

Resolving the Optimal Metric Distortion Conjecture

Vasilis Gkatzelis
Drexel University
gkatz@drexel.edu

Daniel Halpern
University of Toronto
daniel.halpern@mail.utoronto.ca

Nisarg Shah
University of Toronto
nisarg@cs.toronto.edu

Abstract

We study the following metric distortion problem: there are two finite sets of points, V and C , that lie in the same metric space, and our goal is to choose a point in C whose total distance from the points in V is as small as possible. However, rather than having access to the underlying distance metric, we only know, for each point in V , a ranking of its distances to the points in C . We propose algorithms that choose a point in C using only these rankings as input and we provide bounds on their *distortion* (worst-case approximation ratio). A prominent motivation for this problem comes from voting theory, where V represents a set of voters, C represents a set of candidates, and the rankings correspond to ordinal preferences of the voters.

A major conjecture in this framework is that the optimal deterministic algorithm has distortion 3. We resolve this conjecture by providing a polynomial-time algorithm that achieves distortion 3, matching a known lower bound. We do so by proving a novel lemma about matching voters to candidates, which we refer to as the *ranking-matching lemma*. This lemma induces a family of novel algorithms, which may be of independent interest, and we show that a special algorithm in this family achieves distortion 3. We also provide more refined, parameterized, bounds using the notion of α -decisiveness, which quantifies the extent to which a voter may prefer her top choice relative to all others. Finally, we introduce a new randomized algorithm with improved distortion compared to known results, and also provide improved lower bounds on the distortion of all deterministic and randomized algorithms.

1 Introduction

In the metric distortion problem, we are given two finite (potentially non-disjoint) sets V and C that lie in the same metric space. If we had access to the distance metric d over $V \cup C$, we could compute a point $c^* \in C$ that minimizes its total distance to the points in V , i.e., $c^* \in \arg \min_{c \in C} \sum_{i \in V} d(i, c)$. However, rather than having access to the exact distances, we are given, for each point $i \in V$, a ranking σ_i over the points in C by their distance to i (with ties broken arbitrarily). Given only this collection of rankings $\sigma = (\sigma_i)_{i \in V}$, we need to design an algorithm (a social choice rule) that picks a point $\hat{c} \in C$ (possibly in a randomized fashion), and we are interested in its worst-case approximation ratio $\mathbb{E}[\sum_{i \in V} d(i, \hat{c})] / \sum_{i \in V} d(i, c^*)$, where the expectation is over randomization of the algorithm. This worst-case approximation ratio is known as the *distortion* of the algorithm in this framework. Note that the need for approximation stems not from computational constraints, but from the lack of full information. Our main question is: *To what extent can we minimize the total distance using only ordinal information regarding pairwise distances?*

An interesting motivating example for this question stems from voting theory, where V is a set of voters, C is a set of candidates, and the distance between voters and candidates corresponds to the distance of their ideologies with respect to different issues. This model, inspired by spatial models of voting from the political science literature, such as the Downsian proximity model [26], has recently become popular in the computer science literature [12, 47, 44]. An alternative, possibly more concrete, interpretation of these distances can also be derived through the well-studied problem of facility location [21, 36, 25]. In this case V corresponds to a set of customers and C is the set of feasible locations for a new facility to be opened; the distance measure between customers and locations could correspond to the actual physical distance between them, or the time or cost required for a customer to reach each location. In both cases, the most common way to elicit the preferences of the voters, or customers, is in the form of a ranking of the available options (e.g., see [15, Part I] or [42]), and the goal is to pick an outcome that minimizes the total distance.

Centuries of research on voting theory has produced a plethora of voting rules that choose an outcome using the ranked preferences of voters as input. In recent work, starting with Anshelevich et al. [11], researchers have analyzed the performance of these voting rules using the metric distortion framework described above, which enables a quantitative comparison among them. Anshelevich et al. [11] showed that no deterministic algorithm can achieve a distortion better than 3, while a popular voting rule known as Copeland’s rule comes close, by achieving distortion 5. Informally, Copeland’s rule measures the strength of each candidate c as the number of candidates c' defeated by c in a pairwise election (i.e., the candidates c' that are ranked below c by the majority of voters), and picks a candidate with the greatest strength. They also analyzed the distortion of other popular voting rules such as plurality, Borda count, k -approval, and veto, and showed that their distortion increases linearly with the number of voters or candidates. Subsequently, Skowron and Elkind [47] analyzed another popular voting rule, known as single transferable vote (STV), and showed that its distortion grows logarithmically in the number of candidates.

The first deterministic algorithm to break the bound of 5 was very recently proposed by Munagala and Wang [44], who designed a novel deterministic algorithm that achieves a distortion of $2 + \sqrt{5} \approx 4.236$. Their algorithm does not even require access to voters’ full ordinal preferences. It takes as input only the weighted tournament graph, which contains, for every pair of candidates, the number of voters that prefer one over another. They proved that $2 + \sqrt{5}$ is the best distortion of any deterministic algorithm with this limited input, and they conjectured that with full ordinal preferences, the best possible distortion should be 3, outlining a sufficient condition for achieving distortion 3. However, they were unable to prove that there always exists a candidate satisfying their sufficient condition. Independently, Kempe [39] derived the same bound of $2 + \sqrt{5}$ using a linear

programming duality framework, and provided several alternative formulations of the sufficient condition outlined by Munagala and Wang [44]. He used the linear programming duality framework to show that the distortion of two social choice rules, the ranked pairs and Schulze’s rule, grows as the square root of the number of candidates, thus disproving a conjecture by Anshelevich et al. [11] that ranked pairs might have distortion 3 (this conjecture was also earlier disproved by Goel et al. [31]). The main result of our paper is a culmination of this line of work in the form of a novel deterministic algorithm that guarantees the optimal distortion of 3.

1.1 Our Results & Techniques

The technical core of our paper is in Section 3, where we prove the *Ranking-Matching Lemma*, an elementary lemma regarding the existence of fractional perfect matchings within a family of vertex-weighted bipartite graphs induced by an instance of our problem. Specifically, given a problem instance and some candidate a , we get a bipartite graph whose vertices correspond to the set of voters on one side, and the set of candidates on the other. The vertex of each voter i has some non-negative weight $p(i)$ and the vertex of each candidate c has some non-negative weight $q(c)$ such that the total weight of all voters $\sum_{i \in V} p(i)$ and the total weight of all candidates $\sum_{c \in C} q(c)$ are both normalized to 1. An edge (i, c) between the vertices of some voter i and some candidate c (potentially equal to a) exists if and only if candidate a is (weakly) closer voter i than candidate c , indicating that voter i would (weakly) prefer candidate a to c . Intuitively, the more edges this graph has, the more appealing candidate a appears, since more voters would prefer a over other candidates. For example, if the graph is complete, then a must be the top choice of all voters. Following this intuition, we propose a notion of fractional perfect matching in this graph so that, if such a matching exists for the graph of candidate a , then a is an “appealing” candidate. Specifically, this matching suggests that for every other candidate c , there are sufficiently many *distinct* voters who prefer a to c . The Ranking-Matching Lemma shows that for *any* problem instance and *any* weight vectors p and q , there always exists some candidate a whose graph admits a fractional perfect matching, i.e., there always exists at least one appealing candidate.

Armed with this non-trivial result, we then leverage it in order to achieve a good distortion bound. A natural social choice rule would be to select one of the candidates that satisfy the aforementioned matching criterion. Note that the Ranking-Matching Lemma leaves open the choice of weights of voters and candidates (given by p and q). These weights can be used to quantify the relative significance of voters and candidates, and each choice yields a different algorithm. As we discuss in Section 7, some of these algorithms could be of independent interest, but in this paper we focus on the case where the weights of all voters are the same, and those of the candidates are proportional to their plurality score (i.e., the number of voters that rank them first). We observe that for this set of weights, the Ranking-Matching Lemma also implies the existence of an *integral* perfect matching in a different bipartite graph whose vertices on both sides correspond to voters. For a given candidate a , we refer to this graph as the integral domination graph of a and denote it by $G(a)$, and an edge (i, j) between the vertices of voters i and j exists if voter i weakly prefers a to the top choice candidate of voter j .

We define a novel voting rule, PLURALITYMATCHING, which returns a candidate a whose integral domination graph $G(a)$ has a perfect matching, and the Ranking-Matching Lemma implies that such a candidate always exists. We then prove our main result in Section 4, which is that PLURALITYMATCHING achieves the optimal distortion bound of 3.

Our approach is inspired by the recent work of Munagala and Wang [44], who used integral perfect matchings in bipartite graphs to define the *matching uncovered set*. Specifically, given a *pair* of candidates a and b , they define a bipartite graph $G(a, b)$, in which both sides of vertices

correspond to the voters, and edge (i, j) between voters i and j exists if there is some candidate c for which voter i weakly prefers a to c and voter j weakly prefers c to b . The matching uncovered set is the set of candidates a such that for every candidate b the $G(a, b)$ graph has a perfect matching. They proved that every candidate in this set has distortion at most 3, but were unable to answer whether this set is always non-empty. Our integral domination graph $G(a)$ of candidate a is a subgraph of $G(a, b)$ for every candidate b . Thus, a candidate a for which $G(a)$ has a perfect matching is also in the matching uncovered set, and the existence of such a candidate is guaranteed by our Ranking-Matching Lemma, implying that the matching uncovered set is always non-empty.

While the non-emptiness of the matching uncovered set, combined with the argument of Munagala and Wang [44], implies a distortion bound of 3, we provide a direct proof of this bound that is much shorter and simpler because it uses the existence of a perfect matching in the much sparser integral domination graph (thus using the full strength of the Ranking-Matching Lemma). More significantly, as we show in Section 4.1 this proof also allows us to achieve improved distortion bounds for interesting special cases which the matching uncovered set does not meet.

Having settled the question of distortion for the class of deterministic algorithms, an exciting open problem is to also achieve optimal distortion bounds for *randomized* algorithms. Prior to this work, the best known distortion bound achieved by a randomized algorithm was $3 - 2/m$ [38], which is better than what any deterministic algorithm can achieve. In Section 5, we take a step toward a better understanding of the power and limitations of randomization in this context. Specifically, we propose a new randomized algorithm, which we term SMARTDICTATORSHIP, that only depends on the plurality scores of the candidates. We show that it matches the distortion bound of $3 - 2/m$ for general metric spaces, but has improved distortion in special cases. Further, we show that among randomized algorithms that depend only on the plurality scores, SMARTDICTATORSHIP has optimal distortion, even in the aforementioned special cases. We also provide an improved lower bound that applies to randomized algorithms that use full preference rankings.

As the number of candidates grows, no known distortion bound is better than 3, and we know that no randomized algorithm can achieve better than 2; the problem of identifying the optimal bound in $[2, 3]$ for $m \rightarrow \infty$, and the optimal bound for a fixed value of m in general, remains open.

Throughout the paper, our distortion bounds are presented as a function of the number of candidates in the instance, but we also parameterize our bounds using the notion of α -decisiveness due to Anshelevich and Postl [6]. Given a parameter value $\alpha \in [0, 1]$, an instance is α -decisive if every voter’s distance from her top candidate is at most α times her distance from any other candidate. For a motivating example, consider the peer selection problem, where the voters are choosing a candidate among themselves as a representative; in this case, every voter’s distance from herself is zero, so this is a special case of $\alpha = 0$. Among other results, we are the first to establish non-trivial (i.e. greater than 1) lower bounds on distortion for this special case for deterministic as well as randomized algorithms.

The paper is organized as follows. Section 2 presents the model, Section 3 presents the Ranking-Matching Lemma, and Sections 4 and 5 study the distortion of deterministic and randomized algorithms, respectively. Section 6 provides an overview of related work not covered above, Section 7 contains a discussion of our results and open directions, and the appendix contains deferred proofs and additional observations.

2 Model

Let $V = \{1, \dots, n\}$ be a set of n voters, and C be a set of m candidates. Throughout the paper, we use letters i, j, k to index voters, and letters a, b, c to index candidates. Also, we use $\Delta(V)$ and

$\Delta(C)$ to denote the set of probability distributions over V and C , respectively; for simplicity we also use this notation for vectors of nonnegative weights that add up to 1, even when these weights do not correspond to probabilities.

The voters and candidates are located in a common metric space (\mathcal{M}, d) , where $d : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ is the metric (distance function) satisfying the following conditions: (1) Positive definiteness: $\forall x, y \in \mathcal{M}, d(x, y) \geq 0$, and $d(x, y) = 0$ if and only if $x = y$; (2) Symmetry: $\forall x, y \in \mathcal{M}, d(x, y) = d(y, x)$; and (3) Triangle inequality: $\forall x, y, z \in \mathcal{M}, d(x, y) + d(y, z) \geq d(x, z)$

For simplicity of notation, we extend d to operate directly on the voters and the candidates instead of the points of the metric space where they are located. Thus, we will write $d(x, y)$, for $x, y \in V \cup C$, to mean the distance between the points where voters/candidates x and y are located. Given a metric, d , the *social cost* of a candidate c is defined as $\text{SC}(c, d) = \sum_{i \in V} d(i, c)$; we simply write $\text{SC}(c)$ when the metric d is clear from the context.

Preference Profiles. Given sets V, C and a metric d , for each voter $i \in V$ we get a linear order, or *preference ranking*, σ_i over the candidates in C , where i ranks candidates in non-decreasing order of their distance from her. We say that candidate c *defeats* candidate c' in vote i , denoted $c \succ_i c'$, if voter i ranks c above c' ; we say that c *weakly defeats* c' , denoted $c \succsim_i c'$, if $c = c'$ or $c \succ_i c'$. We say that d is consistent with the preference ranking σ_i of voter i , denoted $d \triangleright \sigma_i$, if $d(i, c) \leq d(i, c')$ for all c, c' such that $c \succ_i c'$. Note that candidates that are equidistant to a voter can be ranked arbitrarily by the voter. We refer to a profile of preference rankings, $\sigma = (\sigma_i)_{i \in V}$, as the *preference profile*, and we say that d is consistent with the preference profile σ , denoted $d \triangleright \sigma$, if $d \triangleright \sigma_i$ for each $i \in V$.

For a given σ , let $\text{top}(i) \in C$ denote the candidate ranked highest in σ_i , and let the plurality score of candidate a , i.e., the number of voters who rank a as their top choice, be $\text{plu}(a) = |\{i \in V : \text{top}(i) = a\}|$. With slight abuse of notation, we also write $\text{plu}(D) = \sum_{a \in D} \text{plu}(a)$ for all $D \subseteq C$.

Social Choice Rules and Distortion. Given an instance (V, C, d) , our main goal in this paper is to output a candidate $c \in C$ that minimizes the social cost $\text{SC}(c, d)$. However, instead of having full access to the metric d , we have access to a tuple $\mathcal{E} = (V, C, \sigma)$, which we refer to as an *election*, and all we know is that d is consistent with σ .

Let $\mathcal{L}(C)$ denote the set of linear orders (rankings) over C . A *social choice rule* (or simply a rule) is a function $f : \cup_{n \geq 1} \mathcal{L}(C)^n \rightarrow \Delta(C)$. We say that the rule is *deterministic* if it always returns a distribution with singleton support. Note that we seek rules that operate on preference profiles consisting of any number of votes, but over a fixed set of candidates. The *distortion* of a social choice rule f , denoted $\text{distortion}(f)$, is the worst-case approximation ratio it provides to social cost:

$$\text{distortion}(f) = \sup_{\sigma} \sup_{d : d \triangleright \sigma} \frac{\mathbb{E}[\text{SC}(f(\sigma), d)]}{\min_{c \in C} \text{SC}(c, d)},$$

where the expectation is over possible randomization in the social choice rule. Note that there is no restriction on the number of voters in the preference profile, so we are interested in the distortion as a function of the number of candidates. This is because in real-world applications of voting, it is reasonable to expect the number of candidates, m , to be small — finer distortion bounds as a function of m are therefore important — but the number of votes may be large.

Decisive Voters. We further refine the distortion bounds using the α -decisiveness framework of Anshelevich and Postl [6]. For $\alpha \in [0, 1]$, voter i is called α -decisive if $d(i, \text{top}(i)) \leq \alpha \cdot d(i, c)$, $\forall c \neq \text{top}(i)$. In words, the distance of voter i from her top choice is at most α times her distance from any other candidate. We say that the metric space is α -decisive if every voter is α -decisive.

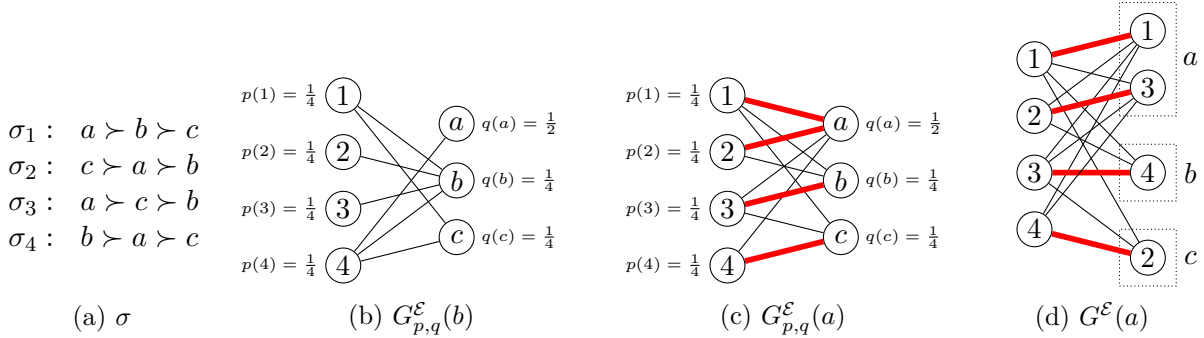


Figure 1: Domination graph examples

For a given value of α , when we analyze the distortion of social choice rules for α -decisive metric spaces, we take the worst case only over metric spaces that are α -decisive and for which $d \triangleright \sigma$.

Note that for $\alpha = 0$ the distance of every voter from her top-choice candidate is zero. This captures settings where each voter is a candidate herself, and ranks herself at the top. For instance, the *peer selection* setting, in which n agents submit ranked votes over themselves and each agent ranks herself first, is a special case of $\alpha = 0$; thus, the upper bounds (and some lower bounds) presented in our paper for $\alpha = 0$ apply to this interesting special case as well. When $\alpha = 1$, we retrieve the general, unrestricted, metric space.

3 Ranking-Matching Lemma

We begin by proving our key contribution, which is an elementary lemma about matching preference rankings, i.e., voters, to candidates. We then use this lemma to define a novel deterministic social choice rule, and in the next section show that this rule has distortion at most 3, thus resolving the open conjecture. First, we need to introduce some terminology.

Definition 1. Given an election $\mathcal{E} = (V, C, \sigma)$, and (normalized) weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$, the (p, q) -domination graph of candidate a is the vertex-weighted bipartite graph $G^{\mathcal{E}}_{p,q}(a) = (V, C, E_a, p, q)$, where $(i, c) \in E_a$ if and only if $a \succ_i c$. Vertex $i \in V$ has weight $p(i)$ and vertex $c \in C$ has weight $q(c)$. We will drop \mathcal{E} from notation when it is clear from the context.

For example, Figure 1a shows the preference rankings of four voters (1, 2, 3, 4) over three candidates (a, b, c). Figure 1b and Figure 1c show the (p, q) -domination graphs of candidates b and a , respectively, with $p(1) = p(2) = p(3) = p(4) = \frac{1}{4}$ and $q(a) = \frac{1}{2}$, $q(b) = q(c) = \frac{1}{4}$.

Definition 2. We say that the (p, q) -domination graph of candidate a , i.e., $G^{\mathcal{E}}_{p,q}(a)$, admits a fractional perfect matching if there exists a weight function $w : E_a \rightarrow \mathbb{R}_{\geq 0}$ such that the total weight of edges incident on each vertex equals the weight of the vertex, i.e., for each $i \in V$ we have $\sum_{c \in C: (i,c) \in E_a} w((i, c)) = p(i)$, and for each $c \in C$ we have $\sum_{i \in V: (i,c) \in E_a} w((i, c)) = q(c)$.

With slight abuse of notation, let us denote $p(S) = \sum_{i \in S} p(i)$ for all $S \subseteq V$ and $q(D) = \sum_{a \in D} q(a)$ for all $D \subseteq C$. Given an election $\mathcal{E} = (V, C, \sigma)$, a candidate $a \in C$, and a subset of voters $S \subseteq V$, define $D_a^{\mathcal{E}}(S) = \{c \in C : a \succ_i c, \exists i \in S\}$ to be the set of candidates that a weakly defeats in at least one vote in S . Note that this is precisely the set of neighbors of S in $G^{\mathcal{E}}_{p,q}(a)$ (for any p and q). A simple generalization of Hall's condition shows that the existence of a fractional perfect matching in $G^{\mathcal{E}}_{p,q}(a)$ is equivalent to the set of neighbors of S having at least as much weight as S itself, for all $S \subseteq V$. For the sake of completeness, we provide a short proof in Appendix A.

Lemma 1. *For any election $\mathcal{E} = (V, C, \sigma)$, weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$, and candidate $a \in C$, $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching if and only if $q(D_a^{\mathcal{E}}(S)) \geq p(S)$ for all $S \subseteq V$, and whether $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching can be checked in strongly polynomial time.*

Using this lemma for the instance in Figure 1, we can verify that $G_{p,q}^{\mathcal{E}}(b)$ does not admit a fractional perfect matching by letting $S = \{2, 3\}$, for which $p(S) = \frac{1}{2}$, and observing that $D_b^{\mathcal{E}}(S) = \{b\}$ with $q(b) = \frac{1}{4}$. On the other hand, $G_{p,q}^{\mathcal{E}}(a)$ does admit a fractional perfect matching, indicated using thick red edges in Figure 1c.

Note that if, in a given election, there exists some candidate a such that $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching, then by Lemma 1, in *any* subset of votes S , a weakly defeats a set of candidates with total weight at least that of S . This intuitively makes candidate a a compelling choice. In fact, in Section 4 we show that for appropriate choices of p and q , this also implies low distortion. However, it is far from obvious that such a candidate a is guaranteed to exist. Our key lemma, presented next, shows that there will always exist such a candidate in *any* election \mathcal{E} and for *any* weight vectors p and q .

While we only need to prove the lemma for the choices of p and q that we use in Section 4 to settle the conjecture, we prove it for all possible choices of p and q for two reasons. First, our proof is inductive, and the proof for given choices of p and q in a given election requires the lemma to hold for *all* choices of p and q in smaller elections. Second, as we discuss in Sections 3.1 and 7, other choices of p and q also lead to interesting social choice rules, which may be of independent interest.

Lemma 2 (RANKING-MATCHING LEMMA). *For any election \mathcal{E} , and weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$, there exists a candidate a whose (p, q) -domination graph $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching, and such a candidate can be computed in strongly polynomial time.*

Proof. Suppose that the lemma is not true. Let $\mathcal{E} = (V, C, \sigma)$ be an election with the smallest number of voters for which the lemma is violated for any weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$. Fix such p and q . By Lemma 1, we have that for each candidate $a \in C$, there exists $S \subseteq V$ for which $q(D_a^{\mathcal{E}}(S)) < p(S)$. We refer to such a set S as a *counterexample for a* . We say that a counterexample for a is *minimal* if no strict subset of it is a counterexample for a . For each $a \in C$, fix an arbitrary minimal counterexample $X_a \subseteq V$.

Consider the set of candidates $C^* = \arg \max_{c \in C} p(X_c)$ whose minimal counterexamples have the largest weight under p . Define a partial order R over C^* , whereby, for $b, c \in C^*$, bRc if and only if $X_b \supseteq X_c$ and $b \succ_i c$ for all $i \in X_c$. Let $a \in C^*$ be an arbitrary maximal element of this partial order. Note that the counterexample for a has the highest weight, and no other candidate b with this property “dominates” a in the sense of being better than a under relation R . Fix this candidate a . For brevity of notation, write $X = X_a$ and $D = D_a^{\mathcal{E}}(X_a)$. We use the standard notation $\bar{X} = V \setminus X$ and $\bar{D} = C \setminus D$. The next lemma shows that for any candidate b that a does not weakly defeat in its minimal counterexample, the minimal counterexamples of a and b are incomparable in the sense that neither is a subset of the other.

Lemma 3. *For each $b \in \bar{D}$, we have $X \setminus X_b \neq \emptyset$ and $X_b \setminus X \neq \emptyset$.*

Proof. Fix $b \in \bar{D}$. By the definition of D , this implies $b \succ_i a$ for all $i \in X$.

First, we show that $X \setminus X_b \neq \emptyset$. Suppose this is false, so $X \subseteq X_b$. Then, we have $p(X_b) \geq p(X)$. However, since $a \in C^* = \arg \max_{c \in C} p(X_c)$, we also have $b \in C^*$. Further, since $X \subseteq X_b$ and $b \succ_i a$ for all $i \in X$, we have bRa . This contradicts the fact that a is a maximal element of C^* under partial order R .

Next, we show that $X_b \setminus X \neq \emptyset$. Suppose this is not true, so $X_b \subseteq X$. Since we have already established that $X \setminus X_b \neq \emptyset$, we know $X_b \neq X$. Hence, $X_b \subsetneq X$. Since X is a minimal counterexample

for a , X_b is not a counterexample for a . Thus, $D_a^\mathcal{E}(X_b) \geq p(X_b)$. However, since $b \succ_i a$ for all $i \in X$, we have that $q(D_b^\mathcal{E}(X_b)) \geq q(D_a^\mathcal{E}(X_b)) \geq p(X_b)$, which contradicts the fact that X_b is a counterexample for b . \square

We now use Lemma 3 to prove Lemma 2. We now make four claims: (1) $D \neq \emptyset$, (2) $q(\overline{D}) > 0$ (which implies $\overline{D} \neq \emptyset$), (3) $p(X) > 0$ (which implies $X \neq \emptyset$), and (4) $\overline{X} \neq \emptyset$. The first claim is true because $a \in D$. For the next two claims, recall that X is a counterexample for a . Hence, $0 \leq q(D) < p(X) \leq 1$, which implies both $q(\overline{D}) > 0$ and $p(X) > 0$. Finally, for the last claim, suppose for contradiction that $\overline{X} = \emptyset$, i.e., $X = V$. Then, $X_b \setminus X = \emptyset$ for all b . Thus, by Lemma 3, $\overline{D} = \emptyset$, which contradicts claim (2).

Next, we show that every candidate $c \in \overline{D}$ weakly defeats candidates of sufficient total weight in its minimal counterexample X_c , i.e., that $q(D_c^\mathcal{E}(X_c))$ is sufficiently large. Later, we will use this to identify a specific candidate $b \in \overline{D}$ and show that $q(D_b^\mathcal{E}(X_b)) \geq p(X_b)$, contradicting the fact that X_b is a counterexample for b . For all $c \in \overline{D}$, we have that

$$\begin{aligned} q(D_c^\mathcal{E}(X_c)) &= q(D_c^\mathcal{E}(X_c) \cap D) + q(D_c^\mathcal{E}(X_c) \cap \overline{D}) \\ &\geq q(D_c^\mathcal{E}(X_c \cap X) \cap D) + q(D_c^\mathcal{E}(X_c \cap \overline{X}) \cap \overline{D}) \\ &\geq q(D_a^\mathcal{E}(X_c \cap X)) + q(D_c^\mathcal{E}(X_c \cap \overline{X}) \cap \overline{D}) \\ &\geq p(X_c \cap X) + q(D_c^\mathcal{E}(X_c \cap \overline{X}) \cap \overline{D}). \end{aligned} \tag{1}$$

Here, the first transition holds because for any set T , sets $T \cap D$ and $T \cap \overline{D}$ form a partition of T . The second transition holds because $D_c^\mathcal{E}(X_c \cap T) \subseteq D_c^\mathcal{E}(X_c)$ for any set T . The third transition holds because $D_a^\mathcal{E}(X_c \cap X) \subseteq D_c^\mathcal{E}(X_c \cap X) \cap D$: every candidate in $D_a^\mathcal{E}(X_c \cap X)$ (i.e. weakly defeated by a in $X_c \cap X$) is both in D (because it is also weakly defeated by a in X) and in $D_c^\mathcal{E}(X_c \cap X)$ (because it is also weakly defeated by c in $X_c \cap X$ as $c \succ_i a$ for all $i \in X$). Finally, for the last transition, note that $X_c \cap X \subsetneq X$ due to Lemma 3. Since X is a minimal counterexample for a , $X_c \cap X$ is not a counterexample for a . Thus, $q(D_a^\mathcal{E}(X_c \cap X)) \geq p(X_c \cap X)$.

Our goal now is to find some candidate $b \in \overline{D}$ and use Equation (1) to derive a contradiction to the fact that X_b is a counterexample for b (that is, show that $q(D_b^\mathcal{E}(X_b)) \geq p(X_b)$). We consider two cases, $p(\overline{X}) = 0$ and $p(\overline{X}) > 0$, and show that for each case we can identify such a candidate $b \in \overline{D}$.

First, suppose $p(\overline{X}) = 0$. Fix an arbitrary $b \in \overline{D}$; such a candidate exists because we know $\overline{D} \neq \emptyset$. Since $p(X) = 1$, we have that $p(X_b \cap X) = p(X_b)$. Substituting this in Equation (1) already yields $q(D_b^\mathcal{E}(X_b)) \geq p(X_b)$, which contradicts the fact that X_b is a counterexample for b .

The more interesting case is when $p(\overline{X}) > 0$, which implies $\overline{X} \neq \emptyset$. In this case, consider the restricted election $\overline{\mathcal{E}} = (\overline{X}, \overline{D}, \sigma|_{\overline{X}, \overline{D}})$, where $\sigma|_{\overline{X}, \overline{D}}$ denotes the preference profile σ restricted to the preferences of the voters in \overline{X} over the candidates in \overline{D} . Note that $\overline{\mathcal{E}}$ is a valid election because $\overline{X} \neq \emptyset$ and $\overline{D} \neq \emptyset$, as shown above. Further, $\overline{\mathcal{E}}$ has fewer voters than \mathcal{E} because $X \neq \emptyset$. Therefore, $\overline{\mathcal{E}}$ must satisfy Lemma 2 for any weight vectors, since \mathcal{E} was chosen as an election with the smallest number of voters for which Lemma 2 fails to hold. In particular, choose the re-normalized weight vectors $p' \in \Delta(\overline{X})$ and $q' \in \Delta(\overline{D})$ as follows:

$$p'(i) = p(i)/p(\overline{X}), \forall i \in \overline{X}, \quad \text{and} \quad q'(c) = q(c)/q(\overline{D}), \forall c \in \overline{D}.$$

Note that these are well-defined because we have already shown that $q(\overline{D}) > 0$, and we are in the case $p(\overline{X}) > 0$. Since Lemma 2 holds for $\overline{\mathcal{E}}$ with weight vectors p' and q' , there exists some candidate

$b \in \bar{D}$ such that $p'(S) \leq q'(D_b^{\bar{\mathcal{E}}}(S))$ for all $S \subseteq \bar{X}$. In particular, for $S = X_b \cap \bar{X}$, we have,

$$p'(X_b \cap \bar{X}) \leq q'(D_b^{\bar{\mathcal{E}}}(X_b \cap \bar{X})) = q'(D_b^{\mathcal{E}}(X_b \cap \bar{X}) \cap \bar{D}),$$

where the last transition holds because $\bar{\mathcal{E}}$ is a restriction of \mathcal{E} to the voters in \bar{X} and the candidates in \bar{D} , which implies $D_b^{\bar{\mathcal{E}}}(X_b \cap \bar{X}) = D_b^{\mathcal{E}}(X_b \cap \bar{X}) \cap \bar{D}$. Expanding the definitions of q' and p' , we get

$$\frac{q(D_b^{\mathcal{E}}(X_b \cap \bar{X}) \cap \bar{D})}{q(\bar{D})} \geq \frac{p(X_b \cap \bar{X})}{p(\bar{X})} \Rightarrow q(D_b^{\mathcal{E}}(X_b \cap \bar{X}) \cap \bar{D}) \geq p(X_b \cap \bar{X}) \cdot \frac{q(\bar{D})}{p(\bar{X})} \geq p(X_b \cap \bar{X}),$$

where the last step follows because X is a counterexample for a , which implies $q(\bar{D}) < p(X)$, i.e., $q(\bar{D}) > p(\bar{X})$. Substituting this bound in Equation (1), we get

$$q(D_b^{\mathcal{E}}(X_b)) \geq p(X_b \cap X) + p(X_b \cap \bar{X}) = p(X_b),$$

which is the desired contradiction to the fact that X_b is a counterexample for b .

Finally, to compute a desired candidate a , we can simply iterate over all m candidates, and check whether the (p, q) -domination graph of each admits a fractional perfect matching, which can be done in strongly polynomial time due to Lemma 1. \square

3.1 Weight Vectors and Integral Matching

Note that Lemma 2 leaves open the choice of weight vectors p and q , and thus induces a novel family of (deterministic) social choice rules: for any given p and q , the corresponding social choice rule returns, on a given election \mathcal{E} , an arbitrary candidate a whose (p, q) -domination graph $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching. We now discuss two interesting choices of p and q . Perhaps the most straightforward choice is that of uniform weight vectors.

Definition 3 (Rule UNIFORMMATCHING). Given an election $\mathcal{E} = (V, C, \sigma)$, UNIFORMMATCHING returns a candidate a (ties broken arbitrarily) whose $(p^{\text{uni}}, q^{\text{uni}})$ -domination graph admits a fractional perfect matching, where $p^{\text{uni}}(i) = 1/n$ for each $i \in V$ and $q^{\text{uni}}(c) = 1/m$ for each $c \in C$.

For any candidate a returned by this rule, Lemma 1 implies that $|D_a^{\mathcal{E}}(S)|/|C| \geq |S|/|V|$ for all $S \subseteq V$. Succinctly, this can be summarized as follows: *In any $x\%$ of the votes, for any x , candidate a weakly defeats at least $x\%$ of the candidates.* To the best of our knowledge, ours is the first result establishing that such a candidate always exists. We believe this could be of independent interest; for a discussion on how this may be connected to another major open conjecture in voting, see Section 7. The rule we are interested in, however, is induced by a different choice of weight vectors.

Definition 4 (Rule PLURALITYMATCHING). Given an election $\mathcal{E} = (V, C, \sigma)$, PLURALITYMATCHING returns a candidate a (ties broken arbitrarily) whose $(p^{\text{uni}}, q^{\text{plu}})$ -domination graph admits a fractional perfect matching, where $p^{\text{uni}}(i) = 1/n$ for each $i \in V$ and $q^{\text{plu}}(c) = \text{plu}(c)/n$ for each $c \in C$.

For any candidate a returned by this rule, Lemma 1 implies that $\text{plu}(D_a^{\mathcal{E}}(S)) \geq |S|$ for all $S \subseteq V$, i.e., in any subset of votes S , candidate a weakly defeats a set of candidates with total plurality score at least $|S|$. Note that, in the domination graphs in Figure 1, the weights used are, in fact, p^{uni} and q^{plu} , which suggests that PLURALITYMATCHING could return candidate a in that case.

With these weight vectors, we can equivalently view the domination graph a bit differently. Instead of each voter i on the left and each candidate c on the right having weights $1/n$ and $\text{plu}(c)/n$, respectively, we can let them have integral weights 1 and $\text{plu}(c)$, respectively. Then, we can replace each node c on the right with weight $\text{plu}(c)$ by $\text{plu}(c)$ many nodes, each representing a unique voter

whose top choice is c , that have weight 1 each and are connected to the same set of nodes on the left as c was. This leads to a bipartite graph whose vertices on both sides correspond to the voters.

Definition 5. Given an election $\mathcal{E} = (V, C, \sigma)$ and a candidate $a \in C$, define, with slight abuse of notation, the *integral domination graph* of candidate a as the bipartite graph $G^\mathcal{E}(a) = (V, V, E_a)$, where $(i, j) \in E_a$ if and only if $a \succ_i \text{top}(j)$.

Returning to the example in Figure 1, the graph $G^\mathcal{E}(a)$ appears in Figure 1d. The only edges missing are $(2, 2)$ and $(4, 4)$, since $\text{top}(2) = c$ and $c \succ_2 a$, and $\text{top}(4) = b$ and $b \succ_4 a$. Comparing the graph $G^\mathcal{E}(a)$ in Figure 1d to the graph $G_{p^{\text{uni}}, q^{\text{plu}}}^\mathcal{E}(a)$ in Figure 1c, note that candidate a was replaced by voters 1 and 3 (who rank a at the top), candidate b was replaced by voter 4, and candidate c was replaced by voter 2 (also indicated in the figure, using dotted boxes).

It is easy to check that Hall's condition for the existence of an (*integral*) *perfect matching* in the integral domination graph is identical to the condition outlined above for the candidate returned by PLURALITYMATCHING. Therefore, a direct consequence of Lemma 2 is that there always exists some candidate a such that $G^\mathcal{E}(a)$ admits a perfect matching, and PLURALITYMATCHING returns one such candidate. A perfect matching in Figure 1d is denoted with thick red lines. Comparing this to Figure 1c, it is easy to see how this integral matching is actually derived directly from the fractional matching in $G_{p^{\text{uni}}, q^{\text{plu}}}^\mathcal{E}(a)$. Note that we could have also chosen edges $(1, 3)$ and $(2, 1)$ instead of $(1, 1)$ and $(2, 3)$ in Figure 1d.

Corollary 1. *Let $\mathcal{E} = (V, C, \sigma)$ be any election. For each candidate $a \in C$, $G^\mathcal{E}(a)$ admits a perfect matching if and only if $G_{p^{\text{uni}}, q^{\text{plu}}}^\mathcal{E}(a)$ admits a fractional perfect matching. Consequently, there exists a candidate $a \in C$ whose integral domination graph $G^\mathcal{E}(a)$ admits a perfect matching, i.e., for which there exists a bijection $M : V \rightarrow V$ satisfying $a \succ_i \text{top}(M(i))$ for each $i \in V$.*

In the next section, we use this guarantee to achieve an optimal distortion bound.

4 The Distortion of Deterministic Social Choice Rules

We begin by proving that PLURALITYMATCHING has distortion at most $2 + \alpha$ for α -decisive metric spaces. Since all metric spaces are 1-decisive, this settles a major conjecture that the optimal deterministic rule has distortion 3 [11, 44], matching a lower bound of 3 established by Anshelevich et al. [11]. The proof crucially leverages the fact that PLURALITYMATCHING returns a candidate whose integral domination graph has a perfect matching (Corollary 1). The proof is rather short, as the heavy lifting is already done in proving the Ranking-Matching Lemma (Lemma 2).

Theorem 1. *For every $m \geq 3$ and $\alpha \in [0, 1]$, PLURALITYMATCHING has distortion $2 + \alpha$ for α -decisive metric spaces.*

Proof. We begin with the upper bound. Let (\mathcal{M}, d) be an α -decisive metric space, and let σ be the induced preference profile. Let a be the candidate returned by PLURALITYMATCHING, and let b be any other candidate. We want to show that $\text{SC}(a) \leq (2 + \alpha) \cdot \text{SC}(b)$.

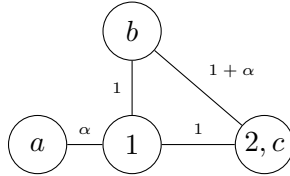
Due to Corollary 1, the integral domination graph $G^\mathcal{E}(a)$ admits a perfect matching $M : V \rightarrow V$ which satisfies $a \succ_i \text{top}(M(i))$ for each $i \in V$. Thus,

$$\begin{aligned} \text{SC}(a) &= \sum_{i \in V} d(i, a) \\ &\leq \sum_{i \in V} d(i, \text{top}(M(i))) \quad (\because a \succ_i \text{top}(M(i)), \forall i \in V) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i \in V} (d(i, b) + d(b, \text{top}(M(i)))) && (\because \text{triangle inequality}) \\
&= \text{SC}(b) + \sum_{i \in V} d(b, \text{top}(M(i))) \\
&= \text{SC}(b) + \sum_{i \in V} d(b, \text{top}(i)) && (\because M \text{ is a perfect matching}) \\
&= \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} d(b, \text{top}(i)) && (\because d(b, b) = 0) \\
&\leq \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} (d(b, i) + d(i, \text{top}(i))) && (\because \text{triangle inequality}) \\
&\leq \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} (d(b, i) + \alpha \cdot d(i, b)) && (\because \alpha\text{-decisiveness}) \\
&\leq \text{SC}(b) + \sum_{i \in V} (d(b, i) + \alpha \cdot d(i, b)) \\
&= (2 + \alpha) \cdot \text{SC}(b).
\end{aligned}$$

Because the metric space, the election, and the choice of b were arbitrary, this establishes that PLURALITYMATCHING has distortion at most $2 + \alpha$.

To show that this analysis is tight for $m \geq 3$, consider an election with two voters ($V = \{1, 2\}$) and three candidates ($C = \{a, b, c\}$). The preference profile is as follows: $\sigma_1 = a \succ b \succ c$ and $\sigma_2 = c \succ b \succ a$. First, note that the integral domination graph of b has two edges: $(1, 2)$ and $(2, 1)$. Both edges together form a perfect matching. Hence, PLURALITYMATCHING may return b . However, consider an α -decisive metric consistent with the preference profile given by the following undirected graph, where the distance between any two points is the shortest distance in the graph.



Note that $\text{SC}(b) = 2 + \alpha$, while $\text{SC}(c) = 1$.¹ Hence, the distortion of PLURALITYMATCHING is at least $2 + \alpha$. For $m > 3$, we can derive the same lower bound by placing the remaining candidates sufficiently far away. \square

Note that in the trivial case of $m = 2$, PLURALITYMATCHING (and every other reasonable voting rule) coincides with the majority rule, which picks the top choice of a majority of voters. This is known to have distortion $1 + 2\alpha$, which is optimal for deterministic rules for every $\alpha \in [0, 1]$ [6].

Corollary 2. *For every m , the distortion of PLURALITYMATCHING is 3, which is optimal across all deterministic social choice rules.*

This resolves the open conjecture by establishing that the optimal distortion of deterministic rules is 3, and can be achieved by a polynomial-time computable combinatorial rule.

While our bound is tight for every $\alpha \in [0, 1]$ when $m = 2$, and for $\alpha = 1$ when $m \geq 3$, one might wonder if it is also tight for $\alpha < 1$ when $m \geq 3$. The best known lower bound for deterministic

¹In this case, c is also a candidate that PLURALITYMATCHING may return, leading one to hope that tie-breaking might help reduce the distortion. However, in Appendix C, we provide a more complex construction which shows that the lower bound effectively holds regardless of the tie-breaking.

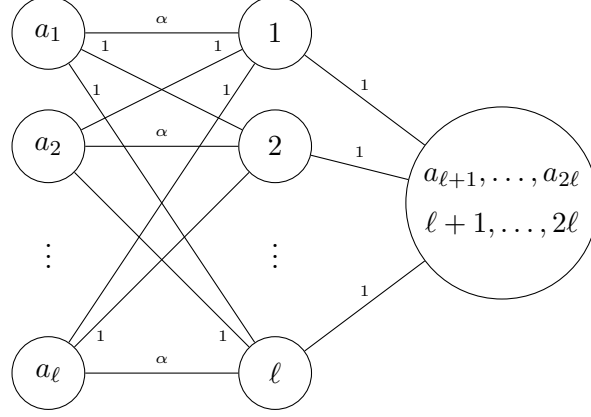


Figure 2: The metric space construction used in the proof of Theorem 2.

rules is $1 + 2\alpha$ [6], which could be significantly lower than our upper bound of $2 + \alpha$. We derive an improved lower bound for deterministic rules, which shows that at least as $m \rightarrow \infty$, our upper bound is tight for every $\alpha \in [0, 1]$. Note that our lower bound coincides with the aforementioned $1 + 2\alpha$ lower bound when m is 2 or 3, but improves as m increases.

Theorem 2. *For every $m \geq 2$ and $\alpha \in [0, 1]$, the distortion of every deterministic social choice rule is at least $2 + \alpha - 2(1 - \alpha)/\lfloor m \rfloor_{\text{even}}$ for α -decisive metric spaces, where $\lfloor m \rfloor_{\text{even}}$ is the largest even integer smaller than or equal to m .*

Proof. We show the lower bound of $2 + \alpha - 2(1 - \alpha)/m$ for every even m ; for odd m , we can simply add an irrelevant candidate that is sufficiently far away from all other candidates and voters.

Assume that $m = 2\ell$ for some $\ell \in \mathbb{N}$. Consider an election $\mathcal{E} = (V, C, \sigma)$, where $V = \{1, \dots, 2\ell\}$, $C = \{a_1, \dots, a_{2\ell}\}$, and the preference profile σ is as follows. Define $A = \{a_1, \dots, a_\ell\}$, and $B = C \setminus A$, $V_A = \{1, \dots, \ell\}$, and $V_B = V \setminus V_A$. For $i \in V_A$, voter i ranks candidate $a_i \in A$ first, followed by the remaining candidates in A in an arbitrary order, followed by the candidates in B in an arbitrary order. For $i \in V_B$, voter i ranks candidate $a_i \in B$ first, followed by the remaining candidates in B in an arbitrary order, followed by the candidates in A in an arbitrary order.

Let f be a deterministic social choice rule. Without loss of generality, suppose f returns a candidate from A ; the argument is symmetric when f returns a candidate from B . We now construct an α -decisive metric space consistent with σ under which every candidate in A has an approximation ratio to social cost at least $2 + \alpha - 2(1 - \alpha)/m$. The metric is represented by the graph in Figure 2

For $i \in V_A$, voter i has distance α from a_i , and 1 from every other candidate. Voters in V_B and candidates in B are located at a single point at distance 1 from each voter in V_A . The distance between any other pair of points is the shortest distance between them in the resulting graph.

It is easy to check that this metric is α -decisive and consistent with σ . Further, we have that for each $a_i \in A$, $\text{SC}(a_i) = 1 \cdot \alpha + (\ell - 1) \cdot 1 + \ell \cdot (1 + \alpha)$, whereas for each $a_i \in B$, $\text{SC}(a_i) = \ell \cdot 1$. Since f returns a candidate from A , its distortion is at least

$$\frac{\alpha + \ell - 1 + \ell \cdot (1 + \alpha)}{\ell} = 2 + \alpha - \frac{1 - \alpha}{\ell} = 2 + \alpha - \frac{2 \cdot (1 - \alpha)}{m},$$

as desired. □

Note that for even m , in the construction used in the proof of Theorem 2 each voter ranks a

distinct candidate first. Hence, the lower bound applies even in the special case of peer selection, in which the set of voters is identical to the set of candidates (and thus $\alpha = 0$).

For odd m , the bound can be slightly improved by having more voters whose top choices are slightly unequally divided between two subgroups of candidates of sizes $\lfloor m/2 \rfloor$ and $\lceil m/2 \rceil$. However, the improvement is marginal, so we avoid it for ease of presentation.

4.1 Connecting PLURALITYMATCHING to The Matching Uncovered Set

As mentioned in Section 1, we devised the notion of integral domination graphs by directly building on the work of Munagala and Wang [44]. They combinatorially defined a set of candidates that they named the matching uncovered set, and proved that every candidate in this set has distortion at most 3, but failed to answer whether this set is always non-empty. To understand how this is connected to our work, let us begin by introducing their definitions.

Definition 6. Given an election \mathcal{E} , the *separation graph* of candidates a and b is defined as the bipartite graph $G^\mathcal{E}(a, b) = (V, V, E_{a,b})$ where $(i, j) \in E_{a,b}$ if there exists a candidate $c \in C$ such that $a \succ_i c$ and $c \succ_j b$. When the election is clear from the context, we will drop it from the notation.

Definition 7. The *matching uncovered set* of an election \mathcal{E} is the set of candidates $a \in C$ such that for every $b \in C$, $G^\mathcal{E}(a, b)$ admits a perfect matching.

Note that the separation graph $G(a, b)$ has an edge from voter i on the left to voter j on the right if a weakly defeats some candidate in i 's vote who weakly defeats b in j 's vote. In comparison, recall that the integral domination graph $G(a)$ has an edge from i to j if a weakly defeats $\text{top}(j)$ in i 's vote. Because $\text{top}(j)$ clearly weakly defeats every candidate in j 's vote, an $i \rightarrow j$ edge in $G(a)$ also exists in $G(a, b)$ for every candidate b . That is, for each a , the integral domination graph $G(a)$ is a subgraph of the separation graph $G(a, b)$ for every b . Thus, seeking a perfect matching in $G(a)$, as PLURALITYMATCHING does, is equivalent to seeking a single set of edges that form a perfect matching in $G(a, b)$ for every b . This is stronger than the requirement imposed by the matching uncovered set, which is simply that each $G(a, b)$ admits a (potentially different) perfect matching. We therefore have the following.

Proposition 1. *PLURALITYMATCHING always returns a candidate in the matching uncovered set. Consequently, the matching uncovered set is always non-empty.*

When $\alpha = 1$, Proposition 1, along with the proof by Munagala and Wang [44] that every candidate in the matching uncovered set has distortion at most 3, is sufficient to establish that PLURALITYMATCHING has distortion at most 3.

The benefits of our alternative analysis of PLURALITYMATCHING in Theorem 1 are threefold. First, our derivation of the upper bound on distortion is much simpler than the tedious derivation by Munagala and Wang [44]. Second, as we show in Appendix D, our analysis seamlessly extends to produce an upper bound on the so-called *fairness ratio* [31], which is not known to be possible with the previous argument. More importantly, our analysis yields an upper bound of $2 + \alpha$ on distortion, which, for $\alpha < 1$, is stronger than the upper bound of 3 established by Munagala and Wang [44]. However, one might wonder whether the stronger guarantees of the Ranking-Matching Lemma are not really necessary to achieve this improved bound. Could it be the case that selecting any candidate from the matching uncovered set already achieves distortion at most $2 + \alpha$, even when $\alpha < 1$? We answer this negatively by providing two different lower bounds in Propositions 2 and 3.

Proposition 2. *For every $m \geq 5$ and $\alpha \in [0, 1]$, selecting an arbitrary candidate from the matching uncovered set yields distortion at least $2.5 + \alpha/2$ for α -decisive metric spaces.*

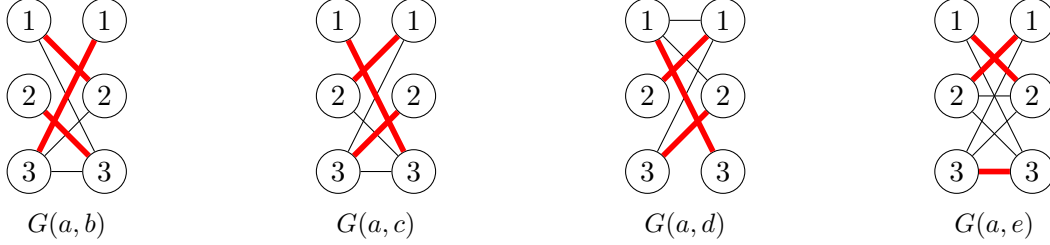


Figure 3: Separation graphs in the proof of Proposition 2. Thick red edges show perfect matchings.

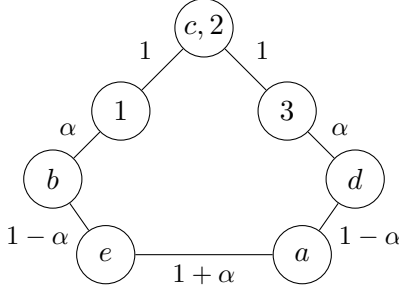


Figure 4: An α -decisive metric space used in the proof of Proposition 2.

Proof. Below, we produce an election with $m = 5$ candidates in which some candidate from the matching uncovered set yields approximation ratio at least $2.5 + \alpha/2$. For $m > 5$, we can simply add dummy candidates that are sufficiently far from all the voters.

Consider an election $\mathcal{E} = (V, C, \sigma)$ with 5 candidates ($C = \{a, b, c, d, e\}$) and 3 voters ($V = \{1, 2, 3\}$) having the following preferences:

$$\begin{aligned} \sigma_1 &: b \succ e \succ c \succ a \succ d \\ \sigma_2 &: c \succ d \succ b \succ a \succ e \\ \sigma_3 &: d \succ a \succ c \succ b \succ e \end{aligned}$$

First, we argue that a is in the matching uncovered set. Figure 3 shows separation graphs $G(a, b)$, $G(a, c)$, $G(a, d)$, and $G(a, e)$, with a perfect matching in each highlighted with thick red edges.

However, consider the metric space shown in Figure 4; as usual, the distance between any pair of points is the shortest distance in the graph shown. It can be checked easily that this metric is α -decisive and consistent with σ . Note that under this metric, $\text{SC}(a) = 5 + \alpha$, whereas $\text{SC}(c) = 2$. Thus, a rule picking a would have distortion at least $(5 + \alpha)/2 = 2.5 + \alpha/2$.

A curious reader may check that in this example, the integral domination graph of a does not admit a perfect matching, while those of c and d do. Thus, PLURALITYMATCHING may pick either of c or d . While c is optimal, d has approximation ratio $2 + \alpha$, matching our upper bound. \square

In fact, for large m , we can prove a better lower bound on the distortion that can be achieved by picking an arbitrary candidate from the matching uncovered set, or even by picking a special candidate from this set that has received significant attention in the voting literature.

Definition 8. Given an election $\mathcal{E} = (V, C, \sigma)$, we say that candidate a is a *Condorcet winner* if for every other candidate b , there are at least $n/2$ voters (a majority) who prefer a to b . A Condorcet winner is not guaranteed to exist. A deterministic social choice rule is called *Condorcet consistent* if it outputs a Condorcet winner on every election in which one exists.

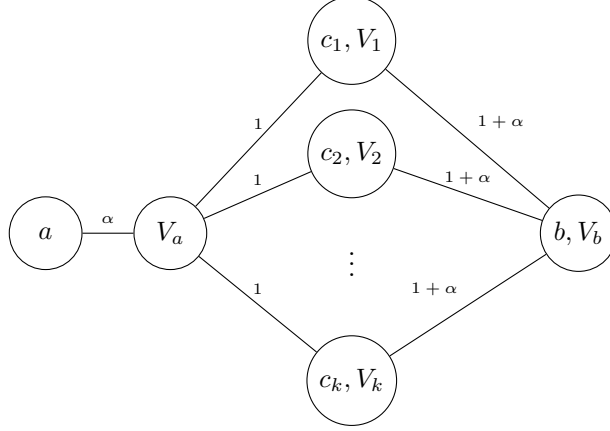


Figure 5: An α -decisive metric space used in the proof of Proposition 3.

It is easy to argue that every Condorcet winner a is in the matching uncovered set.² Thus, any lower bound on distortion that applies to social choice rules that pick a Condorcet winner when one exists (i.e. to Condorcet consistent rules) also applies to picking an arbitrary candidate from the matching uncovered set. The following result establishes such a lower bound.

Proposition 3. *For every $m \geq 2$ and $\alpha \in [0, 1]$, the distortion of every deterministic Condorcet consistent rule is at least $3 - \frac{4(3+\alpha)}{m+2+m\alpha}$ for α -decisive metric spaces. The same lower bound applies to selecting an arbitrary candidate from the matching uncovered set.*

Proof. Let $k = m - 2$. Consider an election $\mathcal{E} = (V, C, \sigma)$, where the set of candidates is $C = \{a, b, c_1, \dots, c_k\}$, and the set of voters is $V = \cup_{c \in C} V_c$ with $|V_a| = 2$, $|V_b| = k$, and $|V_c| = 1$ for each $c \in C \setminus \{a, b\}$. The preference profile σ is as follows.

$$\begin{aligned} \forall i \in V_a, & \quad \sigma_i : a \succ c_1 \succ \dots \succ c_k \succ b \\ \forall i \in V_b, & \quad \sigma_i : b \succ c_1 \succ \dots \succ c_k \succ a \\ \forall r \in [k], i \in V_{c_r}, & \quad \sigma_i : c_r \succ a \succ b \succ c_1 \succ \dots \succ c_{r-1} \succ c_{r+1} \succ \dots \succ c_k \end{aligned}$$

It is easy to check that in this preference profile, a is the unique Condorcet winner. However, consider the metric space shown in Figure 5. As usual, all pairwise distances are the corresponding shortest distances in the graph.

One can check that this is an α -decisive metric space consistent with σ . However, $SC(a) = 2\alpha + k \cdot (1+\alpha) + k \cdot (2+2\alpha) = 3k + (3k+2)\alpha$, whereas $SC(b) = k \cdot (1+\alpha) + 2 \cdot (2+\alpha) = k+4 + (k+2)\alpha$. Hence, the distortion of a rule that selects a on this preference profile is at least

$$\frac{3k + (3k + 2)\alpha}{k + 4 + (k + 2)\alpha} = 3 - \frac{4(3 + \alpha)}{m + 2 + m\alpha},$$

where the last step uses the fact that $k = m - 2$. □

Anshelevich et al. [11] proved that picking a Condorcet winner, when one exists, guarantees 3-approximation to social cost. Thus, if we were only interested in bounding the distortion for general metric spaces ($\alpha = 1$), then being Condorcet consistent would not hurt: indeed, any social

²This is because for every other candidate b , a perfect matching in $G(a, b)$ can be obtained easily by matching the (at most $n/2$) voters on the right who prefer b over a to (some of at least $n/2$) voters on the left who prefer a over b , and matching the remaining voters on the right to the remaining voters on the left arbitrarily.

choice rule can be made Condorcet consistent by forcing it to pick a Condorcet winner on profiles where one exists without increasing its distortion.

However, if our goal was to also achieve improved distortion bounds parametrized by α , Proposition 3 shows that being Condorcet consistent is at odds with this goal. When $m \rightarrow \infty$, the lower bound in Proposition 3 converges to 3 regardless of the value of α . In contrast, recall that PLURALITYMATCHING achieves distortion $2 + \alpha$ for any m . Thus, while PLURALITYMATCHING is not Condorcet consistent,³ this is necessary for it to achieve the improved distortion. Similarly, Propositions 2 and 3 also show that picking an arbitrary candidate from the matching uncovered set yields a worse parametrized distortion bound than that of PLURALITYMATCHING.

5 The Distortion of Randomized Social Choice Rules

In the previous section, we settled the question of determining the optimal distortion of deterministic social choice rules, and showed that it is 3 for all $m \geq 2$. The obvious next step is to study the power of randomization in the metric distortion framework, and investigate whether it allows achieving any lower distortion. In this section, we show that randomization indeed helps, but leave open the question of determining the optimal distortion of randomized rules.

Randomized rules were first considered in this framework by Anshelevich and Postl [6]. They studied a natural rule called RANDOMDICTATORSHIP, which chooses one of the voters uniformly at random, and outputs the top candidate of that voter; in other words, each candidate a is chosen with probability equal to $\text{plu}(a) / \sum_{c \in C} \text{plu}(c)$, or $\text{plu}(a)/n$. They proved that its distortion is $2 + \alpha$.⁴ Given our result (Theorem 1) that PLURALITYMATCHING achieves distortion $2 + \alpha$ deterministically, this raises the question of whether randomization helps.

For general metric spaces, i.e., when $\alpha = 1$, Kempe [38] answered this question by proposing a different randomized voting rule which achieves a slightly better distortion of $3 - 2/m$. This rule, like RANDOMDICTATORSHIP, uses only the plurality scores of the candidates, and it is known that $3 - 2/m$ is the best possible distortion that any rule can achieve for $\alpha = 1$ using only the plurality scores [33, 31].

What about the case of $\alpha < 1$? Anshelevich and Postl [6] proved that a randomized social choice rule, even using the full preference rankings, cannot achieve distortion better than $1 + \alpha$ when $m \geq 2$. Interestingly, this bound is not met by RANDOMDICTATORSHIP even for the seemingly trivial case of $m = 2$ candidates, where each vote is one of only two possible rankings. To bridge this gap, specifically for the case of $m = 2$ candidates (say a and b), they proposed a rule, α -GENERALIZED PROPORTIONAL TO SQUARES, which achieves the optimal distortion of $1 + \alpha$. This rule chooses alternative a with probability

$$\frac{(1 + \alpha) \cdot \text{plu}(a)^2 - (1 - \alpha) \cdot \text{plu}(a) \cdot \text{plu}(b)}{(1 + \alpha) \cdot (\text{plu}(a)^2 + \text{plu}(b)^2) - 2(1 - \alpha) \cdot \text{plu}(a) \cdot \text{plu}(b)}, \quad (2)$$

and b with the remaining probability.

Our main contribution in this section is twofold. We propose a novel randomized social choice rule, SMARTDICTATORSHIP, which achieves distortion at most $2 + \alpha - 2/m$ (Theorem 3). This

³In Appendix B, we present a simple necessary condition for a candidate to be possibly returned by PLURALITYMATCHING, and a simpler example than the one used in the proof of Proposition 3 where a Condorcet winner violates this condition.

⁴Specifically, they proved a distortion bound of $2 + \alpha - 2/n$ for n voters. But since we defined distortion using supremum over preference profiles of all sizes, it becomes $\sup_{n \geq 1} 2 + \alpha - 2/n = 2 + \alpha$ in our framework. That said, the distortion of RANDOMDICTATORSHIP is strictly less than 3 for any finite n , which is impossible for deterministic rules.

matches the distortion of $3 - 2/m$ achieved by Kempe [38] for $\alpha = 1$, but improves when $\alpha < 1$. SMARTDICTATORSHIP, like RANDOMDICTATORSHIP, α -GENERALIZED PROPORTIONAL TO SQUARES, and the rule proposed by Kempe [38], uses only the plurality scores. In fact, for $m = 2$, it coincides with α -GENERALIZED PROPORTIONAL TO SQUARES and achieves the optimal distortion of $1 + \alpha$. Further, we show that no randomized social choice rule which uses only the plurality scores can achieve distortion lower than $2 + \alpha - 2/m$ (Theorem 4). Thus, SMARTDICTATORSHIP achieves the optimal distortion achievable using only plurality scores *for every* α . We also improve upon the $1 + \alpha$ lower bound of Anshelevich and Postl [6] that applies to rules which are allowed to use the full preference rankings (Theorem 5).

Recall that in Theorem 2, we showed that deterministic rules cannot achieve distortion better than $2 + \alpha - 2(1 - \alpha)/m$ for even m . Note that $2 + \alpha - 2/m$, which is the distortion achieved by SMARTDICTATORSHIP, is strictly smaller when $\alpha > 0$. Hence, this establishes a strict separation between deterministic and randomized rules for every $\alpha > 0$ (which was previously only established for $\alpha = 1$), but leaves open the question when $\alpha = 0$.

5.1 Randomized SMARTDICTATORSHIP Rule

We begin by defining SMARTDICTATORSHIP.

Definition 9. Let $\mathcal{E} = (V, C, \sigma)$ be an election. If there exists a candidate a with $\text{plu}(a) \geq (1 + \alpha) \cdot \frac{n}{2}$, SMARTDICTATORSHIP arbitrarily chooses one such candidate with probability 1. Otherwise, it chooses each candidate a with probability proportional to $\text{plu}(a) / \left(n - \frac{2}{1 + \alpha} \cdot \text{plu}(a)\right)$, i.e., with probability equal to

$$\frac{\text{plu}(a)}{\left(n - \frac{2}{1 + \alpha} \cdot \text{plu}(a)\right) \sum_{c \in C} \frac{\text{plu}(c)}{n - \frac{2}{1 + \alpha} \cdot \text{plu}(c)}}. \quad (3)$$

We leave it as an exercise to the reader to verify that the expression in Equation (3) coincides with the expression in Equation (2) when $C = \{a, b\}$. We also note that for $\alpha = 1$, SMARTDICTATORSHIP simply selects each candidate a with probability proportional to $\text{plu}(a) / (n - \text{plu}(a))$. While the distortion of SMARTDICTATORSHIP for $\alpha = 1$, proved in Theorem 3, matches the distortion of the rule proposed by Kempe [38], the two rules are different when $m \geq 3$.

In order to prove the distortion bound for SMARTDICTATORSHIP, we make use of the following results by Anshelevich and Postl [6].

Lemma 4. [6] *Let $\mathcal{E} = (V, C, \sigma)$ be an election induced by an α -decisive metric space with distance metric d . Let c^* be an optimal candidate. Then, the approximation ratio of a randomized social choice rule choosing each candidate a with probability $\Pr[a]$ on this instance is at most*

$$1 + \frac{(1 + \alpha) \sum_{a \in C} \Pr[a] \cdot \left(n - \frac{2}{1 + \alpha} \cdot \text{plu}(a)\right) \cdot d(a, c^*)}{\sum_{a \in C} \text{plu}(a) \cdot d(a, c^*)}.$$

As an intermediate step toward proving this lemma, they also showed the following.

Proposition 4. [6] *Let $\mathcal{E} = (V, C, \sigma)$ be an election induced by an α -decisive metric space with distance metric d . Let c^* be an optimal candidate. Then, for every candidate a ,*

$$\text{SC}(a) \leq \text{SC}(c^*) + \left(n - \frac{2}{1 + \alpha} \cdot \text{plu}(a)\right) \cdot d(a, c^*).$$

We also prove the following lemma that will help us simplify our proof of the distortion bound.

Lemma 5. For any $w \in (0, 1]$, $n \in \mathbb{N}$, and $(x_i)_{i=1}^m \in [0, n/w]^m$ such that $\sum_{i=1}^m x_i = n$,

$$\sum_{i=1}^m \frac{x_i}{n - w \cdot x_i} \geq \frac{m}{m - w}.$$

Proof. We will show that the summation on the left hand side of the lemma's statement is minimized when $x_i = n/m$ for all $i \in \{1, \dots, m\}$, which will imply that

$$\sum_{i=1}^m \frac{x_i}{n - w \cdot x_i} \geq \sum_{i=1}^m \frac{n/m}{n - w \cdot n/m} = \sum_{i=1}^m \frac{1}{m - w} = \frac{m}{m - w}.$$

Let \mathcal{X} denote the set of feasible vectors $x = (x_i)_{i=1}^m$ that minimize the summation on the left. Aiming for a contradiction, assume that for every vector $x \in \mathcal{X}$, there exist two values, x_j and x_k , such that $x_j \neq x_k$. Let x^* be a vector in \mathcal{X} that is leximin-optimal across all vectors in \mathcal{X} . That is, we choose x^* from \mathcal{X} by maximizing the smallest x_i , then breaking ties in favor of higher second-smallest x_i , then in favor of higher third-smallest x_i , and so on. In other words, we sort each vector in a non-decreasing order, and then compare them lexicographically, picking a maximum.

Take any two values x_j^* and x_k^* such that $x_j^* > x_k^*$, and consider their contribution to the sum $\sum_{i=1}^m \frac{x_i^*}{n - wx_i^*}$, which is equal to

$$\frac{x_j^*}{n - wx_j^*} + \frac{x_k^*}{n - wx_k^*} = \frac{(n - wx_k^*)x_j^* + (n - wx_j^*)x_k^*}{(n - wx_j^*)(n - wx_k^*)} = \frac{n(x_j^* + x_k^*) - 2wx_j^*x_k^*}{n^2 - n(x_j^* + x_k^*) + w^2x_j^*x_k^*}. \quad (4)$$

Now, consider an alternative vector $x' = (x'_i)_{i=1}^m$ such that $x'_j = x'_k = (x_j^* + x_k^*)/2$ and $x'_i = x_i^*$ for all $i \neq j, k$. Clearly, these values also satisfy $\sum_{i=1}^m x'_i = n$ since this was satisfied by x^* , only values of j and k changed, and $x'_j + x'_k = x_j^* + x_k^*$. Also, since $x^* \in [0, n/w]^m$ and $x_k^* < x'_j = x'_k < x_j^*$, we have that $x' \in [0, n/w]^m$ (thus x' is a feasible vector) and that x' is strictly better in the leximin-ordering than x^* .

Finally, due to the arithmetic mean-geometric mean inequality,

$$\sqrt{x_j^*x_k^*} \leq \frac{x_j^* + x_k^*}{2} \Rightarrow x_j^*x_k^* \leq \left(\frac{x_j^* + x_k^*}{2}\right)^2 = x'_jx'_k.$$

Therefore, $x_j^* + x_k^* = x'_j + x'_k$ and $x_j^*x_k^* \leq x'_jx'_k$. By observing the right hand side of Equation (4), and noting the fact that the numerator is decreasing in $x_j^*x_k^*$ and the denominator is decreasing in $x_j^*x_k^*$, we conclude that

$$\begin{aligned} \frac{n(x_j^* + x_k^*) - 2wx_j^*x_k^*}{n^2 - n(x_j^* + x_k^*) + w^2x_j^*x_k^*} &\leq \frac{n(x'_j + x'_k) - 2wx'_jx'_k}{n^2 - n(x'_j + x'_k) + w^2x'_jx'_k} \\ \Rightarrow \frac{x_j^*}{n - wx_j^*} + \frac{x_k^*}{n - wx_k^*} &\leq \frac{x'_j}{n - wx'_j} + \frac{x'_k}{n - wx'_k} \\ \Rightarrow \sum_{i=1}^m \frac{x_i^*}{n - wx_i^*} &\leq \sum_{i=1}^m \frac{x'_i}{n - wx'_i}. \end{aligned}$$

Hence, x' is a feasible vector minimizing the summation, and is therefore in \mathcal{X} . Further, it is strictly better than x^* in the leximin ordering, which contradicts the fact that x^* is leximin-optimal among vectors in \mathcal{X} . \square

Using these results, we are now ready to prove the improved distortion bound that we can achieve using SMARTDICTATORSHIP.

Theorem 3. *For every $m \geq 2$ and $\alpha \in [0, 1]$, SMARTDICTATORSHIP has distortion at most $2 + \alpha - 2/m$ for α -decisive metric spaces.*

Proof. Let $\mathcal{E} = (V, C, \sigma)$ be an election. If there exists any candidate a with $\text{plu}(a) \geq (1 + \alpha) \cdot \frac{n}{2}$, note that SMARTDICTATORSHIP arbitrarily chooses one such candidate with probability 1. However, using Proposition 4, we can conclude that for this candidate c , we have $\text{SC}(c) \leq \text{SC}(c^*)$, where c^* is an optimal candidate. Hence, any such candidate is optimal as well, yielding an approximation ratio of 1.

If this is not the case, then we substitute the probability $\Pr[a]$ with which SMARTDICTATORSHIP chooses each candidate a into Lemma 4, and get that the approximation ratio is at most

$$1 + \frac{(1 + \alpha) \sum_{a \in C} \text{plu}(a) d(a, c^*)}{\left(\sum_{c \in C} \frac{\text{plu}(c)}{n - \frac{2}{1+\alpha} \cdot \text{plu}(c)} \right) \cdot \left(\sum_{a \in C} \text{plu}(a) d(a, c^*) \right)} = 1 + \frac{1 + \alpha}{\sum_{c \in C} \frac{\text{plu}(c)}{n - \frac{2}{1+\alpha} \cdot \text{plu}(c)}}.$$

Using Lemma 5 with $w = \frac{2}{1+\alpha}$ and using the fact that $\text{plu}(c) \in [0, (1 + \alpha) \cdot \frac{n}{2})$ for each $c \in C$ and $\sum_{c \in C} \text{plu}(c) = n$, we get that $\sum_{c \in C} \frac{\text{plu}(c)}{n - \frac{2}{1+\alpha} \cdot \text{plu}(c)} \geq \frac{m}{m - \frac{2}{1+\alpha}}$. We therefore conclude that

$$\text{distortion}(\text{SMARTDICTATORSHIP}) \leq 1 + (1 + \alpha) \frac{m - \frac{2}{1+\alpha}}{m} = 2 + \alpha - \frac{2}{m},$$

as desired. \square

A few remarks are in order. First, for $\alpha = 1$ (i.e. general metric spaces), the bound is $3 - 2/m < 3$, which shows a separation between randomized and deterministic rules for every $m \geq 2$. This is despite SMARTDICTATORSHIP only accessing plurality votes, while the deterministic rules being allowed to access the full ranked preference profile.⁵ Second, as mentioned earlier, even as a function of α , the upper bound of $2 + \alpha - 2/m$ for SMARTDICTATORSHIP is still better than the lower bound of $2 + \alpha - 2(1 - \alpha)/m$ for deterministic rules when $m > 3$ is even and $\alpha > 0$. However, whether randomized rules can truly outperform deterministic rules when $\alpha = 0$, or when $\alpha = 1$ but $m \rightarrow \infty$, is still open.

5.2 Lower Bounds

Note that both SMARTDICTATORSHIP and RANDOMDICTATORSHIP access only the plurality scores of the candidates, in contrast to the deterministic PLURALITYMATCHING from the previous section which requires knowledge of the full preference profile. We first show that among randomized rules which only take the plurality scores as input, SMARTDICTATORSHIP is in fact optimal, even as a function of α . This parametrizes the lower bound of $3 - 2/m$ derived in prior work for $\alpha = 1$ [33, 31].

Theorem 4. *For every $m \geq 2$ and $\alpha \in [0, 1]$, no randomized social choice rule with access to only plurality scores has distortion better than $2 + \alpha - 2/m$ for α -decisive metric spaces.*

Proof. Let f be any randomized social choice rule. Fix any m . Consider an instance with $n = m$ voters in which every candidate has plurality score 1, i.e., $\text{plu}(c) = 1$ for each $c \in C$. By the

⁵As mentioned earlier, this was also established by Kempe [38].

pigeonhole principle, there exists a candidate c^* that f picks with probability at most $1/m$ on this instance.

We now construct an underlying metric space consistent in which each candidate has plurality score 1, but choosing c^* with probability at most $1/m$ leads to bad distortion.

For the voter i whose top choice is c^* , we let $d(i, c^*) = 0$ and $d(i, c) = 1$ for every $c \neq c^*$. For every other voter i , we let $d(i, \text{top}(i)) = \alpha$, $d(i, c^*) = 1 + \alpha$, and $d(i, c) = 2 + \alpha$ for all $c \neq \text{top}(i), c^*$. It is easy to check that this forms a metric (let all other distances be the induced “shortest path” distances). Note that the social cost of candidate c^* is $\text{SC}(c^*) = \sum_{i \in V} d(i, c^*) = m - 1$, whereas for any other candidate $b \neq c^*$,

$$\begin{aligned} \text{SC}(b) &= \sum_{i \in V} d(i, b) = \sum_{i \in V: \text{top}(i)=b} d(i, b) + \sum_{i \in V: \text{top}(i)=c^*} d(i, b) + \sum_{i \in V: \text{top}(i) \neq b, c^*} d(i, b) \\ &= \alpha + (1 + \alpha) + (m - 2) \cdot (2 + \alpha). \end{aligned}$$

It is easy to verify that for all $m \geq 2$ we have $\text{SC}(c^*) \leq \text{SC}(b)$ for all $b \neq c^*$. Since f chooses c^* with probability at most $1/m$ and candidates from $C \setminus \{c^*\}$ with the remaining probability of at least $(m - 1)/m$, its distortion is at least

$$\begin{aligned} \frac{\frac{1}{m} \cdot \text{SC}(c^*) + \frac{m-1}{m} \cdot \text{SC}(b)}{\text{SC}(c^*)} &= \frac{(m - 1) + (m - 1)(\alpha + (1 + \alpha) + (m - 2)(2 + \alpha))}{m \cdot (m - 1)} \\ &= \frac{2m + m\alpha - 2}{m} \\ &= 2 + \alpha - \frac{2}{m}, \end{aligned}$$

as desired. □

In terms of general lower bounds, when the randomized rule is given access to the full preference profile, Anshelevich and Postl [6] showed that no randomized rule can achieve distortion better than $1 + \alpha$ for $m \geq 2$. Our next result provides an improved bound, which matches the bound of $1 + \alpha$ for $m = 2$ (which is optimal for $m = 2$), but increases to $(3 + \alpha)/2$ as $m \rightarrow \infty$.

Theorem 5. *For every $m \geq 2$ and $\alpha \in [0, 1]$, no randomized social choice rule has distortion better than $\frac{3+\alpha}{2} - \frac{1-\alpha}{\lfloor m \rfloor_{\text{even}}}$ for α -decisive metric spaces, where $\lfloor m \rfloor_{\text{even}}$ is the largest even integer that smaller than or equal to m .*

Proof. We use the same construction that we used in the proof of Theorem 2. For even m , recall that the construction had two equal-sized groups of candidates A and B . There were two possible consistent α -decisive metric spaces, one in which each candidate in A is optimal and each candidate in B is $2 + \alpha - 2(1 - \alpha)/m$ times worse, and the other in which the roles of A and B are reversed.

Due to the symmetry, assume that a given randomized rule places a total probability of at least 0.5 on candidates from A . Then, in the latter metric space, the rule chooses an optimal candidate with probability at most 0.5 and a candidate with approximation ratio $2 + \alpha - 2(1 - \alpha)/m$ with the remaining probability of at least 0.5. Therefore, its distortion is at least

$$\frac{1}{2} \left(2 + \alpha - \frac{2(1 - \alpha)}{m} \right) + \frac{1}{2} = \frac{3 + \alpha}{2} - \frac{1 - \alpha}{m}.$$

For odd m , we use the same trick as in the proof of Theorem 2: we use the construction for $m - 1$, and adding a candidate sufficiently far away, which leads to the claimed bound. □

One of the most compelling open problems is whether there exists a randomized rule that can break the barrier of $2 + \alpha$ even as $m \rightarrow \infty$. When $m \rightarrow \infty$, the distortion of $2 + \alpha$ can be achieved by both `RANDOMDICTATORSHIP` and `SMARTDICTATORSHIP` in a randomized fashion, as well as deterministically by `PLURALITYMATCHING`. Theorem 5 implies that a distortion bound better than $\frac{3+\alpha}{2}$ would not be possible in this case, leaving some room for improvement. Of particular interest is the case of $\alpha = 1$ for which the best possible distortion is in $[2, 3]$.

In order to achieve a bound better than $2 + \alpha$, we need a rule that is randomized and, as Theorem 4 shows, accesses more information regarding the preference profile than just the plurality scores of the candidates. A tempting thought is that an appropriate (randomized) combination of a known randomized rule with our deterministic `PLURALITYMATCHING` rule, which depends on a lot more information beyond the plurality scores, could be successful. The following lower bound shows that this approach cannot yield improved bounds when `PLURALITYMATCHING` is combined with a randomized rule that only places positive probability on candidates with positive plurality score (as both `RANDOMDICTATORSHIP` and `SMARTDICTATORSHIP` do). This suggests that any successful solution would need to, at least for some instances, assign positive probability to candidates that have zero plurality score and do not satisfy the conditions of the `PLURALITYMATCHING` rule.

Theorem 6. *For any $\alpha \in [0, 1]$, every randomized social choice rule that randomizes over only candidates that could be returned by `PLURALITYMATCHING` and candidates with non-zero plurality score has distortion at least $2 + \alpha - 2/(m - 2)$ for α -decisive metric spaces.*

Proof. We consider an instance with a set of $m = \ell + 2$ candidates, $\{a_1, a_2, \dots, a_{\ell+1}, a^*\}$, and a set of voters that are partitioned into $\ell + 1$ subsets, $V_1, V_2, \dots, V_{\ell+1}$. The set $V_{\ell+1}$ contains a single voter, and each of the other ℓ sets contain exactly k voters, for some arbitrarily large $k \in \mathbb{N}$, so the total number of voters is $n = k\ell + 1$. The distances of the voters from the candidates are implied by shortest path distances in the graph of Figure 6, so for the voters of each group V_i for $i \leq \ell$, their ranking is a_i first, a^* second, and $a_{\ell+1}$ last, with some arbitrary ordering of the remaining candidates in-between a^* and $a_{\ell+1}$. On the other hand, the ranking of the single voter in $V_{\ell+1}$ is $a_{\ell+1}$ first, a^* last, and some arbitrary ordering of the remaining candidates in-between.

Note that candidate a^* has the optimal social cost, with $\text{SC}(a^*) = k\ell + 3 + \alpha$, but the plurality score of this candidate is zero. So, this candidate will never be chosen by a social choice rule that randomizes over only candidates with non-zero plurality score. Also, note that the voter in $V_{\ell+1}$ ranks candidate a^* last, and no voter ranks a^* first. In other words it's *veto score* is less than its plurality score. We show in Lemma 6 in Appendix B that this implies a^* cannot be chosen by `PLURALITYMATCHING`.

Given the symmetry of the instance, it is easy to verify that the social cost among the remaining candidates is at least as much as the social cost of a_1 . The social cost of this candidate is $\text{SC}(a_1) = k\alpha + (\ell - 1)k(2 + \alpha) + 2 = \ell k\alpha + 2(\ell - 1)k + 2$. We therefore conclude that the distortion for this instance will be at least

$$\frac{\text{SC}(a_1)}{\text{SC}(a^*)} = \frac{\ell k\alpha + 2(\ell - 1)k + 2}{k\ell + 3 + \alpha},$$

which for $k \rightarrow \infty$ becomes

$$\frac{\ell\alpha + 2(\ell - 1)}{\ell} = 2 + \alpha - \frac{2}{\ell} = 2 + \alpha - \frac{2}{m - 2}. \quad \square$$

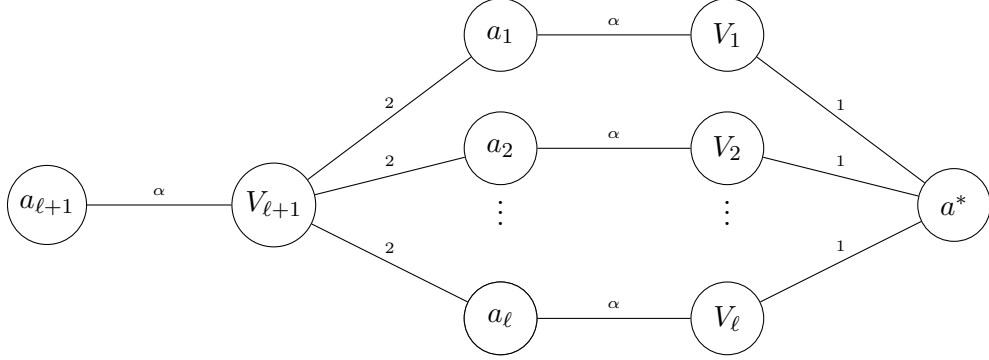


Figure 6: The metric space construction used in the proof of Theorem 6.

6 Other Related Work

The parametrization using α -decisiveness provides an interesting restriction on the underlying metric space, and the special case of $\alpha = 0$ captures settings where the sets of voters and candidates coincide, i.e., $V = C$. This is known as *peer selection*, which involves a set of individuals voting over themselves. Typically, peer selection is studied under binary votes with a focus on strategic manipulation [4], whereas our focus is on determining the desired selection rule (disregarding any manipulations) under ranked votes. The case where $V = C$ is also widely studied in theoretical computer science under the name of the *metric 1-median problem*: given a finite metric space (\mathcal{M}, d) , find a point in \mathcal{M} with the least total distance to all points in \mathcal{M} (think $V = C = \mathcal{M}$). Given all distances, the problem can be solved trivially in $O(|\mathcal{M}|^2)$ time, and a $(1 + \epsilon)$ -approximation can be computed in $O(|\mathcal{M}|/\epsilon^2)$ time through Monte-Carlo search [34, 35]. A different line of inquiry limits the number of pairwise distances the algorithm can query [18, 19, 20]; like in our metric distortion problem, the goal is to find the best approximation given limited information.

Prior work in voting has also considered imposing other well-motivated restrictions on the preferences of the voters, e.g., by considering Euclidean metrics [48, 6], as well as metrics inducing single-peaked [43] or single-crossing [30, 16] preferences. An alternative way of introducing more structure to the voter preferences is to assume the candidates are *representative* of the voters, i.e., drawn i.i.d. from the pool of voters [24, 23]. Note that in this case, each candidate is co-located with some voter, which is reminiscent of the special case $\alpha = 0$ discussed above. Pierczynski and Skowron [45] extend the idea of distortion to approval-based preferences, while Borodin et al. [13] and Filos-Ratsikas et al. [29] extend it to complex voting paradigms such as district elections and primaries. On the other end, prior work has also studied distortion in less constrained settings, e.g., ones that are not even bound by metric constraints [46, 14, 17].

Distortion minimization, at its core, is the simple idea of designing approximation algorithms that have access only to ordinal information about underlying cardinal values. One can therefore leverage the distortion framework to study a diverse set of problems or objectives. As an alternative objective, Fain et al. [28] focus on minimizing *squared distortion* of randomized rules, which is a notion that helps bound both the expectation and the variance of the approximation ratio of a rule at the same time. Goel et al. [31, 32] optimize the fairness ratio, which is a notion that is closely related to and always at least the distortion (in Appendix D, we show that a simple modification of our distortion analysis can also be used to achieve a bound on the fairness ratio of our deterministic rule). The distortion framework has also already been used to study a variety of different problems: Anshelevich and Sekar [7, 8], Abramowitz and Anshelevich [2], and Anshelevich and Zhu [10, 9]

design novel algorithms for problems such as matching, clustering, and facility location using only ordinal information. For the case of matching, recent work by Abebe et al. [1] also evaluated ordinal mechanisms compared to cardinal ones in strategic settings aiming to design truthful mechanisms.

Recent work has also considered moving beyond the cardinal versus ordinal comparison, aiming to get a more refined understanding of the value of information in different settings. Chen et al. [22], Fain et al. [28], Gross et al. [33] are interested in rules that elicit *even less* information from the voters than their full ranked preferences. On the contrary, Abramowitz et al. [3] consider rules that have access to *more* information than just the ranked preferences, specifically the ratio of distances of each voter to every pair of candidates. Extending this idea, Kempe [38], Mandal et al. [40], Mandal et al. [41], and Amanatidis et al. [5] study the explicit tradeoff between communication requirements and distortion of voting rules. More broadly, the *quantitative* evaluation of social choice rules using distortion, arguably a contribution of computer science to voting theory [46, 14], is a recent approach, and stands in stark contrast to centuries-long study of social choice rules through an *axiomatic* approach [15].

7 Discussion

In this paper, we settled the optimal metric distortion conjecture by proving that the optimal distortion of deterministic social choice rules is 3, and introducing a novel polynomial time computable rule PLURALITYMATCHING that achieves this.

While the bound of 3 is tight for general metric spaces, for the restrictive class of α -decisive metric spaces we showed that PLURALITYMATCHING achieves distortion $2 + \alpha$, and there is a lower bound of $2 + \alpha - 2(1 - \alpha)/m$ for all deterministic rules. While this establishes PLURALITYMATCHING as optimal for all α when $m \rightarrow \infty$, and for all m when $\alpha = 1$, it leaves open the question of bridging the bounds when m is finite and $\alpha < 1$.

On the randomized frontier, we proposed SMARTDICTATORSHIP, which achieves distortion $2 + \alpha - 2/m$, improving over the best known bound of $2 + \alpha$ achieved by RANDOMDICTATORSHIP. This also established a separation between the power of randomized and deterministic rules when $\alpha > 0$. However, determining whether randomized rules can outperform deterministic rules for $\alpha = 0$, and analyzing the optimal distortion of randomized rules remain challenging open questions.

Finally, the core strength behind our deterministic PLURALITYMATCHING rule is the novel Ranking-Matching Lemma we prove about matching rankings to candidates. This is an elementary lemma that could be of interest outside of the metric distortion framework as it induces a novel family of voting rules, some of which may have other interesting properties. For example, we defined the UNIFORMMATCHING rule in Section 3, which returns a candidate satisfying the following guarantee: *In any subset of $x\%$ of the votes, for any x , the selected (single) candidate weakly defeats at least $x\%$ of the candidates.*

This guarantee is eerily reminiscent of another major open question in voting theory. Consider selecting a subset, or *committee*, W of k candidates instead of a single candidate. This committee W is said to be in the *core* if no group S of at least n/k voters can find a candidate c^* such that each voter $i \in S$ prefers c^* to every candidate in the selected committee W . It is easy to check that this can be alternatively phrased as: *In any subset of n/k of the votes, the selected k candidates weakly defeat all (i.e. at least m) candidates.*

Thus, instead of finding one candidate that weakly defeats at least m/k candidates in any n/k votes, we want to find k candidates that (collectively) weakly defeat all m candidates in any n/k votes. While it is known that such a committee does not always exist [37], the question remains open if we relax the size of the group from n/k voters to $2n/k$ voters. Similarly, the question of whether

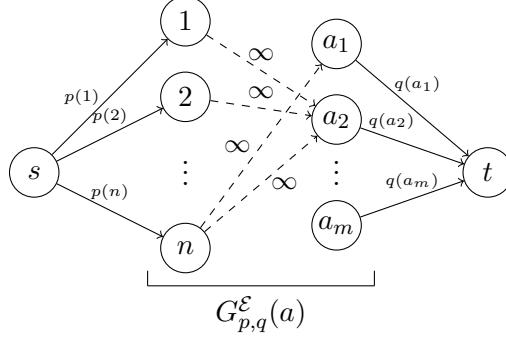


Figure 7: Flow network used to prove a generalization of the Hall’s condition in Lemma 1. Edges from source node s to each voter i has capacity $p(i)$, edges from each candidate c to target t has capacity $q(c)$, and edges from voters to candidates are as in $G_{p,q}^{\mathcal{E}}(a)$ with infinite capacity.

a committee in the core always exists is also open (without any relaxations required) when voters have binary preferences over candidates [27]. It is natural to think of recursively applying rules such as UNIFORMMATCHING or PLURALITYMATCHING to select committees. While the distortion of such a recursive application is easy to bound [32], whether it provides any core-like guarantees is an important question for the future.

Appendix

Below, we provide the missing proofs as well as some additional observations.

A Proof of Lemma 1

Proof. Let $\mathcal{E} = (V, C, \sigma)$ be an election, $p \in \Delta(V)$ and $q \in \Delta(C)$ be weight vectors, and $a \in C$ be a candidate. First, we want to show that $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching if and only if $q(D_a^{\mathcal{E}}(S)) \geq p(S)$ for all $S \subseteq V$.

This can be proved, similarly to the regular Hall’s condition, through the max-flow min-cut theorem. Create a flow network N shown in Figure 7. Note that we have converted all edges of $G_{p,q}^{\mathcal{E}}(a)$ to directed edges from voters to candidates with infinite capacity, added a source node with an outgoing edge of capacity $p(i)$ to each voter i , and added a target node with an incoming edge of capacity $q(c)$ from each candidate c .

We first argue that this network admits a flow of value 1 if and only if $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching. If N has a flow of value 1, then s has an outgoing flow of 1 and t has an incoming flow of 1. But since $\sum_{i \in V} p(i) = \sum_{c \in C} q(c) = 1$, this means the node for each voter i must have a flow of $p(i)$ passing through it, and the node of each candidate c must have a flow of $q(c)$ passing through it. A fractional perfect matching in $G_{p,q}^{\mathcal{E}}(a)$ can now be obtained by using the flow on each (i, c) edge as the weight of the edge. Conversely, if $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching w , then it is straightforward to see that having each (s, i) edge carry a flow of $p(i)$, each (i, c) edge carry a flow of $w(i, c)$, and each (c, t) edge carry a flow of $q(c)$, for every $i \in V$ and $c \in C$, results in a valid flow of value 1.

Next, because p and q are normalized weight vectors, the network has a flow of value 1 if and only if its maximum flow is 1. By the max-flow min-cut theorem, this is equivalent to its minimum

cut capacity being 1. Since $(\{s\}, V \cup C \cup \{t\})$ is a trivial cut with capacity 1, min cut being equal to 1 is equivalent to every cut having capacity at least 1.

We now show that every cut in this network has capacity at least 1 if and only if $q(D_a^\mathcal{E}(S)) \geq p(S)$ for each $S \subseteq V$. For the forward direction, suppose each cut has capacity at least 1. Consider the cut $(\{s\} \cup S \cup D_a^\mathcal{E}(S), (V \setminus S) \cup (C \setminus D_a^\mathcal{E}(S)) \cup \{t\})$. Because $D_a^\mathcal{E}(S)$ is precisely the set of neighbors of S , there are no edges of ∞ capacity in this cut. The only edges part of the cut are those going from s to nodes in $V \setminus S$ and those going from nodes in $D_a^\mathcal{E}(S)$ to t . Thus, the capacity of this cut is $p(V \setminus S) + q(D_a^\mathcal{E}(S))$, which should be at least 1. Using $p(S) = 1 - p(V \setminus S)$, we get the desired condition.

For the reverse direction, suppose $q(D_a^\mathcal{E}(S)) \geq p(S)$ for each $S \subseteq V$. Consider any cut $(\{s\} \cup A, B \cup \{t\})$ of the network. If it contains any edge of ∞ capacity, then its capacity is clearly at least 1. If it does not, then letting $S = A \cap V$, all neighbors of nodes in S must also be in A . Hence, $D_a^\mathcal{E}(S) \subseteq A$. Thus, the cut has capacity at least $p(V \setminus S) + q(D_a^\mathcal{E}(S)) = 1 - p(S) + q(D_a^\mathcal{E}(S)) \geq 1$, as desired.

We have thus established that $G_{p,q}^\mathcal{E}(a)$ admits a fractional perfect matching if and only if the network constructed has max flow 1, if and only if $q(D_a^\mathcal{E}(S)) \geq p(S)$ for each $S \subseteq V$.

Finally, since maximum flow value can be computed in strongly polynomial time, it follows that the existence of a fractional perfect matching in $G_{p,q}^\mathcal{E}(a)$ can also be checked in strongly polynomial time. \square

We remark that while Lemma 1 is stated in terms of (p, q) -domination graphs, it in fact applies to all vertex-weighted bipartite graphs and provides a Hall's-like condition for the existence of a fractional perfect matching in such graphs.

B PLURALITYMATCHING and Condorcet Consistency

In Section 4.1, we showed that PLURALITYMATCHING is not Condorcet consistent, but this is necessary in order for it to achieve a distortion bound of $2 + \alpha$, as no Condorcet consistent rule can achieve this bound (Proposition 3). The preference profile used in the proof of Proposition 3 is an example where PLURALITYMATCHING does not select a Condorcet winner, even when one exists.

Here, we present a simpler and more insightful example in which the same phenomenon occurs. First, we observe an important necessary condition for a candidate to be potentially returned by PLURALITYMATCHING (i.e. for its integral domination graph to have a perfect matching). This condition is easy to check, and we use it elsewhere in the paper to rule out various candidates from being the output of PLURALITYMATCHING.

Recall that $\text{plu}(a)$ denotes the plurality score of a , i.e., the number of voters who rank a first. Define $\text{veto}(a)$ to be the veto score of a , i.e., the number of voters who rank a last.

Lemma 6. *Given an election $\mathcal{E} = (V, C, \mathcal{E})$ and candidate $a \in C$, a can be returned by PLURALITYMATCHING (i.e. its integral domination graph $G^\mathcal{E}(a)$ has a perfect matching) only if $\text{plu}(a) \geq \text{veto}(a)$.*

Proof. Note that in the integral domination graph $G(a)$, every $i \in V$ on the left who ranks a last — there are $\text{veto}(a)$ many such voters on the left — must be matched to some $j \in V$ on the right who ranks a first — there are $\text{plu}(a)$ many such voters on the right. Hence, for a perfect matching to exist, we must have $\text{veto}(a) \leq \text{plu}(a)$, as desired. \square

It is easy to see that a Condorcet winner may not necessarily satisfy this. In fact, it may have zero plurality score and non-zero veto score, as the example below shows.

Example 1. Consider an election \mathcal{E} with a set of 7 voters $V = \{1, \dots, 7\}$, a set of 4 candidates $C = \{a, b, c, d\}$, and the following preference profile σ .

$$\begin{aligned} \sigma_1 = \sigma_2 & : b \succ a \succ c \succ d \\ \sigma_3 = \sigma_4 & : c \succ a \succ b \succ d \\ \sigma_5 = \sigma_6 & : d \succ a \succ b \succ c \\ \sigma_7 & : b \succ c \succ d \succ a \end{aligned}$$

It is easy to check that a is the only Condorcet winner (in fact, for every other candidate, there is a *strict* majority of voters preferring a over that candidate). However, $\text{plu}(a) = 0 < \text{veto}(a) = 1$. Thus, by Lemma 6, a cannot be picked by PLURALITYMATCHING.

C PLURALITYMATCHING and Tie-Breaking

Theorem 1 established that the distortion of PLURALITYMATCHING is $2 + \alpha$ for $m \geq 3$. The tightness of the bound was shown through an example in which some candidate a whose integral domination graph had a perfect matching (and thus *could* be picked by PLURALITYMATCHING) had $2 + \alpha$ approximation to social cost. However, in that example, even the integral domination graph of an optimal candidate had a perfect matching, and therefore could also be picked by PLURALITYMATCHING, leading one to hope that perhaps wisely picking among candidates whose integral domination graphs have a perfect matching may lead to lower distortion. Below, we show that the distortion remains no less than $2 + \alpha$ regardless of the tie-breaking scheme used.

Proposition 5. *For every $m \geq 3$ and $\alpha \in [0, 1]$, every refinement of PLURALITYMATCHING (i.e. every deterministic social choice rule that always returns a candidate whose integral domination graph admits a perfect matching) has distortion $2 + \alpha$.*

Proof. The upper bound follows from Theorem 1. For the lower bound, we construct a family of elections with three candidates and an increasing number of voters such that (i) in each election, there is only a single candidate that could be chosen by PLURALITYMATCHING, and (ii) the approximation ratio of this candidate converges to $2 + \alpha$ as the number of voters grows.

For $k \geq 1$, we construct an election with $2k + 2$ voters ($V = \{1, \dots, 2k + 2\}$) and three candidates ($C = \{a, b, c\}$). The preference profile is as follows: for each $1 \leq i \leq k$, $\sigma_i = a \succ b \succ c$, for each $k + 1 \leq i \leq 2k$, $\sigma_i = c \succ b \succ a$, $\sigma_{2k+1} = b \succ a \succ c$, and $\sigma_{2k+2} = b \succ c \succ a$.

First, note that both a and c have veto score $k + 1$ but plurality score only k . Therefore, by Lemma 6, neither can be chosen by PLURALITYMATCHING. Therefore, PLURALITYMATCHING must pick candidate b . However, consider the α -decisive metric consistent with the preference profile given by the undirected graph in Figure 8; as usual, all pairwise distances are induced by shortest paths.

Note that $\text{SC}(b) = k \cdot (2 + \alpha)$, while $\text{SC}(c) = k + 2 \cdot (2 + \alpha)$. Hence, choosing b yields approximation ratio

$$\begin{aligned} \frac{k(2 + \alpha)}{k + 2(2 + \alpha)} &= \frac{k(2 + \alpha) + 2(2 + \alpha)(2 + \alpha) - 2(2 + \alpha)(2 + \alpha)}{k + 2(2 + \alpha)} \\ &= \frac{(k + 2(2 + \alpha))(2 + \alpha) - 2(2 + \alpha)^2}{k + 2(2 + \alpha)} \\ &= 2 + \alpha - \frac{2(2 + \alpha)^2}{k + 2(2 + \alpha)}. \end{aligned}$$

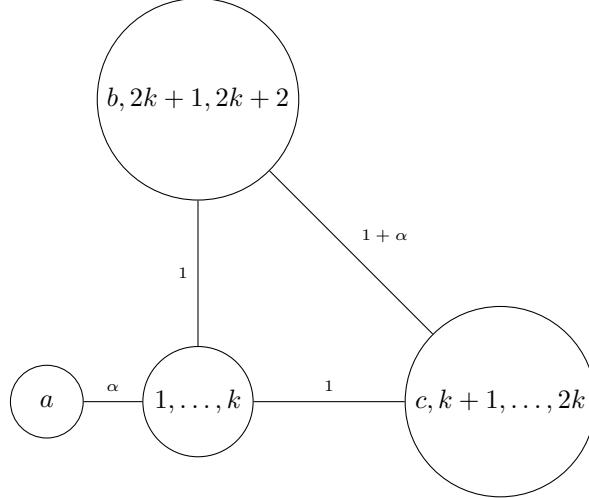


Figure 8: The metric space construction used in the proof of Proposition 5

As k grows, this approaches $2 + \alpha$. Hence, the distortion of f is at least $2 + \alpha$ when $m = 3$. For $m > 3$, we can add remaining candidates sufficiently far away as before, which does not affect the construction. \square

D PLURALITYMATCHING and the Fairness Ratio

Goel et al. [31] observe that given an underlying metric and a chosen candidate, costs may vary widely among the voters. In order to quantify how “fair” a candidate is, they look at a variation of the social cost, defined, for any given $k \in \{1, \dots, n\}$, as the sum of the k largest costs to voters. In particular, given a metric d , candidate $c \in C$, and $k \in \{1, \dots, n\}$, they define

$$\phi_k(c, d) = \max_{S \subseteq V: |S|=k} \sum_{v \in S} d(v, c).$$

Then, they define the *fairness ratio* of a social choice rule f , denoted $\text{fairness}(f)$, as its worst-case approximation ratio to the expected total of the k largest voter costs, for any k . That is,

$$\text{fairness}(f) = \sup_{\sigma} \sup_{d: d \triangleright \sigma} \max_{1 \leq k \leq n} \frac{\mathbb{E}[\phi_k(f(\sigma), d)]}{\min_{c \in C} \phi_k(c, d)}.$$

Since fixing $k = n$ yields the distortion definition, it follows that $\text{distortion}(f) \leq \text{fairness}(f)$ for any social choice rule f . Interestingly, Goel et al. [32] show that $\text{fairness}(f) \leq \text{distortion}(f) + 2$ for any social choice rule f . Hence, the fairness ratio is within an additive factor of 2 from the distortion.

Given that PLURALITYMATCHING has distortion at most $2 + \alpha$ for $m \geq 3$ (Theorem 1), it immediately follows that its fairness ratio is at most $4 + \alpha$. However, a simple modification of the proof of Theorem 1 shows that its fairness ratio is in fact $2 + \alpha$ as well, which is a stronger result than its distortion being $2 + \alpha$.

Proposition 6. *For every $m \geq 2$ and $\alpha \in [0, 1]$, PLURALITYMATCHING has fairness ratio $2 + \alpha$ for α -decisive metric spaces.*

Proof. Since $\text{fairness}(\text{PLURALITYMATCHING}) \geq \text{distortion}(\text{PLURALITYMATCHING})$, a lower bound of $2 + \alpha$ follows directly from Theorem 1. For the upper bound, the proof is nearly identical to the proof of $2 + \alpha$ upper bound on the distortion in Theorem 1. The trick is to observe that when we start from $\phi_k(a, d)$ for the candidate a picked by PLURALITYMATCHING, it consists of the k largest costs of a to some k voters. As we manipulate this quantity to express it in terms of the cost of some other candidate b , we get an expression involving the cost of b to *some* k voters, which is at most $\phi_k(b, d)$. Crucially, we use the fact that there are precisely k voters matched to any k voters in a perfect matching in the integral domination graph. We reproduce the full proof below for completeness.

Let (\mathcal{M}, d) be an α -decisive metric space, and let σ be the induced preference profile. Let a be the candidate returned by PLURALITYMATCHING, and let b be any other candidate. Fix some $k \in \{1, \dots, n\}$. We want to show that $\phi_k(a, d) \leq (2 + \alpha) \cdot \phi_k(b, d)$. Let $S_a, S_b \subseteq V$ be sets of voters of size k that have the k largest distances to candidates a and b , respectively. Due to Corollary 1, the integral domination graph $G^\mathcal{E}(a)$ admits a perfect matching $M : V \rightarrow V$ which satisfies $a \succ_i \text{top}(M(i))$ for each $i \in V$. Thus,

$$\begin{aligned}
\phi_k(a, d) &= \sum_{i \in S_a} d(i, a) \\
&\leq \sum_{i \in S_a} d(i, \text{top}(M(i))) && (\because a \succ_i \text{top}(M(i)), \forall i \in V) \\
&\leq \sum_{i \in S_a} (d(i, b) + d(b, \text{top}(M(i)))) && (\because \text{triangle inequality}) \\
&= \sum_{i \in S_a} d(i, b) + \sum_{i \in S_a} d(b, \text{top}(M(i))) \\
&= \sum_{i \in S_a} d(i, b) + \sum_{i \in M(S_a)} d(b, \text{top}(i)) && (M(S_a) = \{M(i) : i \in S_a\}) \\
&= \sum_{i \in S_a} d(i, b) + \sum_{i \in M(S_a) : \text{top}(i) \neq b} d(b, \text{top}(i)) && (\because d(b, b) = 0) \\
&\leq \sum_{i \in S_a} d(i, b) + \sum_{i \in M(S_a) : \text{top}(i) \neq b} (d(b, i) + d(i, \text{top}(i))) && (\because \text{triangle inequality}) \\
&\leq \sum_{i \in S_a} d(i, b) + \sum_{i \in M(S_a) : \text{top}(i) \neq b} (d(b, i) + \alpha \cdot d(i, b)) && (\because \alpha\text{-decisiveness}) \\
&\leq \sum_{i \in S_a} d(i, b) + \sum_{i \in M(S_a)} (d(b, i) + \alpha \cdot d(i, b)) \\
&\leq (2 + \alpha) \phi_k(b, d),
\end{aligned}$$

where the last inequality follows because $\sum_{i \in S} d(i, b) \leq \phi_k(b, d)$ for any S with $|S| = k$, and M being a perfect matching, we have $|S_a| = |M(S_a)| = k$.

Because the metric space, the election, k , and the choice of b were arbitrary, this establishes that PLURALITYMATCHING has fairness ratio $2 + \alpha$. \square

References

- [1] R. Abebe, R. Cole, V. Gkatzelis, and J. D. Hartline. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2096–2113, 2020.

- [2] B. Abramowitz and E. Anshelevich. Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 894–901, 2018.
- [3] B. Abramowitz, E. Anshelevich, and W. Zhu. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*, pages 3–16, 2019.
- [4] N. Alon, F. Fischer, A. D. Procaccia, and M. Tennenholtz. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 101–110, 2011.
- [5] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, and A. A. Voudouris. Peeking behind the ordinal curtain: improving distortion via cardinal queries. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [6] E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58:797–827, 2017.
- [7] E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 390–396, 2016.
- [8] E. Anshelevich and S. Sekar. Truthful mechanisms for matching and clustering in an ordinal world. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pages 265–278, 2016.
- [9] E. Anshelevich and W. Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. In *Proceedings of the 14th Conference on Web and Internet Economics (WINE)*, pages 3–20, 2018.
- [10] E. Anshelevich and W. Zhu. Tradeoffs between information and ordinal approximation for bipartite matching. *Theory of Computer Systems*, 63(7):1499–1530, 2019.
- [11] E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, 2015.
- [12] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [13] A. Borodin, O. Lev, N. Shah, and T. Strangway. Primarily about primaries. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1804–1811, 2019.
- [14] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [15] F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [16] R. Bredereck, J. Chen, and G. J. Woeginger. A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4):989–998, 2013.

- [17] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.
- [18] C. L. Chang. Some results on approximate 1-median selection in metric spaces. *Theoretical Computer Science*, 426:1–12, 2012.
- [19] C. L. Chang. A lower bound for metric 1-median selection. *Journal of Computer and System Sciences*, 84:44–51, 2017.
- [20] C. L. Chang. Metric 1-median selection: Query complexity vs. approximation ratio. *ACM Transactions on Computation Theory (TOCT)*, 9(4):1–23, 2017.
- [21] M. Charikar and S. Guha. Improved combinatorial algorithms for the facility location and k-median problems. In *Proceedings of the 40th Symposium on Foundations of Computer Science (FOCS)*, pages 378–388, 1999.
- [22] X. Chen, M. Li, and C. Wang. Favorite-candidate voting for eliminating the least popular candidate in metric spaces. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [23] Y. Cheng, S. Dughmi, and D. Kempe. Of the people: voting is more effective with representative candidates. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2017.
- [24] Y. Cheng, S. Dughmi, and D. Kempe. On the distortion of voting with multiple representative candidates. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 973–980, 2018.
- [25] V. Cohen-Addad, M. Pilipczuk, and M. Pilipczuk. A polynomial-time approximation scheme for facility location on planar graphs. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*, pages 560–581, 2019.
- [26] J. M. Enelow and M. J. Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.
- [27] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.
- [28] B. Fain, A. Goel, K. Munagala, and N. Prabh. Random dictators with a random referee: constant sample complexity mechanisms for social choice. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1893–1900, 2019.
- [29] A. Filos-Ratsikas, E. Micha, and A. A. Voudouris. The distortion of distributed voting. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, pages 312–325, 2019.
- [30] J. S. Gans and M. Smart. Majority voting with single-crossing preferences. *Journal of Public Economics*, 59(2):219–237, 1996.
- [31] A. Goel, A. K. Krishnaswamy, and K. Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 287–304, 2017.

- [32] A. Goel, R. Hulett, and A. K. Krishnaswamy. Relating metric distortion and fairness of social choice rules. arXiv:1810.01092, 2018.
- [33] S. Gross, E. Anshelevich, and L. Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2017.
- [34] P. Indyk. Sublinear time algorithms for metric space problems. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC)*, pages 428–434, 1999.
- [35] P. Indyk. *High-dimensional computational geometry*. PhD thesis, Citeseer, 2000.
- [36] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM*, 50(6): 795–824, 2003.
- [37] Z. Jiang, K. Munagala, and K. Wang. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 463–472, 2020.
- [38] D. Kempe. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2087–2094, 2020.
- [39] D. Kempe. An analysis framework for metric voting based on lp duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [40] D. Mandal, A. D. Procaccia, N. Shah, and D. Woodruff. Efficient and thrifty voting by any means necessary. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 7178–7189, 2019.
- [41] D. Mandal, N. Shah, and D. P. Woodruff. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 795–813, 2020.
- [42] V. Marianov, M. Ríos, and M. J. Icaza. Facility location for market capture when users rank facilities by shorter travel and waiting times. *European Journal of Operational Research*, 191(1):32–44, 2008.
- [43] H. Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.
- [44] K. Munagala and K. Wang. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 245–262, 2019.
- [45] G. Pierczynski and P. Skowron. Approval-based elections and distortion of voting rules. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 543–549, 2019.
- [46] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.
- [47] P. Skowron and E. Elkind. Social choice under metric preferences: scoring rules and STV. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 706–712, 2017.

- [48] X. Sui, A. Francois-Nienaber, and C. Boutilier. Multi-dimensional single-peaked consistency and its approximations. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 375–382, 2013.