Journal of Computer and System Sciences 93 (2018) 11–29

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

On discrete preferences and coordination ☆,☆☆

Flavio Chierichetti [a], Jon Kleinberg [b], Sigal Oren [c,*]

[a] Sapienza University of Rome, Italy
[b] Cornell University, United States
[c] Ben-Gurion University of the Negev, Israel

ARTICLE INFO

Article history:
Received 20 March 2017
Received in revised form 14 September 2017
Accepted 1 November 2017
Available online 22 November 2017

Keywords:
Algorithmic game theory
Price of stability

ABSTRACT

An active line of research has considered games played on networks in which payoffs depend on both a player's individual decision and the decisions of her neighbors. A basic question that has remained largely open is to consider games where the players' strategies come from a fixed, discrete set, and where players may have different preferences among the possible strategies. We develop a set of techniques for analyzing this class of games, which we refer to as *discrete preference games*. We parametrize the games by the relative extent to which a player takes into account the effect of her preferred strategy and the effect of her neighbors' strategies, allowing us to interpolate between network coordination games and unilateral decision-making. We focus on the efficiency of the best Nash equilibrium and provide conditions on when the optimal solution is also a Nash equilibrium.

1. Introduction

People often make decisions in settings where the outcome depends not only on their own choices, but also on the choices of the people they interact with. A natural model for such situations is to consider a game played on a graph that represents an underlying social network, where the nodes are the players. Each node's personal decision corresponds to selecting a strategy, and the node's payoff depends on the strategies chosen by itself and its neighbors in the graph [2,4,14].

Coordination and internal preferences A fundamental class of such games involves payoffs based on the interplay between *coordination* — each player has an incentive to match the strategies of his or her neighbors — and *internal preferences* — each player also has an intrinsic preference for certain strategies over others, independent of the desire to match what others are doing. Trade-offs of this type come up in a very broad collection of situations, and it is worth mentioning several that motivate our work here.

- In the context of opinion formation, a group of people or organizations might each possess different internal views, but they are willing to express or endorse a "compromise" opinion so as to be in closer alignment with their network neighbors.

☆ This work has been supported in part by a Simons Investigator Award, a Google Research Grant, an ARO MURI grant, and NSF grants IIS-0910664, CCF-0910940, IIS-1016099 and a Microsoft Research Fellowship.
☆☆ A preliminary version of this paper was published at EC'13.

* Corresponding author.
E-mail addresses: flavio@di.uniroma1.it (F. Chierichetti), kleinber@cs.cornell.edu (J. Kleinberg), orensi@cs.bgu.ac.il (S. Oren).

https://doi.org/10.1016/j.jcss.2017.11.002

- Questions involving technological compatibility among firms tend to have this trade-off as a fundamental component: firms seek to coordinate on shared standards despite having internal cost structures that favor different solutions.
- Related to the previous example, a similar issue comes up in cooperative facility location problems, where firms have preferences for where to locate, but each firm also wants to locate near the firms with which it interacts.

A line of work beginning in the mathematical social sciences has considered versions of this question — often motivated by the first class of examples above, concerned with opinion formation — where the possible strategies correspond to a continuous space such as $\mathbb{R}^d$ [6,10]. This makes it possible for players to adopt arbitrarily fine-grained “average” strategies from among any set of options, and most of the dynamics and equilibrium properties of such models are driven by this type of averaging. In particular, dynamics based on repeated averaging have been shown in early work to exhibit nice convergence properties [6], and more recent work including by two of the authors has developed bounds on the relationship between equilibria and social optima [1].

Discrete preferences In many settings that exhibit a tension between coordination and individual preferences, however, there is no natural way to average among the available options. Instead, the alternatives are drawn from a fixed discrete set — for example, there is only a given set of available technologies for firms to choose among, or a fixed set of political candidates to endorse or vote for. On a much longer time scale, there is always the possibility that additional options could be created to interpolate between what’s available, but on the time scale over which the strategic interaction takes place, the players must choose from among the discrete set of alternatives that is available.

Among a small fixed set of players, coordination with discrete preferences is at the heart of a long line of games in the economic theory literature — perhaps the most primitive example is the classic *Battle of the Sexes* game, based on a pedagogical story in which one member of a couple wants to see movie A while the other wants to see movie B, but both want to go to a movie together. This provides a very concrete illustration of a set of payoffs in which the (two) players have (i) conflicting internal preferences (A and B respectively), (ii) an incentive to arrive at a compromise, and (iii) no way to “average” between the available options.

But essentially nothing is known about the properties of the games that arise when we consider such a payoff structure in a network context. Even the direct generalization of Battle of the Sexes (BoS) to a graph is more or less unexplored in this sense — each node plays a copy of BoS on each of its incident edges, choosing a single strategy A or B for use in all copies, incurring a cost from miscoordination with neighbors and an additional fixed cost when the node’s choice differs from its inherent preference. Indeed, as some evidence of the complexity of even this formulation, note that the version in which each node has an intrinsic preference for A is equivalent to the standard network coordination game, which already exhibits rich graph-theoretic structure [14]. And beyond this, of course, lies the prospect of such games with larger and more involved strategy sets.

Formalizing discrete preference games In this paper, we develop a set of techniques for analyzing this type of discrete preference games on a network, and establish tight bounds on the efficiency of the best Nash equilibrium[1] for several important families of such games.

To formulate a general model for this type of game, we start with an undirected graph $G = (V, E)$ representing the network on the players, and an underlying finite set L of strategies. Each player $i \in V$ has a *preferred strategy* $s_i \in L$, which is what i would choose in the absence of any other players. Finally, there is a metric $d(\cdot,\cdot)$ on the strategy set L — that is, a distance function satisfying (i) $d(i,i) = 0$ for all i, (ii) $d(i, j) = d(j,i)$ for all i, j, and (iii) $d(i, j) \le d(i,k) + d(k, j)$ for all i, j and k. For $i, j \in L$, the distance $d(i, j)$ intuitively measures how “different” i and j are as choices; players want to avoid choosing strategies that are at large distance from either their own internal preference or from the strategies chosen by their neighbors.

Each player’s objective is to minimize her cost: for a fixed parameter $\alpha \in [0, 1)$, the cost to player i when players choose the strategy vector $z = \langle z_j : j \in V\rangle$ is

$$c_i(z) = \alpha \cdot d(s_i, z_i) + \sum_{j\in N(i)} (1-\alpha) \cdot d(z_i, z_j),$$

where $N(i)$ is the set of neighbors of i in G. The parameter α essentially controls the extent to which players are more concerned with their preferred strategies or their network neighbors; we will see that the behavior of the game can undergo qualitative changes as we vary α.

We say that the above formulation defines a *discrete preference game*. As standard in game theory, we will be interested in Nash equilibria of this game. These are strategy profiles in which no player can reduce its cost by deviating to a different strategy. Note that the network version of Battle of the Sexes described earlier is essentially the special case in which $|L| = 2$, and network coordination games are the special case in which $|L| = 2$ and $\alpha = 0$, since then players are only concerned with matching their neighbors. The case in which $d(\cdot,\cdot)$ is the distance metric among nodes on a path is also interesting to

[1] Formally, we prove bounds on the price of stability which we later define.

focus on, since it is the discrete analogue of the one-dimensional space of real-valued opinions from continuous averaging models [1,6] — consider for example the natural scenario in which a finite number of discrete alternatives in an election are arranged along a one-dimensional political spectrum. The case of two strategies and $\alpha = 1/2$ is one that was studied in [5]. The focus of [5] was mainly the speed of convergence to equilibrium under various dynamics.

We also note that discrete preference games belong to the well-known framework of graphical games, which essentially consist of games in which the utility of every player depends only on the actions of its neighbors in a network. The interested reader is referred to the relevant chapter in [15] and the references within. In this context, Gottlob et al. proposed a generalization of Battle of the Sexes (BoS) to a graphical setting [9], but their formulation was much more complex than our starting point, with their questions correspondingly focused on existence and computational complexity, rather than on the types of performance guarantees we will be seeking.

For any discrete preference game, we will see that it is possible to define an exact potential function, and hence these games possess pure Nash equilibria [13].

Price of stability in discrete preference games We can also ask about the *social cost* of a strategy vector $z = \langle z_j : j \in V \rangle$, defined as the sum of all players' costs:

$$c(z) = \sum_{i \in V} \alpha \cdot d(s_i, z_i) + 2 \sum_{(i,j) \in E} (1-\alpha) \cdot d(z_i, z_j).$$

We refer to a strategy profile y minimizing the social cost function as an optimal solution. We note that the problem of minimizing the social cost is an instance of the *metric labeling problem*, in which we want to assign labels to nodes in order to minimize a sum of per-node costs and edge separation costs [3,12].

Since an underlying motivation for studying this class of games is the tension between preferred strategies and agreement on edges, it is natural to study its consequences on the social cost via the ratio between the social cost of a Nash equilibrium and the social cost of the optimal solution. Two particular measures that are commonly studied are the price of anarchy (the ratio between the social cost of the worst Nash equilibrium and the optimal solution) and the price of stability (the ratio between the social cost of the best Nash equilibrium and the optimal solution). The price of anarchy is in fact too severe a measure for this class of games; indeed, as we discuss in the next section, it is already unbounded for the well-studied class of network coordination games that our model contains as a special case.

We therefore consider the price of stability, which turns out to impart a rich structure to the problem. The price of stability is also natural in terms of the underlying examples discussed earlier as motivation; in most of these settings, it makes sense to propose a solution — for example, a compromise option in a political setting or a proposed set of technology choices for a set of interacting firms — and then to see if it is stable with respect to equilibrium.

Overview of results As a starting point for reference, observe that network coordination games (where players are not concerned with their preferred strategies) clearly have a price of stability of 1: the players can all choose the same strategy and achieve a cost of 0. But even for a general discrete preference game with two strategies — i.e. Battle of the Sexes on a network — the price of stability is already more subtle, since the social optimum may have a more complex structure (as a two-label metric labeling problem, and hence a minimum cut problem).

We begin by giving tight bounds on the maximum possible price of stability in the two-strategy case as a function of the parameter α. The dependence on α has a complex non-monotonic character; in particular, the price of stability is equal to 1 for all instances if and only if $\alpha \le 1/2$ or $\alpha = 2/3$, and more generally the price of stability as a function of α displays a type of "saw-tooth" behavior with infinitely many local minima in the interval $[0, 1]$. Our analysis uses a careful scheduling of the best-response dynamics so as to track the updates of players toward a solution with low social cost.

Above we also mentioned the distance metric of a path as a case of interest in opinion formation. We show that when $\alpha \le 1/2$, the price of stability for instances based on such metrics is always 1, by proving the stronger statement that in fact the price of stability is always 1 for any discrete preference game based on a tree metric. Our analysis for tree metrics involves considering how players' best responses lie at the medians of their neighbors' strategies in the metric, and then developing combinatorial techniques for reasoning about the arrangement of these collections of medians on the underlying tree.

Like path metrics, tree metrics are also relevant to motivating scenarios in terms of opinion formation, when individuals classify the space of possible opinions according to a hierarchical structure rather than a linear one. To take one example of this, consider students choosing a major in college, where each student has an internal preference and an interest in picking a major that is similar to the choices of her friends. The different subjects roughly follow a hierarchy — on top we might have science, engineering, and humanities; under science we can have for example biology, physics, and other areas; and under biology we can have subjects including genetics and plant breeding. This setting fits our model since each person has some internal inclination for a major, but still it is arguably the case that a math major has more in common in her educational experience with her computer science friends than with her friends in comparative literature.

The two families of instances described above (two strategies and tree metrics) both have price of stability equal to 1 when $\alpha \le 1/2$. But the price of stability can be greater than 1 for more general metrics when $\alpha \le 1/2$. It is not hard to show (as we do in the next section) that the price of stability is always at most 2 for all α, and we match this bound by constructing and analyzing examples, based on perturbations of uniform metrics, showing that the price of stability can be arbitrarily close to 2 when $\alpha = 1/2$.

Finally, we consider a generalization of our model of discrete preference games for $\alpha = 1/2$, which we term an *anchored preference game*. Suppose that nodes are partitioned into two types: there are *fixed nodes* i that have a preferred strategy anchored on a particular value s_i, and there are *strategic nodes* that have no preferred strategy, so the cost of such a node i is purely the term $\sum_{j \in N(i)} d(z_i, z_j)$. Only the strategic nodes choose strategies, and only they take part in the definition of the equilibrium. For $\alpha = 1/2$ this generalizes the main model we consider because of the following reduction: given an instance of a discrete preference game, we can take each node i in the instance and make it a strategic node in the generalized instance by eliminating its preferred strategy s_i, and adding a new fixed node i' to the instance that has preferred strategy s_i and is connected only to node i by an edge (i, i'). In this way i', which is non-strategic, plays the role of i's preferred strategy.

Anchored preference games are also of interest in their own right, and not just as a generalization of discrete preference games, since they can model settings where certain parts of the network represent unmodifiable constraints — for example, how technological compatibility might have to take into account certain unchangeable factors, or how individuals adapting their opinions to each other might also be taking into account agents such as media sources that are in effect outside the immediate strategic environment. We generalize our results for tree metrics to the case of anchored preference games, parameterizing the price of stability by the maximum number of fixed nodes among the neighbors of any strategic node.

2. Preliminaries

Recall that in a discrete preference game played on a graph $G = (V, E)$ with strategy set L, each player $i \in V$ has a preferred strategy $s_i \in L$. The cost incurred by player i when all players choose strategies $z = \langle z_j : j \in V \rangle$ is

$$c_i(z) = \alpha \cdot d(s_i, z_i) + \sum_{j \in N(i)} (1 - \alpha) \cdot d(z_i, z_j).$$

The social cost of the game is the sum of all the players' costs:

$$c(z) = \sum_{i \in V} \alpha \cdot d(s_i, z_i) + 2 \sum_{(i,j) \in E} (1 - \alpha) \cdot d(z_i, z_j).$$

Another quantity that is useful to define is the contribution of player i to the social cost — by this we quantify both the cost player i is exhibiting and the cost it is inflicting on its neighbors:

$$c(z_i | z_{-i}) = \alpha \cdot d(s_i, z_i) + 2 \sum_{j \in N(i)} (1 - \alpha) \cdot d(z_i, z_j).$$

As is standard, we denote by z_{-i} the strategy vector z without the ith coordinate. We note that for the optimal solution y we have that for every player i the strategy y_i minimizes $c(y_i | y_{-i})$.

It is not hard to show that discrete preference games are potential games. This means that a pure Nash equilibrium always exists:

Claim 2.1. *Discrete preference games always admit a pure Nash equilibrium.*

Proof. We will show that the following function is an exact potential function:

$$\phi(z) = \alpha \sum_{i \in V} d(z_i, s_i) + (1 - \alpha) \sum_{(i,j) \in E} d(z_i, z_j).$$

To see why, note that: $\phi(z_i, z_{-i}) - \phi(z'_i, z_{-i}) =$

$$\alpha \cdot d(z_i, s_i) + (1 - \alpha) \sum_{j \in N(i)} d(z_i, z_j) - \Big(\alpha \cdot d(z'_i, s_i) + (1 - \alpha) \sum_{j \in N(i)} d(z'_i, z_j)\Big)$$
$$= c_i(z_i, z_{-i}) - c_i(z'_i, z_{-i}).$$

Thus, discrete preference games are potential games and hence a pure Nash equilibrium always exists. □

We show that the class of discrete preference games includes instances for which the price of anarchy (PoA) is unbounded. The idea here is to devise instances in which all the players share the same preferred strategy, however there exists an equilibrium in which all the players play a different strategy. This type of equilibrium, has a natural interpretation in our motivating contexts. In technology adoption, it corresponds to convergence on a standard that no firm individually

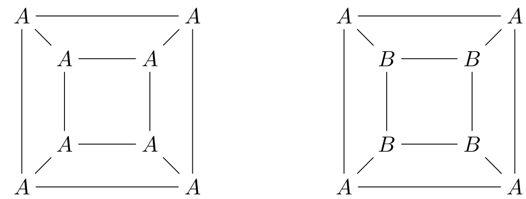

(a) *An optimal solution .* (b) *A Nash equilibrium .*

Fig. 1. An instance illustrating that the PoA can be unbounded even when the players do not have a preferred strategy (i.e., $\alpha = 0$).

wants, but which is hard to move away from once it has become the consensus. In opinion formation, it corresponds essentially to a kind of "superstitious" belief that is universally expressed, and hence is hard for people to outwardly disavow even though they prefer an alternate opinion. In the next claim, we formalize this intuition to construct for every value of $\alpha < 1$, an instance for which the PoA is unbounded.

Claim 2.2. *For any $\alpha < 1$ there exists an instance for which the price of anarchy is unbounded.*

Proof. Assume the strategy space contains two strategies A and B, such that $d(A, B) = 1$. For any $0 < \alpha < 1$ we consider a clique of size $\lceil \frac{\alpha}{1-\alpha} \rceil + 1$ in which all players' preferred strategy is A and show it is an equilibrium for all the players to play strategy B. To see why, observe that if the rest of the players play strategy B, then player i's cost for playing strategy A is $(1 - \alpha) \cdot \lceil \frac{\alpha}{1-\alpha} \rceil$ which is at least α. Since the cost of player i for playing strategy B is α we have that it is an equilibrium for all players to play strategy B. The PoA of such an instance is unbounded as the cost of the equilibrium in which all players play strategy B is strictly positive but the cost of the optimal solution is 0.

To show that the PoA can be unbounded for $\alpha = 0$, a slightly different instance is required, which will be familiar from the literature of network coordination games. When players do not have a preference the optimal solution is clearly for all players to play the same strategy, as such a solution has a cost of 0. However, Fig. 1 depicts an instance for which there exists a Nash equilibrium in which not all the players play the same strategy and hence the cost of this equilibrium is strictly positive. □

We note that the worst equilibrium in the previous instances is not a strong equilibrium.[2] Thus, if all the players could coordinate a joint deviation to strategy A they can all benefit. A natural question is what happens if we restrict ourselves to the worst strong Nash equilibrium (strong PoA), in which a simultaneous deviation by a set of players is allowed. Unfortunately, as the next claim demonstrates the strong PoA can still be quite high:

Claim 2.3. *For any $\alpha < 1$ there exists an n-player instance for which the strong price of anarchy is $\Theta(n)$.*

Proof. Let $k = \lfloor \frac{\alpha}{\alpha-1} \rfloor$ and consider a clique of size $n >> k$ in which $k+1$ players prefer strategy B and the rest of the players prefer strategy A. Consider the strategy profile in which all players play strategy B. The social cost of this strategy profile is $\alpha \cdot (n - k - 1)$ as each of the $n - k - 1$ players with preferred strategy A exhibits a cost of α for playing strategy B. The social cost of the optimal solution in which all the players play strategy A is $\alpha(k + 1)$. To see why this strategy profile is a strong Nash equilibrium, we first observe that no player with preferred strategy B can take part in a deviating coalition as their current cost is 0 and their cost in any other strategy profile will be positive. Since all the players with preferred strategy B cannot be a part of any deviating coalition. Given this, the cost of any player (with preferred strategy A) participating in a deviating coalition is at least $(1 - \alpha) \cdot (k + 1)$. By our choice of k this is greater than α, the player's cost in the strategy profile where all the players play B. □

As we just showed both the PoA and the strong PoA can be very high, and hence for the remainder of the paper we focus on the qualities of the best Nash equilibrium, trying to bound the price of stability (PoS). We begin by showing that the price of stability is bounded by 2. This is done by using the exact potential function we presented in Claim 2.1

Claim 2.4. *The price of stability is upper bounded by 2.*

Proof. Note that by Claim 2.1 the following function is an exact potential function for discrete preference games:

$$\phi(z) = \alpha \sum_{i \in V} d(z_i, s_i) + (1 - \alpha) \sum_{(i,j) \in E} d(z_i, z_j).$$

[2] A strategy profile is a strong Nash equilibrium if there is no subset of players S that can coordinate a deviation and by doing so strictly increase the utility of at least a single player in the coalition while keeping the utility of the rest of the players the same.

Denote by x the global minimizer of the potential function and by y the optimal solution. By definition x is an equilibrium and it provides a 2-approximation to the optimal social cost since $c(x) \leq 2\phi(x) \leq 2\phi(y) \leq 2c(y)$. □

Lastly, we provide some conditions under which the price of stability is 1. These conditions will be used later in the paper where we identify cases in which the price of stability is 1. Denote by $C_i(z)$ and $SC_i(z)$ the strategies of player i that minimize $c_i(z)$ and $c(z_i|z_{-i})$ respectively. We show that if for every player i the intersection of the two sets $C_i(z)$ and $SC_i(z)$ is always non-empty then the price of stability is 1:

Claim 2.5. *If for every player i and strategy vector z, $SC_i(z) \cap C_i(z) \neq \emptyset$, then $PoS = 1$.*

Proof. Let Y denote the set of optimal solutions and recall the potential function $\phi(\cdot)$ used in the proof of Claim 2.4. Now, consider the optimal solution y for which the value of the potential function is minimal (i.e., $y = \arg\min_{y' \in Y} \phi(y')$). If y is also a Nash equilibrium then we are done. Else, there exists a node i that can strictly reduce its cost by performing a best response. By our assumption, node i has a best response strategy x_i, such that $x_i \in SC_i(y) \cap C_i(y)$. The fact that $x_i \in SC_i(y)$ implies that the change in strategy of player i does not affect the social cost. Therefore, (x_i, y_{-i}) is also an optimal solution and $\phi(y) > \phi(x_i, y_{-i})$, in contradiction to the assumption that y is the optimal solution minimizing $\phi(\cdot)$. □

3. The case of two strategies: battle of the sexes on a network

We begin by considering the subclass of instances in which the players only have two different strategies A and B. Without loss of generality we assume that $d(A, B) = 1$. Recall that s_i is the preferred strategy of player i. We denote by $N_j(i)$ the set of i's neighbors playing strategy j and by $\bar{s}_i$ the strategy opposite to s_i (for example, if $s_i = A$ then $\bar{s}_i = B$). When the strategy space contains only two strategies, a player's best response is to pick a strategy which is the weighted majority of its own preferred strategy and the strategies played by its neighbors. The next two observations formalize this statement and a similar statement regarding a player's strategy minimizing the social cost:

Observation 3.1. Strategy s_i minimizes player i's cost ($c_i(z)$) if and only if:

$$(1-\alpha)N_{\bar{s}_i}(i) \leq \alpha + (1-\alpha)N_{s_i}(i) \Longrightarrow N_{\bar{s}_i}(i) \leq \frac{\alpha}{1-\alpha} + N_{s_i}(i).$$

Where $(1-\alpha)N_{\bar{s}_i}(i)$ is the cost of player i for playing strategy s_i and $\alpha + (1-\alpha)N_{s_i}(i)$ is player i's cost for playing strategy $\bar{s}_i$.

Observation 3.2. Strategy s_i minimizes the contribution of player i to the social cost ($c(z_i|z_{-i})$) if and only if:

$$2(1-\alpha)N_{\bar{s}_i}(i) \leq \alpha + 2(1-\alpha)N_{s_i}(i) \Longrightarrow N_{\bar{s}_i}(i) \leq \frac{\alpha}{2(1-\alpha)} + N_{s_i}(i).$$

Where $2(1-\alpha)N_{\bar{s}_i}(i)$ is the contribution of player i to the social cost for playing strategy s_i and $\alpha + 2(1-\alpha)N_{s_i}(i)$ is player i's contribution to the social cost for playing strategy $\bar{s}_i$.

By the two observations below we get the following simple but useful corollary:

Corollary 3.3. *If s_i minimizes $c(z_i|z_{-i})$ then it also minimizes $c_i(z)$.*

We present a simple best response order that results in a Nash equilibrium after a linear number of best responses. We will later see how this order can be used to bound the PoS.

Lemma 3.4. *Starting from some initial strategy vector, the following best response order results in a Nash equilibrium:*

1. *While there exists a player that can reduce its cost by changing its strategy to A, let it do a best response. Once there is no such player continue to step 2.*
2. *While there exists a player that can reduce its cost by changing its strategy to B, let it do a best response.*

Proof. To see why the resulting strategy vector is a Nash equilibrium, first observe that any player that by step 2 any player that plays strategy A clearly minimizes its cost by doing so. Now, consider some player i that plays strategy B. Denote by $N^j_A(i)$ and $N^j_B(i)$ the number of i's neighbor playing at the end of step j of the algorithm strategy A and B respectively. Since at step 2 players only switch from strategy A to strategy B we have that: $N^2_B(i) \geq N^1_B(i)$ and $N^2_A(i) \leq N^1_A(i)$. Note that at the end of step 1 of the algorithm player i minimized its cost by playing strategy B. This means that depending on the preferred strategy of player i either $\alpha + (1-\alpha)N^1_B(i) \leq (1-\alpha)N^1_A(i)$ or $(1-\alpha)N^1_B(i) \leq \alpha + (1-\alpha)N^1_A(i)$. By the observation that $N_B^2(i) \geq N_B^1(i)$ and $N_A^2(i) \leq N_A^1(i)$ we have that at the end of step 2 either $\alpha + (1-\alpha)N_B^2(i) \leq (1-\alpha)N_A^2(i)$ or $(1-\alpha)N_B^2(i) \leq \alpha + (1-\alpha)N_A^2(i)$ and hence player i still minimizes its cost by playing strategy B. □

Next, we characterize the values of α for which the price of stability is 1.

Claim 3.5. *For $\alpha \leq \frac{1}{2}$ or $\alpha = \frac{2}{3}$, any instance admits an optimal solution which is also a Nash equilibrium ($PoS = 1$).*

Proof. Recall that we use $C_i(z)$ and $SC_i(z)$ to denote the strategies of player i that minimize $c_i(z)$ and $c(z_i|z_{-i})$ respectively. By Claim 2.5 we have that if for every strategy vector z it holds that $SC_i(z) \cap C_i(z) \neq \emptyset$ then the price of stability is 1. By Corollary 3.3 we have that if $s_i \in SC_i(z)$ then $s_i \in C_i(z)$ and hence the claim holds. We are left with showing that for $\alpha \leq \frac{1}{2}$ or $\alpha = \frac{2}{3}$, if $SC_i(z) = \bar{s}_i$ then $\bar{s}_i \in C_i(z)$. Since in this case $\bar{s}_i$ is the unique minimizer of the social cost function, Observation 3.2 implies that:

$$N_{\bar{s}_i}(i) > \frac{\alpha}{2(1-\alpha)} + N_{s_i}(i)$$

On the other hand, by Observation 3.1 player i is satisfied with playing $\bar{s}_i$, if:

$$N_{\bar{s}_i}(i) \geq \frac{\alpha}{1-\alpha} + N_{s_i}(i)$$

Observe that for $\alpha \leq 1/2$ both $\frac{\alpha}{2(1-\alpha)}$ and $\frac{\alpha}{1-\alpha}$ are less than or equal to 1. As $N_{\bar{s}_i}(i)$ and $N_{s_i}(i)$ are clearly integers this implies that whenever $N_{\bar{s}_i}(i) > \frac{\alpha}{2(1-\alpha)} + N_{s_i}(i)$ it is also the case that $N_{\bar{s}_i}(i) \geq \frac{\alpha}{1-\alpha} + N_{s_i}(i)$. Furthermore, for the case of $\alpha = 2/3$, we have that $\frac{\alpha}{2(1-\alpha)} = 1$ implying that $N_{\bar{s}_i}(i) \geq 2 + N_{s_i}(i)$ and $\frac{\alpha}{1-\alpha} = 2$ hence again y_i is also player i's best response. □

As we will later see, for $\frac{1}{2} < \alpha < 1$ and $\alpha \neq \frac{2}{3}$ it is not always that case that the optimal solution is also a Nash equilibrium. In the following theorem we compute the ratio between the optimal solution and a Nash equilibrium for values of α in the interval above. This is done by performing the sequence of best responses that Lemma 3.4 prescribes and analyzing the social cost of the resulting Nash equilibrium.

Theorem 3.6. *For $\frac{1}{2} < \alpha < 1$, $PoS \leq 2\left\lceil \frac{\alpha}{1-\alpha} - 1 \right\rceil \cdot \frac{1-\alpha}{\alpha}$.*

Proof. Let x be the equilibrium achieved by the sequence described in Lemma 3.4 starting from an optimal solution y. We assume that a player performs a best response only when it can strictly decrease its cost by doing so, thus, since $|L| = 2$, we only consider cases in which the player's best response is unique. In the following Lemma we bound the increase in the social cost inflicted by the players' unique best responses:

Lemma 3.7. *Let x_i be player i's* unique *best response when the rest of the players play z_{-i}:*

1. *If $x_i = \bar{s}_i$ then $c(\bar{s}_i, z_{-i}) - c(s_i, z_{-i}) \leq \alpha - 2(1-\alpha)\left\lfloor \frac{\alpha}{1-\alpha} + 1 \right\rfloor$.*
2. *If $x_i = s_i$ then $c(s_i, z_{-i}) - c(\bar{s}_i, z_{-i}) \leq -\alpha + 2(1-\alpha)\left\lceil \frac{\alpha}{1-\alpha} - 1 \right\rceil$.*

Proof. Notice we are only considering the effect player i's strategy has on the social cost, thus we have that: $c(\bar{s}_i, z_{-i}) - c(s_i, z_{-i}) = c(\bar{s}_i|z_{-i}) - c(s_i|z_{-i})$. This implies that:

$$\begin{aligned} c(\bar{s}_i, z_{-i}) - c(s_i, z_{-i}) &= \alpha + 2(1-\alpha)N_{s_i}(i) - 2(1-\alpha)N_{\bar{s}_i}(i) \\ &= \alpha - 2(1-\alpha)(N_{\bar{s}_i}(i) - N_{s_i}(i)). \end{aligned}$$

For proving statement (1) we observe that since $\bar{s}_i$ is player i's unique best response then by Observation 3.1 we have that $N_{\bar{s}_i}(i) > \frac{\alpha}{1-\alpha} + N_{s_i}(i)$. Since, $N_{\bar{s}_i}(i)$ and $N_{s_i}(i)$ are integers this implies that $N_{\bar{s}_i}(i) - N_{s_i}(i) \geq \lfloor \frac{\alpha}{1-\alpha} + 1 \rfloor$ and the bound is achieved.

For proving statement (2), observe that $c(s_i, z_{-i}) - c(\bar{s}_i, z_{-i}) = -\alpha + 2(1-\alpha)(N_{\bar{s}_i}(i) - N_{s_i}(i))$. Now since s_i is player i's best response we have by Observation 3.1 that: $N_{\bar{s}_i}(i) < \frac{\alpha}{1-\alpha} + N_{s_i}(i)$. Similarly to the previous bound this implies that $N_{\bar{s}_i}(i) - N_{s_i}(i) \leq \lceil \frac{\alpha}{1-\alpha} - 1 \rceil$ as required. □

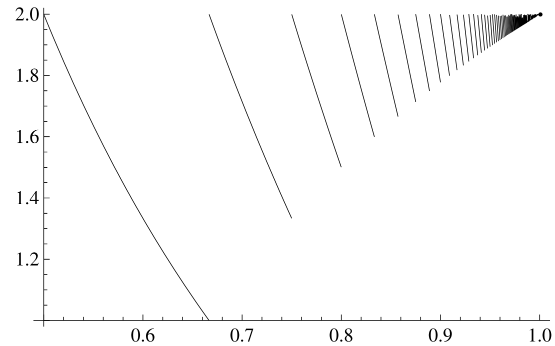

Fig. 2. The tight upper bound on the PoS for two strategies as a function of α for the range $\frac{1}{2} < \alpha < 1$.

Notice that by statement (1) of Lemma 3.7 a node that changes its strategy to a strategy different than its preferred strategy can only reduce the social cost. Also, note that if a node changed its strategy in step 1 to A and in step 2 to B its total contribution to the social cost is non-positive. The reason is that in one of these changes the player changed its strategy from s_i to $\bar{s_i}$ and in the other from $\bar{s_i}$ to s_i. The effect of these two changes on the social cost sums up to at most $2(1-\alpha)\left(\lceil\frac{\alpha}{1-\alpha}\rceil - \lfloor\frac{\alpha}{1-\alpha}\rfloor - 2\right) \le 0$. Thus we can ignore such changes as well.

The only nodes that are capable of increasing the social cost by performing a best response are nodes that play in the optimal solution a different strategy than their preferred strategy ($y_i = \bar{s}_i$). By definition the number of such nodes equals exactly $\sum_{i \in V} d(y_i, s_i)$ as $d(y_i, s_i) = 1$ if $y_i \ne s_i$ and 0 otherwise. Statement (2) of Lemma 3.7 guarantees us that each of these nodes can increase the social cost by at most $-\alpha + 2(1-\alpha)\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil$. Thus, we get the following bound:

$$\begin{aligned}
c(x) &\le c(y) + \left(-\alpha + 2(1-\alpha)\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil\right)\sum_{i\in V} d(y_i, s_i) \\
&= \alpha\sum_{i\in V} d(y_i, s_i) + 2(1-\alpha)\sum_{(i,j)\in E} d(y_i, y_j) + \left(-\alpha + 2(1-\alpha)\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil\right)\sum_{i\in V} d(y_i, s_i) \\
&= 2(1-\alpha)\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil\sum_{i\in V} d(y_i, s_i) + 2(1-\alpha)\sum_{(i,j)\in E} d(y_i, y_j) \\
&\le 2\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil \cdot \frac{1-\alpha}{\alpha} \cdot \left(\alpha\sum_{i\in V} d(y_i, s_i) + 2(1-\alpha)\sum_{(i,j)\in E} d(y_i, y_j)\right). \\
&= 2\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil \cdot \frac{1-\alpha}{\alpha} \cdot c(y)
\end{aligned}$$

Thus, we get that the price of stability is $2\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil \cdot \frac{1-\alpha}{\alpha}$ as required. □

It is interesting to take a closer look at the upper bound on the PoS we computed (as we will see later this bound is tight). In Fig. 2 we plot the upper bound on the PoS as a function of α. We can see that as α approaches 1 the PoS approaches 2 and also that for any $k \ge 2$, as ε approaches 0, the PoS of $\alpha = \frac{k-1}{k} + \varepsilon$ also approaches to 2. This uncharacteristic saw-like behavior of the PoS originates from the fact that for every value of α the maximal PoS is achieved by a star graph. This is proved in the following claim.

Claim 3.8. *For any $1/2 < \alpha < 1$, $\alpha \ne 2/3$ there exists an instance achieving a price of stability of $2\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil \cdot \frac{1-\alpha}{\alpha}$.*

Proof. Consider a star consisting of $\left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil$ peripheral nodes that prefer strategy A and a central node that prefers strategy B. It is easy to see that in any Nash equilibria all the peripheral nodes play their preferred strategy A as $\alpha > (1-\alpha)$. In the optimal solution the central node plays strategy A for a cost of α. However, this is not a Nash equilibrium since for playing strategy B it exhibits a cost of $(1-\alpha) \cdot \left\lceil\frac{\alpha}{1-\alpha} - 1\right\rceil < \alpha$. Thus, the central node prefers to play its preferred strategy B and the total cost of the unique Nash equilibrium is $2(1-1\alpha)\left\lceil \frac{\alpha}{1-\alpha} - 1 \right\rceil$ while the cost of the optimal solution is only α. $\square$

Corollary 3.9. *As n goes to infinity the price of stability of a star with n nodes for $\alpha = \frac{n}{n+1}$ is approaching 2.*

4. Tree metrics

In the previous section we have seen that even when there are only two strategies in the game (Battle of the Sexes on a network), for at least some values of $\alpha > \frac{1}{2}$, the PoS can be close to 2. These bounds carry over to larger strategy spaces since an instance can always use only two strategies from the strategy space. However, for $\alpha \le \frac{1}{2}$ the PoS for the Battle of the Sexes on a network is 1, so a natural question is how bad the PoS can be once we have more strategies in the space. This is the question we deal with for the rest of the paper.

We begin by considering the case in which the distance function on the strategy set is a tree metric, defined as the shortest-path metric among the nodes in a tree. (As such, tree metrics are a special case of graphic metrics, in which there is a graph on the elements of the space and the distance between every two elements is defined to be the length of the shortest path between them in the graph.) We show that if the distance function is a tree metric then the price of stability is 1 for any rational $\alpha \le \frac{1}{2}$. Formally, we will show that:

Theorem 4.1. *If the distance metric is a tree metric then for rational $\alpha \le \frac{1}{2}$, there exists an optimal solution which is also a Nash equilibrium ($PoS = 1$).*

Recall that we use $C_i(z)$ and $SC_i(z)$ to denote the strategies of player i that minimize $c_i(z)$ and $c(z_i|z_{-i})$ respectively. By Claim 2.5 we have that if for every strategy vector z it holds that $SC_i(z) \cap C_i(z) \neq \emptyset$ then the price of stability is 1. Our goal now is to show that the conditions of Claim 2.5 hold for a tree metric when $\alpha \le \frac{1}{2}$. Our first step is relating the strategies that consist of a player's best response (or social cost minimizer) and the set of medians of a node-weighted tree.

Definition 4.2 *(Medians of a tree).* Given a tree T where the weight of node v is denoted $w(v)$, the set of T's medians is $M(T) = \arg\min_{u \in V}\{\sum_{v \in V} w(v) \cdot d(u,v)\}$.

Definition 4.3. Given a network G, a tree metric T, a strategy vector z, a player i and non-negative integers q and r, we construct a tree $T_{i,z}(q,r)$ with the same edges as T and the following node weights:

- $w(s_i) = q + r \cdot |\{j \in N(i) | z_j = s_i\}|$ for the node representing player i's preferred strategy.
- $w(v) = r \cdot |\{j \in N(i) | z_j = v\}|$ for every other node v.

Next we show that for $\alpha = \frac{a}{a+b}$, every player i and strategy vector z, it holds that $M(T_{i,z}(a,b)) = C_i(z)$. To see why, observe that by construction we have $M(T_{i,z}(a,b)) =$

$$\arg\min_{u \in V}\left\{\left(a + b \cdot |\{j \in N(i)|z_j = s_i\}|\right) \cdot d(u,s_i) + \sum_{v \neq s_i \in V} b \cdot |\{j \in N(i)|z_j = v\}| \cdot d(u,v)\right\}$$

$$= \arg\min_{u \in V}\left\{a \cdot d(u,s_i) + b \sum_{j \in N(i)} d(u,z_j)\right\} = C_i(z).$$

Similarly, it is easy to show that $M(T_{i,z}(a,2b)) = SC_i(z)$. Thus, to show that $SC_i(z) \cap C_i(z) \neq \emptyset$ it is sufficient to show that $T_{i,z}(a,b)$ and $T_{i,z}(a,2b)$ share a median. This is done by using the following proposition:

Proposition 4.4. *Let T_1 and T_2 be two node-weighted trees with integer weights and the same edges and nodes, then:*

1. *If there exists a node v, such that for every $u \neq v \in V$, we have $w_1(u) = w_2(u)$ and for v we have $|w_1(v) - w_2(v)| = 1$, then T_1 and T_2 share a median.*
2. *If T_1 and T_2 share a median, then it is also a median of their union $T_1 \cup T_2$. Where the union of T_1 and T_2 is a tree with the same nodes and edges where the weight of node v is $w_{1+2}(v) = w_1(v) + w_2(v)$.*

Proof. The proof builds on the highly tractable structure of medians in trees developed in early work; see [7,8,11] and the references therein.

In particular Kariv and Hakimi [11] proved the following useful claim (for completeness we provide a proof of this claim in Appendix A):

Claim 4.5. *Consider a tree T and let $w(V) = \sum_{v \in V} w(v)$ denote the weight of the tree. A node u is a median of T if and only if the weight of each connected component of $T - u$ is at most $w(V)/2$.*

Let $w_1(V)$ and $w_2(V)$ be the sum of weights of all nodes in T_1 and T_2 respectively. For the first statement, we distinguish between two cases based on the parity of $w_1(V)$:

- $w_1(V)$ is odd. Let u be a median of T_1, then, in this case by Claim 4.5 the weight of each component in $T_1 - u$ is at most $w_1(V)/2 - 1/2$. Thus for the same median u in T_2 the weight of each component in $T_2 - u$ is at most $w_1(V) + 1/2 = w_2(V)/2$. Therefore by Claim 4.5 u is still a median.
- $w_1(V)$ is even. This implies that $w_2(V) = w_1(V)+1$ is odd. Consider a median u' of T_2. Then the size of each connected component in $T_2 - u'$ is at most $w_1(V)/2 + 1/2$, since $w_1(V)$ is even, this implies that the weight of each connected component is bounded by $w_1(V)/2$ and therefore u' is also a median of T_1.

The proof of the second statement is pretty straight forward: let u be a median of both T_1 and T_2. This implies that in T_1 the weight of every connected component in $T_1 - u$ is at most $w_1(T_1)/2$ and in T_2 the weight of every connected component in $T_2 - u$ is at most $w_2(T_2)/2$. Hence, in $T_1 \cup T_2$ the weight of every connected component in $T_1 \cup T_2 - u$ is at most $w_1(T_1)/2 + w_2(T_2)/2 = w_{1+2}(T_1 \cup T_2)/2$. Thus, u is also a median of $T_1 \cup T_2$. □

Given Proposition 4.4 we can now show that $T_{i,z}(a, b)$ and $T_{i,z}(a, 2b)$ share a median:

Lemma 4.6. *For $\alpha = \frac{a}{a+b} \le \frac{1}{2}$, every player i and strategy vector z, $M(T_{i,z}(a, b)) \cap M(T_{i,z}(a, 2b)) \neq \emptyset$.*

Proof. We first handle the case that $a < b$ (i.e., $\frac{a}{a+b} < \frac{1}{2}$) and then handle the case that $a = b$. Observe that by statement 1 of Proposition 4.4 we have that $T_{i,z}(0, 1)$ and $T_{i,z}(1, 1)$ share a median. As medians are invariant to scaling this implies that $T_{i,z}(0, b-a)$ and $T_{i,z}(a, a)$ also share a median. Next, by statement 2 of Proposition 4.4 we have that any median they share is also a median of their union $T_{i,z}(a, b)$; let us denote this median by u. As medians are invariant to scaling we have that since u is the median of $T_{i,z}(0, b-a)$ it is also the median of $T_{i,z}(0, b)$. Now, as u is a median of $T_{i,z}(0, b)$ and $T_{i,z}(a, b)$, it is also a median of $T_{i,z}(a, 2b)$ by applying Proposition 4.4 again.

For the case of $a = b$ we need to show that $T_{i,z}(a, a)$ and $T_{i,z}(a, 2a)$ share a median. Observe that $T_{i,z}(a, a)$ and $T_{i,z}(0, a)$ share a median since $T_{i,z}(0, 1)$ and $T_{i,z}(1, 1)$ share a median. Now, by statement 2 of Proposition 4.4 we have that the median they share is also the median of their union $T_{z,z}(a, a)$ which completes the proof. □

With this we conclude the proof of Theorem 4.1.

5. Lower bounds on the price of stability in non-tree metrics

In some sense tree metrics are the largest class of metrics for which the optimal solution is always a Nash equilibrium for $\alpha \le \frac{1}{2}$. The next example demonstrates that even when the distance metric is a simple cycle the PoS can be as high as $\frac{4}{3}$ for $\alpha = \frac{1}{2}$.

Example 5.1. Consider a metric which is a cycle of size $3k + 1$, for some integer $k \ge 1$. Let A, B, C be three strategies in this strategy space such that $d(A, B) = k$, $d(A, C) = k$ and $d(B, C) = k + 1$. Consider an instance composed of a path of three nodes v_1, v_2 and v_3 such that each of the nodes v_1 and v_3 is part of a different clique of size $3k$. The preferred strategies of all the nodes in the clique of v_1 is B, the preferred strategy of node v_2 is A and the preferred strategy of all the nodes in the clique of node v_3 is C and let $\alpha = \frac{1}{2}$.

Consider the strategy profile in which all the nodes play their preferred strategies. Note that in this case, it is clearly an equilibrium for each of the nodes in the cliques to play its preferred strategy as its cost for switching a strategy will be at least $\frac{1}{2}k + \frac{1}{2}k(3k-1)$ and its current cost is at most $\frac{1}{2}k$ (the only nodes exhibiting this cost are v_1 and v_3). Finally, consider v_2, its cost for playing A is $\frac{1}{2} \cdot 2k = k$. On the other hand, its cost for playing any other strategy x which is between A and B (including B) on the cycle is

$$\frac{1}{2}(d(x, A) + d(x, B) + d(x, C)) = \frac{1}{2}\Big(d(A, B) + min\{d(x, A) + d(A, C), d(x, B) + d(B, C)\}\Big),$$

which is greater than k. Similarly one can show that the central player prefers strategy A over any strategy x.

The total cost of the Nash equilibrium is $2k$. Note that this is the best Nash equilibrium since the cost of any solution in which some of the nodes in a clique play a strategy different than their preferred strategy is at least $3k$. In the optimal solution the central node should play strategy B (or C) for a total cost of $(1/2)k + 2 \cdot (1/2)(k+1) = (3/2)k + 1$. Thus we have that the price of stability approaches $4/3$ as k approaches infinity.[3]

Note that this lower bound of $4/3$ is achieved on an instance in which the lowest-cost Nash equilibrium and the socially optimal solution differ only in the strategy choice of a single player. In Appendix B we show that in such cases, where the difference between these two solutions consists of the decision of just a single player, $4/3$ is the maximum possible price of stability for $\alpha = \frac{1}{2}$. More generally, for these cases, we show an upper bound of $\dfrac{2}{2-\alpha}$ for $\alpha < \frac{1}{2}$. We leave open the question of whether or not this is a tight bound for $\alpha < 1/2$.

As Illustrated by Example 5.1 even in very simple non-tree metrics, the price of stability can be greater than 1. We now give a set of stronger lower bounds, using a more involved family of constructions. First, we give an asymptotically tight lower bound of 2 when $\alpha = \frac{1}{2}$, and then we adapt this construction to give non-trivial lower bounds for all $0 < \alpha < \frac{1}{2}$.

5.1. Price of stability for $\alpha = \frac{1}{2}$

The following instance illustrates that the PoS for $\alpha = \frac{1}{2}$ can be arbitrarily close to 2. The network we consider is composed of a path of n nodes and each of the end points is connected to a single node of a clique of size n^2. We assume that the preferred strategy of node i on the path is s_i, the preferred strategy of all nodes in the leftmost clique is s_0, and the preferred strategy of all nodes in the rightmost clique is s_{n+1}. The following is a sketch of the network:

Since all the s_i's are distinct we use them also as names for the different possible strategies. To get a lower bound on the price of stability, we compare between the cost of two solutions: (i) a solution in which there is exactly one edge such that its two endpoints play different strategies, and (ii) the solution in which each player is playing its preferred strategy. We refer to the first solution as a *bi-consensus solution*. The cost of a bi-consensus solution is an upper bound on the optimal solution. We choose the distance metric such that the gap between the costs of the two solutions is close to 2 and that solution (ii) is the best Nash equilibrium. The distance metric we choose is a perturbed uniform metric. The fact that the metric is almost uniform enables to get a gap of at most 2 between the costs of the two solutions while the perturbations serve only to make sure that solution (ii) is indeed the best Nash equilibrium. We formally define the distance metric as follows: for $i > j$, we have $d(s_i, s_j) = 1 + (i - j - 1)\varepsilon$. (When $i < j$, we simply use $d(s_i, s_j) = d(s_j, s_i)$.)

In Claim 5.2 below we show that indeed in the best Nash equilibrium all players play their preferred strategies. As by definition $d(s_i, s_{i+1}) = 1$, the cost of this equilibrium is $c(s) = \frac{1}{2} \cdot 2 \sum_{i=0}^{n} d(s_i, s_{i+1}) = n + 1$. On the other hand, consider the bi-consensus solution in which for some node i all the nodes up till node i choose strategy s_0 and all the nodes from node $i+1$ choose strategy s_{n+1}. The cost of such assignment is

$$\frac{1}{2}\sum_{j=1}^{i} d(s_0, s_j) + 2 \cdot \frac{1}{2} \cdot d(s_0, s_{n+1}) + \frac{1}{2}\sum_{j=i+1}^{n} d(s_j, s_{n+1}) = n + \varepsilon \cdot \left(\sum_{j=1}^{i} j + \sum_{j=1}^{i}(n-j)\right) + (1 + n\varepsilon)$$
$$= \frac{1}{2}n + 1 + \varepsilon \cdot O(n^2)$$

Therefore for $\varepsilon = n^{-3}$ as n goes to infinity the price of stability goes to 2. We now show that the strategy profile in which each player plays its preferred strategy is indeed the best Nash equilibrium:

Claim 5.2. *In the previously defined instance the best Nash equilibrium is for each player to play its preferred strategy.*

Proof. We first show that in the best Nash equilibrium all the nodes in the two cliques should play their preferred strategy. Consider for example the clique with the preferred strategy s_0, it is easy to see that regardless of which strategies all the nodes outside this clique play it is an equilibrium for all the nodes in the clique to play s_0. Thus, to show that in the best equilibrium the players in the clique play s_0 it suffices to show that the contribution of the nodes of the clique to the social cost is minimized when they all play s_0. To see why this is the case, assume that $n^2 - r$ players play the strategy s_0 and r players play some other strategy. Observe that in this case the contribution of the nodes in the clique to the social cost is at least $\frac{1}{2}r + r \cdot (n^2 - r)$. This is because the r nodes exhibit a cost of at least $\frac{1}{2}$ for playing a different strategy than their preferred strategy and there are at least $r \cdot (n^2 - r)$ edges that the players at their endpoints play different strategies. On the other hand, if all nodes in the left clique play s_0 their total contribution to the social cost will be at most $1 + n\varepsilon$ since they only node contributing to the social cost is the one connected to the path and the cost associated with this edge is at

[3] The example can be easily extended to general values of α to achieve a lower bound of $\frac{4(1-\alpha)}{4-5\alpha}$ on the PoS.

most $\frac{1}{2}(1+n\varepsilon)$ for each of its endpoints. As $1+n\varepsilon < \frac{1}{2}r + r\cdot(n^2-r)$ for any $r \geq 1$ we have that indeed in the best Nash equilibrium all players in each clique play the same strategy.

Next, we show that in any equilibrium in which all the nodes in the cliques play their preferred strategies the rest of the nodes in the graph play their preferred strategies as well. To this end, assume towards a contradiction that there is an equilibrium in which one of the players plays a strategy different than its preferred strategy. Pick the rightmost node i that plays a strategy s_j such that $j < i$. If such a node does not exist pick the leftmost node i that plays a strategy s_j such that $j > i$. In either case we reach a contradiction by applying the following Lemma:

Lemma 5.3. *Let s_a and s_b be the strategies played by player i's neighbors such that $a \leq b$. If $a \leq i \leq b$, then player i's best response is to play strategy s_i.*

Proof. We first note that the cost i exhibits for playing a strategy $j \notin \{i, a, b\}$ is $3 + O(n\cdot\varepsilon)$ whereas the cost of playing the preferred strategy is $2 + O(n\cdot\varepsilon)$. Hence player i will always play a strategy $j \in \{i, a, b\}$.

Next, observe that player i prefers to play strategy s_i over strategy $s_a \neq s_i$ whenever $d(s_i, s_a) + d(s_i, s_b) < d(s_i, s_a) + d(s_a, s_b)$ implying that $d(s_i, s_b) < d(s_a, s_b)$. This condition holds according to our assumptions since $1 + (b-i-1)\varepsilon = d(s_i, s_b) < d(s_a, s_b) = 1 + (b-a-1)\varepsilon$. For the same reason, player i prefers strategy s_i over $s_b \neq s_i$ since $1 + (i-a-1) = d(s_i, s_a) < d(s_a, s_b) = 1 + (b-a-1)$ under the lemma's assumptions. □

Finally, to see that a contradiction is reached, consider the case in which there exists a player i that prefers a strategy s_j such that $j < i$ and pick player i to be the rightmost such player. Denote by s_a and s_b the strategies that its left and right neighbors are playing. By Lemma 5.3 as in equilibrium player i plays strategy $s_j \neq s_i$, we have that either $a, b > i$ or $a, b < i$. Clearly, it will be a best response for player i to play a strategy s_j such that $j < i$ only in the latter case where $a, b < i$. In particular, this implies that player $i+1$ is also playing a strategy $j' < i+1$ in contradiction to the assumption that player i is the rightmost such player. □

5.2. Extension for $\alpha < \frac{1}{2}$

We extend the construction in Section 5.1 to $0 < \alpha < \frac{1}{2}$ by defining the following metric: for $i > j$, let $d(s_i, s_j) = 1 + (i-j-1)\left(\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right)$. We consider the same family of instances defined in Section 5.1 with the small change that we increase the size of the cliques to $\lceil n^2/\alpha \rceil$. While this extension is valid for any value of $\alpha < \frac{1}{2}$, as we will later see, the lower bounds it produces are getting more loose as α is getting closer to 0. Nonetheless, the main merit of this extension is that it qualitatively shows that as α is approaching $1/2$, the price of stability is approaching 2.

The source of inefficiency for our bounds for small values of α is that as α is getting closer to 0 the perturbations to the uniform distance metric are getting more and more pronounced. For example, already for $\alpha = 1/3$ we consider perturbations of scale of $1/2$. Roughly speaking, these large perturbations significantly increase the cost of a bi-consensus solution making the bi-consensus solution better than the best Nash equilibrium only for very small values of n (for example, for $\alpha \leq \frac{1}{3}$ we have that $n = 2$). As instances with smaller size n have an inherent smaller gap between the costs of a bi-consensus solution and the best Nash equilibrium, the resulting lower bound on the price of stability is fairly loose and establishing tighter lower bound for smaller values of α remains an open question.

We now turn to showing that in the best Nash equilibrium for any instance constructed by the extended construction each player plays its preferred strategy:

Claim 5.4. *In every instance of the extended family the best Nash equilibrium is for each player to play its preferred strategy.*

The proof operates very similarly to the proof of Claim 5.2. The main difference is in the proof of the Lemma analogous to Lemma 5.3 which becomes more involved. For this reason, we only provide the proof of the Lemma:

Lemma 5.5. *Let s_a and s_b be the strategies played by player i's neighbors such that $a \leq b$. If $a \leq i \leq b$, then player i's best response is to play strategy s_i.*

Proof. For this proof we use an equivalent distance function which will turn out to be easier to work with: $d(s_i, s_j) = 1 + (|i-j|-1)\left(\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right)$. Also, we only present the proof for the case that $a < i < b$, the proof for the cases that $a = i$ or $a = b$ is very similar. Denote, i's best response by s_j. As in Lemma 5.3, we first show that player i prefers to play strategy s_i over playing any strategy s_j such that $j \neq a, b$. The cost of player i for playing s_i is:

$$(1-\alpha)\left(2 + (|i-a|-1+|b-i|-1)\cdot\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right) = 2(1-\alpha) + (|b-a|-2)\cdot(1-2\alpha)(1+\varepsilon)$$

The cost of playing strategy s_j such that $j \neq a, b$ is:

$$\alpha(1 + (|i-j|-1)\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)) + (1-\alpha)\left(2 + (|j-a|-1+|j-b|-1)\cdot\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right)$$
$$\geq 2-\alpha+(|b-a|-2)\cdot(1-2\alpha)(1+\varepsilon).$$

The last transition is due to the fact that by the triangle inequality $|j-a|+|j-b| \geq |b-a|$. Thus we conclude that player i's best response can only be s_i, s_a or s_b. Next we consider strategies s_a and s_b. By writing the cost of playing each one of these strategies it is easy to see that these costs are greater than the costs for playing s_i.

The cost of playing strategy $s_a \neq s_i$ is:

$$\alpha(1 + (|i-a|-1)\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)) + (1-\alpha)\left(1 + (|b-a|-1)\cdot\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right)$$
$$= 1 + (\frac{\alpha}{1-\alpha}(|i-a|-1)+|b-a|-1)\cdot(1-2\alpha)(1+\varepsilon).$$
$$= 2(1-\alpha)+\varepsilon\cdot(1-2\alpha)+\left(\frac{\alpha}{1-\alpha}(|i-a|-1)+|b-a|-2\right)\cdot(1-2\alpha)(1+\varepsilon).$$

The cost of playing strategy $s_b \neq s_i$ is:

$$\alpha(1 + (|b-i|-1)\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)) + (1-\alpha)\left(1 + (|b-a|-1)\cdot\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)\right)$$
$$= 1 + (\frac{\alpha}{1-\alpha}(|b-i|-1)+|b-a|-1)\cdot(1-2\alpha)(1+\varepsilon).$$
$$= 2(1-\alpha)+\varepsilon\cdot(1-2\alpha)+(\frac{\alpha}{1-\alpha}(|b-i|-1)+|b-a|-2)\cdot(1-2\alpha)(1+\varepsilon)$$

This concludes the proof as we have shown that player i's best response it to play its preferred strategy. □

5.2.1. Lower bound on the price of stability

As we mentioned earlier, to get a lower bound on the price of stability, we would like to simulate the technique we used in the proof for $\alpha = \frac{1}{2}$ to compare between the cost of a bi-consensus solution and the cost of the best Nash equilibrium. Recall that the cost of a bi-consensus solution is an upper bound on the optimal solution, and hence computing the ratio between the best Nash equilibrium and a best bi-consensus solution gives a lower bound on the PoS achieved by instances defined above.

Observe that in the best bi-consensus solution nodes $i \in [1 \ldots \lfloor n/2 \rfloor]$ play strategy s_0 and nodes $i \in [\lfloor n/2+1 \rfloor \ldots n]$ play strategy s_{n+1}. Intuitively, this is the best bi-consensus solution as it equalizes between the costs of nodes playing s_0 and the costs of nodes playing s_{n+1}. We denote this solution by b and compute its cost as follows:

$$c(b) = \alpha\Big(\sum_{i=1}^{\lfloor n/2\rfloor}(1+(i-1)(\frac{1-2\alpha}{1-\alpha}(1+\varepsilon))) + \sum_{i=\lfloor n/2\rfloor+1}^{n}(1+(n-i)(\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)))\Big)$$
$$+ 2(1-\alpha)(1+(n+1-0-1)(\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)))$$
$$\leq \alpha\cdot n + \alpha\frac{1}{4}(n-1)^2\cdot\frac{1-2\alpha}{1-\alpha}(1+\varepsilon)+2(1-\alpha)+2n(1-2\alpha)(1+\varepsilon).$$

Where the last transition is due to the fact that:

$$\sum_{i=1}^{\lfloor n/2\rfloor}(i-1)+\sum_{i=\lfloor n/2\rfloor+1}^{n}(n-i) \leq \sum_{i=1}^{\lfloor n/2\rfloor-1} i + \sum_{i=1}^{\lfloor n/2\rfloor} i = (\lfloor n/2\rfloor-1)\cdot\lfloor n/2\rfloor+\lfloor n/2\rfloor$$
$$= \lfloor n/2\rfloor^2 = \frac{1}{4}(n-1)^2.$$

The cost of the best Nash equilibrium x is simply $c(x) = 2(1-\alpha)(n+1)$. As we previously observed, as α gets closer to 0 the perturbations to the uniform distance metric are getting more pronounced and the distance metric resembles more the L1 metric. As a result, the cost of the bi-consensus solution increases and hence the largest gap between a bi-consensus solution and the best Nash equilibrium is achieved for an intermediate value of n. By taking the first derivative of the function $\frac{c(x)}{c(b)}$ with respect to n and comparing it to 0, we get that the maximum PoS is achieved for $n = \lceil\frac{2\alpha+2\sqrt{2-7\alpha+6\alpha^2}-1}{1-2\alpha}\rceil$ or $n = \lfloor\frac{2\alpha+2\sqrt{2-7\alpha+6\alpha^2}-1}{1-2\alpha}\rfloor$.

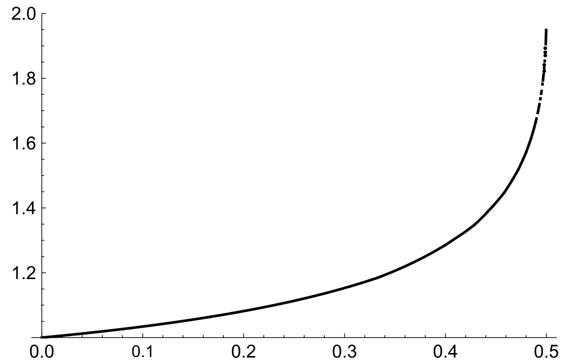

Fig. 3. The PoS achievable by a path.

In Fig. 3 we plot the lower bound on the PoS that can be achieved by these instances. We see that indeed as α approaches $1/2$ the PoS approaches 2.

6. The anchored preference game

In this final section, we consider the following extension of discrete preference games, with $\alpha = \frac{1}{2}$. We assume that nodes are partitioned into two types: F and S. Nodes in F are *fixed nodes* that always play their preferred strategy. The cost that a node $i \in F$ incurs for playing strategy s_i is 0. The nodes in S are *strategic nodes* that have no preferred strategy, so the cost of a node $i \in S$ is purely the term $\sum_{j \in N(i)} d(z_i, z_j)$. We use $F(i)$ and $S(i)$ to denote the fixed and strategic neighbors of node i respectively. Only the strategic nodes choose strategies, and only they take part in the definition of the equilibrium. The social cost for this model is:

$$c(z) = \sum_{\substack{(i,j) \in E; \\ i \in S; j \in F}} d(z_i, s_j) + 2 \sum_{\substack{(i,j) \in E; \\ i,j \in S}} d(z_i, z_j).$$

As noted in the introduction, this model, which we call an *anchored preference game*, generalizes the model from the previous sections (for $\alpha = 1/2$) because of the following reduction: given an instance of a discrete preference game, we can take each node i in the instance and make it a strategic node in the generalized instance by eliminating its preferred strategy s_i, and adding a new fixed node i' to the instance that has preferred strategy s_i and is connected only to node i by an edge (i, i'). In this way i', which is non-strategic, plays the role of i's preferred strategy.

Given that discrete preference games with $\alpha = 1/2$ are (due to the reduction) a special case of anchored preference games with exactly one fixed neighbor per node, it becomes natural to study the price of stability of anchored preference games parametrized by k, the maximum number of fixed nodes adjacent to any strategic nodes. In the next claim we generalize our result on tree metrics by showing that the price of stability for anchored preference games with tree metrics is 1 provided $k \leq 2$. The proof is essentially a generalization of Theorem 4.1.

Proposition 6.1. *If the distance function is a tree metric and $k \leq 2$, then the optimal solution is also a Nash equilibrium.*

Proof. Similar to the proof of Theorem 4.1, we define $C_i(z)$ and $SC_i(z)$ to be the strategies of player i that minimize $c_i(z) = \sum_{j \in F(i)} d(z_i, s_j) + \sum_{j \in S(i)} d(z_i, z_j)$ and $c(z_i | z_{-i}) = \sum_{j \in F(i)} d(z_i, s_j) + 2 \sum_{j \in S(i)} d(z_i, z_j)$ respectively. Our goal now is to show that the set of best responses $C(i)$ and the set of local improvements $SC(i)$ always intersect. By Claim 2.5 this implies that the PoS is always 1.

Observe that $C_i(z) \cap SC_i(z) \neq \varnothing$ when the number of fixed neighbors node i has is 0 (by definition) and when the number of fixed neighbors node i has is 1 (by Lemma 4.6). Thus, the only case that is left to handle is when player i has exactly two fixed neighbors. We denote their preferred strategies by s_1 and s_2. Just as in the proof of Theorem 4.1 we show that $C_i(z)$ and $SC_i(z)$ coincide with the set of medians for some trees we define next. For this purpose we define the tree $T_{i,z}(q_1, q_2, r)$:

Definition 6.2. Given a metric tree T, a strategy vector z and a strategic player i we denote by $T_{i,z}(q_1, q_2, r)$ the tree with the same nodes and edges as T and the following node weights:

- $w(s_1) = q_1 + r \cdot |\{j \in S(i) | z_j = s_1\}|$ for the preferred strategy of the first fixed neighbor.
- $w(s_2) = q_2 + r \cdot |\{j \in S(i) | z_j = s_2\}|$ for the preferred strategy of the second fixed neighbor.
- $w(v) = r \cdot |\{j \in S(i) | z_j = v\}|$ for every other node v.

It is not hard to see that for a player i with two fixed neighbors $C(i) = M(T_{i,z}(1,1,1))$ and $SC(i) = M(T_{i,z}(1,1,2))$. Recall that by Claim 4.5 the set of medians and separators of trees coincide. Therefore, it suffices to show that the set of separators of $T_{i,z}(1,1,1)$ and $T_{i,z}(1,1,2)$ intersect. By Proposition 4.4 we have that $T_{i,z}(2,1,2)$ and $T_{i,z}(1,1,2)$ share a separator u. We will next show that u is also a separator of $T_{i,z}(2,2,2)$. As separators are invariant to scaling this implies that u is a separator of $T_{i,z}(1,1,1)$ as well, proving that indeed $T_{i,z}(1,1,1)$ and $T_{i,z}(1,1,2)$ intersect.

We now show that *any* node u which is a separator of $T_{i,z}(2,1,2)$ is a separator of $T_{i,z}(2,2,2)$. Observe that the sum of weights in $T_{i,z}(2,1,2)$ is odd; this implies that the size of each connected component of $T_{i,z}(2,1,2) - u$ is at most $w(T_{i,z}(2,1,2))/2 - 1/2$. After increasing the weight of a single node by exactly 1 the weight of each connected component is at most $w(T_{i,z}(2,1,2))/2 + 1/2 = w(T_{i,z}(2,2,2))/2$. Thus, u is also a separator of $T_{i,z}(2,2,2)$. □

Lastly, we use Proposition 6.1 to show that for tree metrics and $k > 2$ the PoS is bounded by $\frac{2(k-1)}{k}$. The idea behind this upper-bound is for each strategic node to pick some of its fixed neighbors ($k-2$ to be exact) and replace each fixed node with preferred strategy s_j with a large clique of strategic nodes each is connected to a single fixed node with preferred strategy s_j. By replacing some of a node fixed neighbors with strategic neighbors, we create a new instance such that the number of fixed neighbors each node has is at most 2. This implies we can use Proposition 6.1 and get that the new instance admits an optimal solution which is also a Nash equilibrium. As the two instances are quite similar, it is easy to see that this same Nash equilibrium is also a Nash equilibrium of the original instance. Also, since the reduction increases the social cost of each solution by at most a factor of 2, we reach the desired bound of $\frac{2(k-1)}{k}$. We later show that this bound is tight.

Claim 6.3. *If the distance function is a tree metric and each strategic node has at most k fixed neighbors, then $PoS \le \frac{2(k-1)}{k}$. Furthermore, the bound is tight.*

Proof. Let y be the optimal solution for a graph G with a total of n nodes ($|F| + |S| = n$), and let $c_G(y)$ denote the social cost of this solution. We can assume without loss of generality that each fixed node has at most one strategic neighbor (if necessary we can simply duplicate fixed nodes). Denote by x the best Nash equilibrium for this instance.

We construct a new graph G' from G in which each strategic node will have at most two fixed neighbors. For each strategic node i in G, we pick the $|F(i)| - 2$ fixed neighbors of i with preferred strategies which are closest to y_i in distance. We replace each such fixed neighbor j with a clique of size $\lceil 2c_G(y) \rceil$ consisting of strategic nodes. Each of the strategic nodes in the clique is connected to a fixed node with preferred strategy s_j. We connect node i to one of the nodes in the clique. Since j had i as its only strategic neighbor in G, at the end of this process there is a single node from the original graph connected to each clique.

Let $c_{G'}(\cdot)$ be the social cost function for G'. Since each strategic node has at most two fixed neighbors in G', Proposition 6.1 implies that there exists an optimal solution y' for G' which is also a Nash equilibrium in G'. Observe that in any optimal solution of G' all the nodes in the new cliques we created play their preferred strategies, since the cost of any other strategy vector is at least $\lceil 2c_G(y) \rceil$.

We now define two further solutions. Let y'_G be a strategy vector for G in which the nodes common to G and G' play their strategies in y', and for each fixed node j replaced by a clique, j plays its preferred strategy. (Above we argued that all the nodes in the clique replacing j will play this same preferred strategy in G'.) We observe that y'_G is a well-defined solution to the anchored preference game on G, since all fixed nodes play their preferred strategy. Moreover, y'_G is an equilibrium in G. Therefore we have that $c_G(x) \le c_G(y'_G)$. Now, since $c_G(z) \le c_{G'}(z)$ and by construction y' and y'_G are practically the same, we have that $c_G(y'_G) \le c_{G'}(y')$.

Let $y_{G'}$ denote the strategy vector for G' that agrees with y on the nodes of G, and in which the nodes of each clique in G' play the preferred strategy of the fixed node j they replaced. Since y' is an optimal solution for G' we have that $c_{G'}(y') \le c_{G'}(y_{G'})$. Thus, we have that $c_G(x) \le c_{G'}(y_{G'})$.

For simplicity of presentation we define $F_i = \sum_{j \in F(i)} d(y_i, s_j)$ to be the cost of player i associated with its fixed neighbors in the graph G. Since in G' the cost of the $|F(i)| - 2$ fixed neighbors which are closest to y_i was doubled, it is now at most $(1 + \frac{|F(i)|-2}{|F(i)|})F_i$.

$$
\begin{aligned}
PoS(G) = \frac{c_G(x)}{c_G(y)} \le \frac{c_{G'}(y_{G'})}{c_G(y)} &\le \frac{\sum_{i\in S}\sum_{j\in S(i)} d(y_i, y_j) + \sum_{i\in S}(1 + \frac{|F(i)|-2}{|F(i)|}) \cdot F_i}{\sum_{i\in S}\sum_{j\in S(i)} d(y_i, y_j) + \sum_{i\in S} F_i} \\
&\le \frac{\sum_{i\in S}\sum_{j\in S(i)} d(y_i, y_j) + (1 + \frac{k-2}{k})\sum_{i\in S} \cdot F_i}{\sum_{i\in S}\sum_{j\in S(i)} d(y_i, y_j) + \sum_{i\in S} F_i} \\
&\le \frac{2k-2}{k}. \quad \square
\end{aligned}
$$

To see that the bound for $k > 2$ is tight, consider a star in which the central node i is connected to k fixed nodes that prefer strategy A. i is also connected to $k-1$ strategic nodes that are connected to one another (forming a clique of size $k-1$) and each one is connected to k fixed nodes that prefer strategy B. Observe that in the best Nash equilibrium node i

plays strategy A and the rest of the strategic nodes play strategy B. The social cost of this equilibrium is $2(k-1)$. However, in the optimal solution node i also plays strategy B which reduces the social cost to k.

7. Discussion

In this paper we study the price of stability of discrete preferences games. We show that for the special case of two strategies which is a generalization of the well known battle of the sexes game on a network, the price of stability for $\alpha \leq 1/2$ is 1. Interestingly, for $\alpha \geq 1/2$ the tight bound on the PoS has an uncharacteristic saw-like behavior with peaks in which the PoS is converging to 2 as α approaches $\frac{k-1}{k}$ for some integer value of k. The results for $\alpha \leq 1/2$ carry over to richer strategy spaces when the distance function on the strategy set is a tree metric. We do not know if it is still the case that for $\alpha \geq 1/2$ the price of stability exhibit this saw-like behavior or possibly for richer strategy spaces the PoS could be arbitrarily close to 2 for any value of $\alpha \geq 1/2$. For general tree metrics, we present a construction that achieves a price of stability arbitrarily close to 2 for $\alpha \leq 1/2$. However, we suspect that this extended construction for $\alpha < 1/2$ is not tight.

We demonstrate that even when all the nodes in the graph put the same weight on their own opinion (the same value of α) the price of stability changes considerably as we vary α and the distance metric. Nevertheless, it can be interesting to consider extensions such as weighted graphs or instances in which each player has a different value of α_i. For the case of $\alpha = 1/2$ and integer weighted graph, [5] show that it still the case that the PoS is 1. When every player has a different value of α_i it is not hard to construct examples in which the price of stability is arbitrary close to 3/2 even though for every i, $\alpha_i < 1/2$. The rough idea is to construct a star instance in which half of the peripheral nodes have preferred strategy A and have a very low value of α_i and half have preferred strategy B and have $\alpha = 1/2$. All peripheral nodes are connected to cliques of appropriate size to make sure that in any equilibrium they will play their preferred strategy. The central node has $\alpha_i = 1/2$ and preferred strategy B hence in the unique Nash equilibrium it will play B.

Our main focus in this paper was on the game theoretic properties of the game and the efficiency of equilibria. On the computational aspect, for two strategies we presented a linear time best response algorithm to compute a Nash equilibrium. Discrete preference games are instances of the metric labeling problem. It is known that the metric labeling problem is NP-hard even with a uniform distance metric on the labels [12]. This implies, for example, that in discrete preference games for $\alpha \leq 1/2$ and a tree-metric computing the best Nash equilibrium is NP-hard, as in this case we have showed that the best Nash equilibrium is also an optimal solution.

Acknowledgments

We thank Shahar Dobzinski for valuable comments.

Appendix A. On medians, separators and w-centroids in trees

Recall that we denote the weight of a tree by $w(V) = \sum_{v \in V} w(v)$. The following two related concepts were defined and studied in early work; see [7,8,11] and the references therein.

Definition A.1. A separator of a tree T is a node u such that the weight of each connected component of $T - u$ is at most $w(V)/2$.

Definition A.2. A w-centroid of a tree T is a node u that minimizes the size of the largest component of $T - u$.

Kariv and Hakimi [11] showed that a node u is a w-centroid of a tree T if and only if u is a separator of T. Then they used this to show that u is a w-centroid if and only if it is a median of T. Here we provide a proof showing that u is a median of T if and only if it is a separator of T:

Claim A.3. *A node u is a median of a tree T if and only if it is a separator of T.*

Proof. Let u be a median of a tree T, and assume towards a contradiction that it is not a separator; that is, there exists a component of $T - u$ of weight strictly greater than $w(V)/2$. Let v be the neighbor of u in this component. Consider locating the median at v. This reduces the distance to a total node weight of at least $w(V)/2 + 1/2$ by 1, and increases the distance to less than a total node weight of $w(V)/2$ by 1. Hence the sum of all distances decreases, and this contradicts the fact that u is a median. Thus, every median of a tree is also a separator.

To show that any separator is also a median, we show that for any two separators u_1 and u_2 it holds that $\sum_{v \in V} w(v) \cdot d(u_1, v) = \sum_{v \in V} w(v) \cdot d(u_2, v)$. Since, we know that there exists a median which is a separator this will imply that any separator is a median.

Denote by C_1 the subtree its root is u_1 and by C_2 the subtree its root is u_2. Observe that C_1 and C_2 are disjoint subtrees. Also, denote the connected component that includes the rest of the nodes in the tree by C. Since u_1 is a separator it holds that $w(C_2) + w(C) \leq w(V)/2$. This in turn implies that $w(C_1) \geq w(V)/2$. Similarly, since u_2 is a separator it holds that $w(C_1) + w(C) \le w(V)/2$. This in turn implies that $w(C_2) \ge w(V)/2$. Therefore, it has to be the case that $w(C) = 0$, $w(C_1) = w(C_2) = w(V)/2$. We next show this implies that $\sum_{v\in V} w(v)\cdot d(u_1, v) = \sum_{v\in V} w(v)\cdot d(u_2, v)$. Observe that:

$$\begin{aligned}\sum_{v\in V} w(v)\cdot d(u_1, v) &= \sum_{v\in C_1} w(v)\cdot d(u_1, v) + \sum_{v\in C_2} w(v)\cdot \big(d(u_1, u_2) + d(u_2, v)\big)\\ &= \sum_{v\in C_1} w(v)\cdot d(u_1, v) + \sum_{v\in C_2} w(v)\cdot d(u_2, v) + w(C_2)\cdot d(u_1, u_2),\end{aligned}$$

and similarly that:

$$\sum_{v\in V} w(v)\cdot d(u_2, v) = \sum_{v\in C_2} w(v)\cdot d(u_2, v) + \sum_{v\in C_1} w(v)\cdot d(u_1, v) + w(C_1)\cdot d(u_1, u_2).$$

Thus, we have that $\sum_{v\in V} w(v)\cdot d(u_1, v) = \sum_{v\in V} w(v)\cdot d(u_2, v)$ and the claim follows. □

Appendix B. Upper bound on the price of stability for a special case

In this section we provide an upper bound on the Price of Stability for instances in which the lowest-cost Nash equilibrium and the socially optimal solution differ only in the strategy choice of a single player. Recall, that this was the case in Example 5.1. We show that 4/3 is the maximum possible price of stability for $\alpha = \frac{1}{2}$. More generally we show that for these cases $\dfrac{2}{2-\alpha}$ is the maximum possible price of stability for $\alpha < \frac{1}{2}$.

By the definition of the model, a player's strategy only affects its cost and the cost of its neighbors. Recall that we denote this part of the social cost by $c(z_i|z_{-i}) = \alpha\cdot d(s_i, z_i) + 2(1-\alpha)\cdot \sum_{j\in N(i)} d(z_i, z_j)$. We now prove the following claim.

Claim B.1. *Let $\alpha \le \frac{1}{2}$. Fix an optimal solution y which is not a Nash equilibrium and let player i be a player that can reduce its cost by playing x_i. Then $\dfrac{c(x_i|y_{-i})}{c(y_i|y_{-i})} < \dfrac{2}{2-\alpha}$.*

Proof. Since x_i is player i's best response, then $\alpha\cdot d(s_i, x_i) + (1-\alpha)\sum_{j\in N(i)} d(x_i, y_j) < c_i(y)$. By rearranging the terms we get that $(1-\alpha)\sum_{j\in N(i)} d(x_i, y_j) < c_i(y) - \alpha\cdot d(s_i, x_i)$. This in turn implies that

$$\begin{aligned}c(x_i|y_{-i}) = \alpha\cdot d(s_i, x_i) + 2(1-\alpha)\sum_{j\in N(i)} d(x_i, y_j) &< \alpha\cdot d(s_i, x_i) + 2\big(c_i(y) - \alpha\cdot d(s_i, x_i)\big)\\ &= 2c_i(y) - \alpha\cdot d(s_i, x_i).\end{aligned}$$

Thus, we have that

$$\frac{c(x_i|y_{-i})}{c(y_i|y_{-i})} < \frac{2c_i(y) - \alpha\cdot d(s_i, x_i)}{c_i(y) + (1-\alpha)\cdot \sum_{j\in N(i)} d(y_i, y_j)}.$$

If $\sum_{j\in N(i)} d(y_i, y_j) \ge c_i(y)$ then $c(y_i|y_{-i}) \ge c_i(y) + (1-\alpha)c_i(y) = (2-\alpha)c_i(y)$ and the claim follows. Else, we show that $d(s_i, x_i) > d(s_i, y_i) - \sum_{j\in N(i)} d(y_i, y_j)$ in Lemma B.2 below; this in turn implies that

$$\begin{aligned}c(x_i|y_{-i}) < 2c_i(y) - \alpha\cdot d(s_i, x_i) &\le 2c_i(y) - \alpha\big(d(s_i, y_i) - \sum_{j\in N(i)} d(y_i, y_j)\big)\\ &= c_i(y) + \sum_{j\in N(i)} d(y_i, y_j).\end{aligned}$$

This brings us to the following bound:

$$\frac{c(x_i|y_{-i})}{c(y_i|y_{-i})} < \frac{c_i(y) + \sum_{j\in N(i)} d(y_i, y_j)}{c_i(y) + (1-\alpha)\cdot \sum_{j\in N(i)} d(y_i, y_j)} \qquad = 1 + \frac{\alpha \sum_{j\in N(i)} d(y_i, y_j)}{c_i(y) + (1-\alpha)\cdot \sum_{j\in N(i)} d(y_i, y_j)}.$$

Recall that by our assumption $c_i(y) > \sum_{j\in N(i)} d(y_i, y_j)$, this implies that

$$c_i(y) + (1-\alpha)\cdot \sum_{j\in N(i)} d(y_i, y_j) > (2-\alpha)\sum_{j\in N(i)} d(y_i, y_j)$$

and the claim follows. □

Finally, we present the proof of Lemma B.2 that was used in the proof of Claim B.1 above.

Lemma B.2. *Let $\alpha \le \frac{1}{2}$. Fix an optimal solution y which is not a Nash equilibrium and let player i be a player that can reduce its cost by playing x_i. Then: $d(s_i, x_i) > d(s_i, y_i) - \sum_{j \in N(i)} d(y_i, y_j)$.*

Proof. Note the following: by the triangle inequality for any player j it holds that: $d(x_i, y_j) \ge d(s_i, y_j) - d(s_i, x_i)$ and $d(s_i, y_j) \ge d(s_i, y_i) - d(y_i, y_j)$. By combining the two together we have that $d(x_i, y_j) \ge d(s_i, y_i) - d(y_i, y_j) - d(s_i, x_i)$. This gives us the following lower bound on $c_i(x_i, y_{-i})$:

$$\begin{aligned} c_i(x_i, y_{-i}) &= \alpha \cdot d(s_i, x_i) + (1-\alpha) \sum_{j \in N(i)} d(x_i, y_j) \\ &\ge \alpha \cdot d(s_i, x_i) + (1-\alpha) \sum_{j \in N(i)} \Big(d(s_i, y_i) - d(y_i, y_j) - d(s_i, x_i) \Big) \\ &= \alpha \cdot d(s_i, x_i) + (1-\alpha) \cdot |N(i)| \cdot \big(d(s_i, y_i) - d(s_i, x_i) \big) - (1-\alpha) \sum_{j \in N(i)} d(y_i, y_j). \end{aligned}$$

Since x_i minimizes player i's cost it has to be the case that: $c_i(x_i, y_{-i}) < c_i(y)$. Thus the following inequality holds:

$$\begin{aligned} &\alpha \cdot d(s_i, x_i) + (1-\alpha)|N(i)| \cdot \big(d(s_i, y_i) - d(s_i, x_i) \big) - (1-\alpha) \sum_{j \in N(i)} d(y_i, y_j) \\ &\quad < \alpha \cdot d(s_i, y_i) + (1-\alpha) \sum_{j \in N(i)} d(y_i, y_j). \end{aligned}$$

After some rearranging we get that:

$$d(s_i, x_i) > d(s_i, y_i) - \frac{2(1-\alpha)}{(1-\alpha) \cdot |N(i)| - \alpha} \sum_{j \in N(i)} d(y_i, y_j)$$

which implies that the claim holds whenever $\frac{2(1-\alpha)}{(1-\alpha) \cdot |N(i)| - \alpha} \le 1$. For $\alpha \le \frac{1}{2}$, this latter bound occurs for $|N(i)| \ge 3$.

The case of $|N(i)| \le 2$ is handled separately and requires we use the assumption that y is an optimal solution. Denote i's neighbors by j and k. Then:

$$\alpha \cdot d(s_i, y_j) + 2(1-\alpha) \cdot d(y_j, y_k) \ge \alpha \cdot d(s_i, y_i) + 2(1-\alpha) \big(d(y_i, y_j) + d(y_i, y_k) \big).$$

By the triangle inequality the previous inequality implies that $d(s_i, y_j) \ge d(s_i, y_i)$. When combining this with the fact that $d(x_i, y_j) \ge d(s_i, y_j) - d(s_i, x_i)$ we get that $d(x_i, y_j) \ge d(s_i, y_i) - d(s_i, x_i)$; similarly we get for k that $d(x_i, y_k) \ge d(s_i, y_i) - d(s_i, x_i)$. Therefore,

$$\begin{aligned} c_i(x_i, y_{-i}) &= \alpha \cdot d(s_i, x_i) + (1-\alpha) \big(d(x_i, y_j) + d(x_i, y_k) \big) \\ &\ge \alpha \cdot d(s_i, x_i) + 2(1-\alpha) \big(d(s_i, y_i) - d(s_i, x_i) \big), \end{aligned}$$

and since x_i is player i's best response it has to be the case that:

$$\alpha \cdot d(s_i, x_i) + 2(1-\alpha) \big(d(s_i, y_i) - d(s_i, x_i) \big) < \alpha \cdot d(s_i, y_i) + (1-\alpha) \big(d(y_i, y_j) + d(y_i, y_k) \big).$$

After some rearranging we get that:

$$(2 - 3\alpha) d(s_i, y_i) - (1-\alpha) \big(d(y_i, y_j) + d(y_i, y_k) \big) < (2 - 3\alpha) d(s_i, x_i).$$

By dividing both sides of the inequality by $(2 - 3\alpha)$ we get that $d(s_i, x_i) > d(s_i, y_i) - \sum_{j \in N(i)} d(y_i, y_j)$ holds whenever $\frac{1-\alpha}{2-3\alpha} \le 1$. This completes the proof since for $\alpha \le \frac{1}{2}$ it is always the case that $\frac{1-\alpha}{2-3\alpha} \le 1$. □

References

[1] David Bindel, Jon Kleinberg, Sigal Oren, How bad is forming your own opinion?, in: Proc. 52nd IEEE Symposium on Foundations of Computer Science, 2011.

[2] Larry Blume, The statistical mechanics of strategic interaction, Games Econ. Behav. 5 (1993) 387–424.

[3] Yuri Boykov, Olga Veksler, Ramin Zabih, Fast approximate energy minimization via graph cuts, in: Proc. 7th International Conference on Computer Vision, 1999, pp. 377–384.

[4] Glenn Ellison, Learning, local interaction, and coordination, Econometrica 61 (1993) 1047–1071.

[5] Diodato Ferraioli, Paul W. Goldberg, Carmine Ventre, Decentralized dynamics for finite opinion games, in: Proceedings of the Algorithmic Game Theory: 5th International Symposium, SAGT 2012, Barcelona, Spain, October 22–23, 2012, vol. 7615, Springer, 2012, p. 144.

[6] Noah E. Friedkin, Eugene C. Johnsen, Social influence and opinions, J. Math. Sociol. 15 (3–4) (1990) 193–205.

[7] Bezalel Gavish, Suresh Sridhar, Computing the 2-median on tree networks in $o(n \lg n)$ time, Networks 26 (1995) 305–317.
[8] A.J. Goldman, Optimal center location in simple networks, Transp. Sci. 5 (1971) 212–221.
[9] Georg Gottlob, Gianluigi Greco, Francesco Scarcello, Pure Nash equilibria: hard and easy games, J. Artif. Intell. Res. (2005) 357–406.
[10] Matthew O. Jackson, Social and Economic Networks, Princeton University Press, 2008.
[11] Oded Kariv, S. Louis Hakimi, An algorithmic approach to network location problems. II: the p-medians, SIAM J. Appl. Math. 37 (3) (1979) 539–560.
[12] Jon M. Kleinberg, Éva Tardos, Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov random fields, J. ACM 49 (5) (2002) 616–639.
[13] Dov Monderer, Lloyd S. Shapley, Potential games, Games Econ. Behav. 14 (1996) 124–143.
[14] Stephen Morris Contagion, Rev. Econ. Stud. 67 (2000) 57–78.
[15] Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay V. Vazirani, Algorithmic Game Theory, Cambridge University Press, New York, NY, USA, 2007.