

Resolving the Optimal Metric Distortion Conjecture

Vasilis Gkatzelis
Computer Science Department
Drexel University
Philadelphia, USA
gkatz@drexel.edu

Daniel Halpern
Computer Science Department
University of Toronto
Toronto, Canada
daniel.halpern@mail.utoronto.ca

Nisarg Shah
Computer Science Department
University of Toronto
Toronto, Canada
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 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.

Keywords-distortion; metric; voting;

I. 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 [1], has recently become popular in the computer science literature [2, 3, 4]. An alternative, possibly more concrete, interpretation of these distances can also be derived through the well-studied problem of facility location [5, 6, 7]. 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 [8, Part I] or [9]), 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. [10], 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. [10] showed that no deterministic

For a full version of the paper, see arxiv.org/abs/2004.07447

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 [3] 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 [4], 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 [11] 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 [4]. 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. [10] that ranked pairs might have distortion 3 (this conjecture was also earlier disproved by Goel et al. [12]). 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.

A. Our Results & Techniques

The technical core of our paper is in Section III, 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 VI, 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 IV, which is that PLURALITYMATCHING achieves the optimal distortion bound of 3.

Our approach is inspired by the recent work of Muna-

gala and Wang [4], 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 [4], 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 IV-A 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 $3 - 2/m$ [13], which is better than what any deterministic algorithm can achieve. In Section V, 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 [14]. 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 II presents the model, Section III presents the Ranking-Matching Lemma, and Sections IV and V study the distortion of deterministic and randomized algorithms, respectively. Section VI contains a discussion of our results and open directions, and the appendix of the full version contains deferred proofs and additional observations.

II. 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 $SC(c, d) = \sum_{i \in V} d(i, c)$; we simply write $SC(c)$ 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 [14]. 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. 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.

III. 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 $G_{p,q}^{\mathcal{E}}(a) = (V, C, E_a, p, q)$, where $(i, c) \in E_a$ if and only if $a \succsim_i c$. 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., $G_{p,q}^{\mathcal{E}}(a)$, *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 $D_a^{\mathcal{E}}(S) = \{c \in C : a \succsim_i c, \exists i \in S\}$ 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 $G_{p,q}^{\mathcal{E}}(a)$ (for any p and q). A simple generalization of Hall's condition shows that the existence of a fractional perfect matching in $G_{p,q}^{\mathcal{E}}(a)$ 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 the appendix of the full version.




Figure 1: Domination graph examples

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 IV 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 IV 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 III-A and VI, 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 . ■

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

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

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:

$$\begin{aligned} p'(i) &= p(i)/p(\overline{X}), \forall i \in \overline{X} \\ q'(c) &= q(c)/q(\overline{D}), \forall c \in \overline{D}. \end{aligned}$$

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

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

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

$$\begin{aligned} \frac{q(D_b^\mathcal{E}(X_b \cap \overline{X}) \cap \overline{D})}{q(\overline{D})} &\geq \frac{p(X_b \cap \overline{X})}{p(\overline{X})} \\ \Rightarrow q(D_b^\mathcal{E}(X_b \cap \overline{X}) \cap \overline{D}) &\geq p(X_b \cap \overline{X}) \cdot \frac{q(\overline{D})}{p(\overline{X})} \geq p(X_b \cap \overline{X}), \end{aligned}$$

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

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

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. ■

A. 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 VI. 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.

IV. 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 [4, 10], matching a lower bound of 3 established by Anshelevich et al. [10]. 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,

$$\begin{aligned}
\text{SC}(a) &= \sum_{i \in V} d(i, a) \\
&\leq \sum_{i \in V} d(i, \text{top}(M(i))) \\
&\leq \sum_{i \in V} (d(i, b) + d(b, \text{top}(M(i)))) \\
&= \text{SC}(b) + \sum_{i \in V} d(b, \text{top}(M(i))) \\
&= \text{SC}(b) + \sum_{i \in V} d(b, \text{top}(i)) \\
&= \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} d(b, \text{top}(i)) \\
&\leq \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} (d(b, i) + d(i, \text{top}(i))) \\
&\leq \text{SC}(b) + \sum_{i \in V: \text{top}(i) \neq b} (d(b, i) + \alpha \cdot d(i, b)) \\
&\leq \text{SC}(b) + \sum_{i \in V} (d(b, i) + \alpha \cdot d(i, b)) \\
&= (2 + \alpha) \cdot \text{SC}(b),
\end{aligned}$$

where the second transition holds because $a \succ_i \text{top}(M(i))$ for all $i \in V$, the third and the seventh transition hold follow from the triangle inequality, the fifth transition holds because M is a perfect matching, the sixth transition holds because $d(b, b) = 0$, and the penultimate transition holds due to α -decisiveness.

Because the metric space, the election, and the choice of b were arbitrary, this establishes that PLURALITYMATCHING has distortion at most $2 + \alpha$. A proof of tightness of this bound is given in the full version. ■

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 $\alpha \in [0, 1]$ [14].

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 rules is $1 + 2\alpha$ [14], 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. The proof is deferred to the full version.

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 .*

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.

A. Connecting PLURALITYMATCHING to The Matching Uncovered Set

As mentioned in Section I, we devised the notion of integral domination graphs by directly building on the work of Munagala and Wang [4]. 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 [4] 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 [4]. Second, as we show in the appendix of the full version, our analysis seamlessly extends to produce an upper bound on the so-called *fairness ratio* [12], 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 [4]. 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. Proofs of these are deferred to the full version.

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.

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.

It is easy to argue that every Condorcet winner a is in

the matching uncovered set.¹ 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.

Anshelevich et al. [10] 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 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,² 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.

V. 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. Proofs in this section are deferred to the full version.

¹This 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.

²In the appendix of the full version, 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.

Randomized rules were first considered in this framework by Anshelevich and Postl [14]. 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 $2 + \alpha$.³ 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 [13] 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 [12, 15].

What about the case of $\alpha < 1$? Anshelevich and Postl [14] 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 matches the distortion of $3 - 2/m$ achieved by Kempe [13] for $\alpha = 1$, but improves when $\alpha < 1$. `SMARTDICTATORSHIP`, like `RANDOMDICTATORSHIP`, α -`GENERALIZED PROPORTIONAL TO SQUARES`, and the rule proposed by Kempe [13], 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`

³Specifically, 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.

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 [14] 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 $\alpha = 0$.

A. 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 [13], the two rules are different when $m \geq 3$.

Using helpful lemmas proved by Anshelevich and Postl [14], we derive 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.

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.⁴ 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 $m > 3$

⁴As mentioned earlier, this was also established by Kempe [13].

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.

B. 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 $\alpha = 1$ [12, 15].

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.*

In terms of general lower bounds, when the randomized rule is given access to the full preference profile, Anshelevich and Postl [14] 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 smaller than or equal to m .*

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.*

VI. 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 III, 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 [16], 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 a committee in the core always exists is also open (without any relaxations required) when voters have binary preferences over candidates [17]. 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 [18], whether it provides any core-like guarantees is an important open question.

ACKNOWLEDGMENT

Vasilis Gkatzelis was supported by NSF grant CCF-1755955. Nisarg Shah was partially supported by an NSERC Discovery Grant.

REFERENCES

- [1] J. M. Enelow and M. J. Hinich, *The spatial theory of voting: An introduction*. CUP Archive, 1984.
- [2] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron, “Approximating optimal social choice under metric preferences,” *Artificial Intelligence*, vol. 264, pp. 27–51, 2018.
- [3] 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)*, 2017, pp. 706–712.
- [4] 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)*, 2019, pp. 245–262.
- [5] 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)*, 1999, pp. 378–388.
- [6] 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*, vol. 50, no. 6, pp. 795–824, 2003.
- [7] 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)*, 2019, pp. 560–581.
- [8] F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, Eds., *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [9] 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*, vol. 191, no. 1, pp. 32–44, 2008.
- [10] 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)*, 2015, pp. 777–783.
- [11] D. Kempe, “An analysis framework for metric voting based on lp duality,” in *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [12] 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)*, 2017, pp. 287–304.
- [13] D. Kempe, “Communication, distortion, and randomness in metric voting,” in *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 2087–2094.
- [14] E. Anshelevich and J. Postl, “Randomized social choice functions under metric preferences,” *Journal of Artificial Intelligence Research*, vol. 58, pp. 797–827, 2017.
- [15] 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)*, 2017, pp. 544–550.
- [16] Z. Jiang, K. Munagala, and K. Wang, “Approximately stable committee selection,” in *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, 2020, pp. 463–472.
- [17] 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)*, 2018, pp. 575–592.
- [18] A. Goel, R. Hulett, and A. K. Krishnaswamy, “Relating metric distortion and fairness of social choice rules,” arXiv:1810.01092, 2018.