constraints that involve agent i and that are satisfied at profile $\mathbf{s}$. Note that $w_i(\mathbf{s}) = V - V(\mathbf{s}) + E - E(i) - E(\mathbf{s}) + E(\mathbf{s}, i)$. We next show that, for each agent i, there is a monotone non-decreasing function F_i such that the resulting game $\mathcal{G}'$ is better-response equivalent to $\mathcal{G}$. It suffices to show that, if $(\mathbf{s}, \mathbf{s}')$ is a dimension-i directed edge, then $F_i(w_i - w_i(\mathbf{s})) > F_i(w_i - w_i(\mathbf{s}'))$; if instead $(\mathbf{s}, \mathbf{s}')$ is a (good or bad) undirected edge, then $F_i(w_i - w_i(\mathbf{s})) = F_i(w_i - w_i(\mathbf{s}'))$.

Consider first a directed edge $(\mathbf{s}, \mathbf{s}')$ and let us denote by X and Y the super-nodes of the endpoints. By construction, we have $W(X) = V(\mathbf{s}) + E(\mathbf{s}) > V(\mathbf{s}') + E(\mathbf{s}') = W(Y)$ which, together with the observation that $E(\mathbf{s}, i) = E(\mathbf{s}', i)$, implies $w_i(\mathbf{s}) < w_i(\mathbf{s}')$. Therefore, the directed edge $(\mathbf{s}, \mathbf{s}')$ imposes the constraint that F_i is not constant in the interval $[w_i - w_i(\mathbf{s}'), w_i - w_i(\mathbf{s})]$.

If $(\mathbf{s}, \mathbf{s}')$ is a good edge, then $\mathbf{s}$ and $\mathbf{s}'$ belong to the same super-node and thus $V(\mathbf{s}) + E(\mathbf{s}) = V(\mathbf{s}') + E(\mathbf{s}')$, which implies that $w_i(\mathbf{s}) = w_i(\mathbf{s}')$.

So, it remains to consider a bad edge $(\mathbf{s}, \mathbf{s}')$ with $w_i(\mathbf{s}) < w_i(\mathbf{s}')$. Then, function F_i must be constant in the interval $[w_i - w_i(\mathbf{s}'), w_i - w_i(\mathbf{s})]$ and this is possible if and only if for every directed dimension-i edge $(\mathbf{z}, \mathbf{w})$, the interval $[w_i - w_i(\mathbf{w}), w_i - w_i(\mathbf{z})]$ is not entirely contained in $[w_i - w_i(\mathbf{s}'), w_i - w_i(\mathbf{s})]$. Indeed, we now prove that $w_i(\mathbf{z}) < w_i(\mathbf{s})$. Let X and Z be the super-nodes corresponding to $\mathbf{s}$ and $\mathbf{z}$, respectively. Then, $w_i(\mathbf{z}) - w_i(\mathbf{s}) = W(X) - W(Z) - E(\mathbf{s}, i) + E(\mathbf{z}, i) = W(X) - W(Z) - 2^n$, where we have used that $(\mathbf{s}, i)$ is a bad edge (hence, $E(\mathbf{s}, i) = 2^n$) and $(\mathbf{z}, i)$ is a directed edge (hence, $E(\mathbf{z}, i) = 0$). Since $W(X) - W(Z) \leq \mathtt{dist}(X, Z) \leq 2^n - 1$, we have $w_i(\mathbf{z}) < w_i(\mathbf{s})$, as desired. The proof of Theorem 4 is now complete.

Let us add a technical remark on the above proof. Observe that there exist no or-constraint of cardinality $n-1$ for $n = 2$. However, the above proof can be adjusted in order to work with belief constraints in place of or-constraints of cardinality $n - 1$. Note that, in this case, we need to satisfy the further requirement according to which at most one belief constraint for each agent has a positive weight. However, from Claim 5, it follows that at most one bad edge can exist in this case. If no bad edge exists, the proof above does not require any change. If exactly one bad edge exists, we assume without loss of generality that 1 is the agent involved in this edge. Then, it can be verified that the proof above holds even if agent 2 assigns weight 0 to both her belief constraints, and agent 1 assigns weights 0 and M.

4.2 Impossibility of Isomorphism

Even if structural equivalence between two-strategy generalized ordinal potential games and GDP games is certainly important, one may still wonder whether a stronger result (in the sense of the following definition) is possible.

Definition 6. *Two games $\mathcal{G} = ([n], (S_i)_{i \in [n]}, (c_i)_{i \in [n]})$ and $\mathcal{G}' = ([n], (S_i)_{i \in [n]}, (c'_i)_{i \in [n]})$ with the same set of profiles S are called* isomorphic *if for all profiles $\mathbf{s} \in S$, $c_i(\mathbf{s}) = c'_i(\mathbf{s})$.*

In particular, given a two-strategy game that admits a generalized ordinal potential, is there an *isomorphic* generalized discrete preference game? Unfortunately, this is not possible. To demonstrate the impossibility, we will consider the following simple two-agent game with two strategies per agent:

	0	1
0	1,1	2,0
1	1,0	0,1

The first value in each cell is the cost of the row agent and the second one is the cost of the column agent. It can be easily verified that the following function is a generalized ordinal potential function for this game: $\Phi(00) = 3$, $\Phi(01) = 2$, $\Phi(11) = 1$ and $\Phi(10) = 0$. In order to build an isomorphic generalized discrete preference game we need to define a non-decreasing function F_1 so that the conditions $F_1(w_1 - w_1(01)) > F_1(w_1 - w_1(00))$ and $F_1(w_1 - w_1(10)) > F_1(w_1 - w_1(11))$ are satisfied. These conditions simply require that the row agent strictly prefers profile 00 to profile 01 and profile 11 to profile 10. From the monotonicity of F_1, it must be $w_i(01) < w_i(00)$ and $w_i(10) < w_i(11)$. Recall that $w(\mathbf{s})$ is the sum of the weights of all constraints satisfied at $\mathbf{s}$. For $b \in \{0, 1\}$, let $\beta_i(b)$ be the weight of the belief constraint of agent i if it is satisfied by b, and 0 otherwise. Since there are only two agents, each of them is involved in all constraints, besides the belief constraints of the other agent. Using this notation, the above inequalities can be expressed as

$$w(01) - \beta_2(1) < w(00) - \beta_2(0), \tag{3}$$

$$w(10) - \beta_2(0) < w(11) - \beta_2(1). \tag{4}$$

For the column agent, we need a function F_2 so that the conditions $F_2(w_2 - w_2(00)) > F_2(w_2 - w_2(01))$ and $F_2(w_2 - w_2(11)) > F_2(w_2 - w_2(10))$ are satisfied. The same arguments as above allow us to rewrite these requirements as

$$w(00) < w(01), \tag{5}$$

$$w(11) < w(10). \tag{6}$$

Hence, in order to satisfy (3) and (5) we need that $\beta_2(1) > \beta_2(0)$. But, in order to satisfy (4) and (6) we need that $\beta_2(0) > \beta_2(1)$. We have obtained the desired contradiction.

5 Epilogue

We believe that our findings could open new research lines for understanding the performances and the complexity of equilibria. Even though (approximate) equilibria and convergence issues have been extensively studied in exact potential games — e.g., see [Awerbuch *et al.*, 2008; Bhalgat *et al.*, 2010; Caragiannis *et al.*, 2014; 2011; Christodoulou *et al.*, 2012; Fabrikant *et al.*, 2004] — generalized ordinal potential games are much less understood. Whether generalized discrete preference games could play a role analogous to that of congestion games in this direction certainly deserves investigation.

We also remark that some of our results can be extended to more than two strategies by opportunely choosing a metric for the distances among opinions. Anyway, we refrain to introduce new, possibly unrealistic, metrics, and look forward for the definition of a natural and largely-approved metric.

Acknowledgments. This work was partially supported by COST Action IC1205 on "Computational Social Choice". Ioannis Caragiannis was partially supported by Caratheodory research grant E.114 from the University of Patras. Diodato Ferraioli was partially supported by the "GNCS – INdAM".

References

[Acemoglu and Ozdaglar, 2011] D. Acemoglu and A. Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.

[Anshelevich *et al.*, 2008] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

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

[Awerbuch *et al.*, 2008] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *EC*, pages 264–273, 2008.

[Balcan *et al.*, 2013] M.-F. Balcan, A. Blum, and Y. Mansour. Circumventing the price of anarchy: leading dynamics to good behavior. *SIAM Journal on Computing*, 42(1):230–264, 2013.

[Bhalgat *et al.*, 2010] A. Bhalgat, T. Chakraborty, and S. Khanna. Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In *EC*, pages 73–82, 2010.

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

[Bindel *et al.*, 2011] D. Bindel, J. Kleinberg, and S. Oren. How bad is forming your own opinion? In *FOCS*, pages 55–66, 2011.

[Bonzon *et al.*, 2006] E. Bonzon, M.C. Lagasquie-Schiex, J. Lang, and B. Zanuttini. Boolean games revisited. In *ECAI*, volume 141, pages 265–269, 2006.

[Caragiannis *et al.*, 2011] I. Caragiannis, A. Fanelli, N. Gravin, and A. Skopalik. Efficient computation of approximate pure Nash equilibria in congestion games. In *FOCS*, pages 532–541, 2011.

[Caragiannis *et al.*, 2014] I. Caragiannis, A. Fanelli, and N. Gravin. Short sequences of improvement moves lead to approximate equilibria in constraint satisfaction games. In *SAGT*, pages 49–60, 2014.

[Chierichetti *et al.*, 2013] F. Chierichetti, J. Kleinberg, and S. Oren. On discrete preferences and coordination. In *EC*, pages 233–250, 2013.

[Christodoulou *et al.*, 2012] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27, 2012.

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

[Easley and Kleinberg, 2010] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.

[Even-Dar *et al.*, 2007] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3), 2007.

[Fabrikant *et al.*, 2004] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC*, pages 604–612, 2004.

[Ferraioli *et al.*, 2012] D. Ferraioli, P. W. Goldberg, and C. Ventre. Decentralized dynamics for finite opinion games. In *SAGT*, pages 144–155, 2012.

[Friedkin and Johnsen, 1990] N. E. Friedkin and E. C. Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3–4):193–205, 1990.

[Harrenstein *et al.*, 2014] P. Harrenstein, P. Turrini, and M. Wooldridge. Hard and soft equilibria in boolean games. In *AAMAS*, pages 845–852. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

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

[Mossel and Tamuz, 2014] E. Mossel and O. Tamuz. Opinion exchange dynamics. *arXiv preprint arXiv:1401.4770*, 2014.

[Peleg, 2002] D. Peleg. Local majorities, coalitions and monopolies in graphs: A review. *Theoretical Computer Science*, 282:231–257, 2002.

[Pryymak *et al.*, 2012] O. Pryymak, A. Rogers, and N. R. Jennings. Efficient opinion sharing in large decentralised teams. In *AAMAS*, pages 543–550, 2012.

[Tsang and Larson, 2014] A. Tsang and K. Larson. Opinion dynamics of skeptical agents. In *AAMAS*, pages 277–284, 2014.

[Wong *et al.*, 2014] L. Z. Wong, T. Q. S. Quek, and M. Padilla. Ordinal potential functions for network selection in heterogeneous wireless networks. In *ICASSP*, pages 6122–6126, 2014.

[Wooldridge *et al.*, 2013] M. Wooldridge, U. Endriss, S. Kraus, and J. Lang. Incentive engineering for boolean games. *Artificial Intelligence*, 195:418–439, 2013.