

The Metric Distortion of Multiwinner Voting

Ioannis Caragiannis,¹ Nisarg Shah,² Alexandros A. Voudouris³

¹Department of Computer Science, Aarhus University

²Department of Computer Science, University of Toronto

³School of Computer Science and Electronic Engineering, University of Essex
iannis@cs.au.dk, nisarg@cs.toronto.edu, alexandros.voudouris@essex.ac.uk

Abstract

We extend the recently introduced framework of metric distortion to multiwinner voting. In this framework, n agents and m alternatives are located in an underlying metric space. The exact distances between agents and alternatives are unknown. Instead, each agent provides a ranking of the alternatives, ordered from the closest to the farthest. Typically, the goal is to select a single alternative that approximately minimizes the total distance from the agents, and the worst-case approximation ratio is termed distortion. In the case of multiwinner voting, the goal is to select a committee of k alternatives that (approximately) minimizes the total cost to all agents. We consider the scenario where the cost of an agent for a committee is her distance from the q -th closest alternative in the committee. We reveal a surprising trichotomy on the distortion of multiwinner voting rules in terms of k and q : The distortion is unbounded when $q \leq k/3$, asymptotically linear in the number of agents when $k/3 < q \leq k/2$, and constant when $q > k/2$.

Introduction

The most canonical problem in voting theory is to aggregate ranked preferences of n individual agents over a set of m alternatives to reach a collective decision. Examples of such decisions include selecting a single alternative (single-winner voting), selecting k out of m alternatives for a fixed $k > 1$ (multiwinner voting), and selecting a set of costly alternatives subject to a budget constraint (participatory budgeting). In centuries of research on voting, perhaps the most prominent approach to designing voting rules has been the axiomatic approach, in which one fixes several qualitative axioms and seeks voting rules satisfying them. Unfortunately, this approach has often led to impossibility results, such as Arrow’s impossibility theorem (Arrow 1951).

To circumvent this, Procaccia and Rosenschein (2006) proposed the (utilitarian) *distortion* framework for analyzing, comparing, and designing single-winner voting rules. Under this framework, the ordinal preferences expressed by agents are viewed as proxies for their underlying cardinal utilities, and the goal of a voting rule is to optimize the worst-case approximation ratio (distortion) of the social

welfare (the total utility of the agents). Under minimal assumptions, this framework offers a *quantitative* comparison of voting rules. It has been used successfully to analyze the distortion of well-known methods (Caragiannis and Procaccia 2011) and to identify voting rules with optimal distortion (Boutilier et al. 2015). The framework has also been extended to multiwinner voting (Caragiannis et al. 2017) and participatory budgeting (Benadè et al. 2017), under the assumption that the utility of an agent for a set of alternatives is the maximum and the sum of her utilities for the alternatives in the set, respectively.

Anshelevich, Bhardwaj, and Postl (2015) built on this idea to propose the *metric distortion* framework, in which agents and alternatives are embedded in an underlying metric space, and the *cost* of an agent for an alternative is the distance between them. An agent still ranks the alternatives, but now in a non-decreasing order of their distance from her. Instead of maximizing the social welfare, the goal is now to minimize the social cost (the total cost of the agents). Like in the utilitarian case, scholars have analyzed the metric distortion of prominent voting rules (Skowron and Elkind 2017; Goel, Krishnaswamy, and Munagala 2017; Munagala and Wang 2019; Kempe 2020a) and identified rules with optimal metric distortion (Gkatzelis, Halpern, and Shah 2020). For a detailed overview, we refer the reader to the survey by Anshelevich et al. (2021).

The goal of our work is to extend the metric distortion framework to multiwinner voting, where the objective is to select a subset of k alternatives (committee) for a given $k > 1$. To the best of our knowledge, the only prior works to address multiwinner voting with general metric costs are those of Goel, Hulett, and Krishnaswamy (2018) and Chen, Li, and Wang (2020). Both focus on a model in which an agent’s cost for a committee is the sum of her distances to the alternatives in the committee. In a sense, this assumes that an agent cares equally about all the alternatives chosen in the committee.

However, in many applications, an agent may care only about a few alternatives in the committee, typically the ones she prefers more. For example, when parliament members are chosen in a political election, each voter may associate with just one (or a few) of the elected candidates as her representative(s). Similarly, if a city builds parks at multiple locations, each resident may only be able to access a few parks

closest to her. Motivated by these applications, we consider the case where an agent’s cost for a committee of k alternatives is her distance to the q -th closest alternative in the committee, for a given $q \leq k$. Note that for $q = 1$, this is the distance to the closest alternative in the committee, whereas for $q = k$, it is the distance to the farthest one. Our main research question is:

What is the optimal distortion for selecting a committee of k alternatives under this cost model? How does it depend on the relation between q and k ? Can the optimal distortion be achieved via efficient voting rules?

Before proceeding further, note that another reason why it may not be desirable to model an agent’s cost for a committee as the *sum* of her distances to the alternatives in the committee is that the optimal committee can suffer from the tyranny of the majority; that is, it may consist entirely of the alternatives liked by the majority and include none liked by a minority. This is somewhat reflected by the fact that repeatedly applying a single-winner voting rule is known to work well in this model (Goel, Hulett, and Krishnaswamy 2018). Due to this, prior work on multiwinner voting, both in the distortion literature (Caragiannis et al. 2017) and elsewhere (Chamberlin and Courant 1983; Monroe 1995; Procaccia, Rosenschein, and Zohar 2008; Lu and Boutilier 2011), aims to ensure that there is *at least one* alternative in the committee that every agent likes well (which corresponds to $q = 1$ in our model). An interesting byproduct of our cost model is that an agent’s submitted ranking of individual alternatives can be used to determine her ordinal preferences over committees; we use this fact to derive some of our results.

Our Contributions

Recall that n is the number of agents and m is the number of alternatives. We reveal a surprising trichotomy on the best possible distortion for multiwinner voting.

When $q \leq k/3$, the distortion is unbounded. This holds even for randomized voting rules, and even when n and m are only linear in k . When $q \in (k/3, k/2]$, the best possible distortion is $\Theta(n)$. Here, the upper bound is obtained via a novel voting rule that is deterministic and efficient, while the lower bound holds even for randomized voting rules and when m is linear in k or q . Finally, when $q > k/2$, the best possible distortion is 3 for deterministic rules, and between 2 and 3 for randomized rules. For this case, we show that the costs of agents for committees satisfy the triangle inequality. Hence, we can reduce multiwinner voting to single-winner voting by treating each committee as a separate alternative and applying any single-winner voting rule, allowing us to borrow known best possible distortion bounds from the single-winner case to the multiwinner one. This is where we use the fact that under our cost model, we can deduce an agent’s ordinal preferences over committees from her provided ranking of individual alternatives.

However, this reduction is not efficient as we need to apply the single-winner voting rule on an instance with $\binom{m}{k}$ alternatives. To that end, we also provide an efficient reduc-

tion. We show that there exists an agent such that the committee consisting of the k alternatives she prefers the most has social cost no worse than 3 times the optimal. Thus, we can apply any single-winner voting rule with distortion ρ on a reduced instance with only n committees (one committee per agent) as alternatives and obtain a multiwinner rule with distortion at most 3ρ . In particular, by applying the PLURALITYMATCHING rule of Gkatzelis, Halpern, and Shah (2020), which is known to achieve the best possible distortion of 3 in the single-winner case, we obtain distortion at most 9 in polynomial time. We also show that an efficient reduction of this type cannot be used to achieve distortion better than 5.207.

Related Work

The works of Goel, Hulett, and Krishnaswamy (2018) and Chen, Li, and Wang (2020) are the most closely related to ours as they consider metric distortion for multiwinner voting. As mentioned before, both focus on a setting where the cost of an agent for a committee is the sum of her distances to the alternatives in the committee.¹ Goel, Hulett, and Krishnaswamy (2018) show that selecting a committee by repeatedly applying any single-winner voting rule with distortion δ yields a distortion of at most δ . Since the best possible single-winner distortion is known to be 3 (Gkatzelis, Halpern, and Shah 2020), this implies that a distortion of 3 can be achieved for any k in this model. Chen, Li, and Wang (2020) directly present a voting rule achieving distortion 3 for the case of $k = m - 1$ and prove this to be best possible.

Jaworski and Skowron (2020) consider a model where voters (agents) and candidates (alternatives) have preferences over a set of *binary issues*, which induce the preferences of the voters over the candidates; specifically, voters rank candidates based on the number of issues they agree on. The elected committee uses majority voting to decide on each issue, and the cost of a voter is the number of issues for which the decision differs from her preferred outcome. This can be viewed as a metric distortion model, but with the specific Hamming distance metric. Also, the cost of a voter for a committee depends on the locations of the candidates in the committee, and not just on their distances to her. Meir, Sandmirskiy, and Tennenholtz (2021) consider a similar model where voters are also candidates, and show that sortition — picking k of the voters uniformly at random — leads to low distortion.

The distortion of randomized rules has also received significant attention. For single-winner voting, the best possible distortion under the utilitarian model is known to be $\tilde{\Theta}(\sqrt{m})$ (Boutilier et al. 2015), while that under the metric model still remains a challenging open question (Anshelevich and Postl 2017; Kempe 2020b). Caragiannis et al. (2017) provide bounds on the best possible distortion of randomized multiwinner voting rules under the utilitarian model, and our work does so under the metric model.

More broadly, there is a huge literature on multiwinner voting that focuses on desiderata other than distortion, such

¹Chen, Li, and Wang (2020) focus only on the case of $k = m - 1$, but also consider another cost model.

as proportional representation (Aziz et al. 2017; Peters and Skowron 2020), committee diversity (Bredereck et al. 2018), monotonicity or consistency axioms (Elkind et al. 2017), explainability (Peters et al. 2021), etc. We refer the interested reader to the book chapter by Faliszewski et al. (2017) for an overview.

Preliminaries

For $p \in \mathbb{N}$, define $[p] = \{1, \dots, p\}$. An instance of our problem is given by the tuple $I = (N, A, d, k, q)$, where:

- N is a set of $n \geq 2$ agents.
- A is a set of $m \geq 2$ alternatives.
- d is a pseudometric over $N \cup A$ with $d(x, y)$ denoting the *distance* between $x, y \in N \cup A$. Being a pseudometric, d satisfies, for all $x, y, z \in N \cup A$, $d(x, x) = 0$, $d(x, y) = d(y, x)$, and the *triangle inequality* $d(x, y) \leq d(x, z) + d(z, y)$. Since our framework only uses distances between agents and alternatives (and not between two agents or between two alternatives), we use the following equivalent formulation of the triangle inequality (Goel, Krishnaswamy, and Munagala 2017): $d(i, x) \leq d(i, y) + d(j, y) + d(j, x)$ for all agents $i, j \in N$ and alternatives $x, y \in A$.²
- k and q are positive integers such that $1 \leq q \leq k < m$.

Every agent $i \in N$ ranks the alternatives based on her distances from them, from smallest (most preferable) to largest (least preferable), breaking ties arbitrarily; that is, the pseudometric d induces the ordinal preferences of agent i given by a *ranking* \succ_i over the alternatives such that $x \succ_i y$ implies $d(i, x) \leq d(i, y)$. We refer to $\succ_d = (\succ_i)_{i \in N}$ as the *preference profile*.

For any $S \subseteq A$ with $|S| \geq q$, we denote by $\text{top}_{i,q}(S)$ the set of q most preferred alternatives of agent i in S . A *committee* $C \subseteq A$ is a set of alternatives of size exactly equal to k . We define the q -cost of agent i for C , denoted $c_{i,q}(C|d)$, to be the distance of i from her q -th closest alternative in C : $c_{i,q}(C|d) = \max_{x \in \text{top}_{i,q}(C)} d(i, x)$. The q -social cost of committee C , denoted $\text{SC}_q(C|d)$, is then defined as the total q -cost of the agents for C : $\text{SC}_q(C|d) = \sum_{i \in N} c_{i,q}(C|d)$.

A (randomized) *multiwinner voting rule* f takes as input a preference profile \succ_d and outputs a distribution over committees; we say that f is deterministic if the distribution it returns has singleton support. Given $k \in \mathbb{N}$ and $q \in [k]$, the (k, q) -distortion of f is the worst-case ratio, over all possible instances with these parameters, between the (expected) q -social cost of the committee chosen by f and the minimum possible q -social cost of any committee, i.e.,

$$\sup_{I=(N,A,d,k,q):|A|=m>k} \frac{\mathbb{E}[\text{SC}_q(f(\succ_d)|d)]}{\min_{C \subseteq A: |C|=k} \text{SC}_q(C|d)}.$$

Our goal is to design multiwinner rules with as low (k, q) -distortion as possible. Due to lack of space, some proofs are omitted.

²When proving lower bounds, we will often design a worst-case pseudometric d by embedding agents and alternatives in the 1D Euclidean space and taking the Euclidean distance between them.

agent	ranking
v_1	$X_1 \succ X_2 \succ Y_1 \succ Y_2 \succ Y_3$
v_2	$X_2 \succ X_1 \succ Y_1 \succ Y_2 \succ Y_3$
u_1	$Y_1 \succ Y_2 \succ Y_3 \succ X_2 \succ X_1$
u_2	$Y_2 \succ Y_1 \succ Y_3 \succ X_2 \succ X_1$
u_3	$Y_3 \succ Y_2 \succ Y_1 \succ X_2 \succ X_1$

Table 1: An example of the preference profile used in the proof of Theorem 1, when $k = 8$ and $q = 2$.

To simplify notation, we drop q and d whenever they are clear from the context. In particular, we will use $c_i(C)$ instead of $c_{i,q}(C|d)$, $\text{SC}(C)$ instead of $\text{SC}_q(C|d)$, and $\text{top}_i(S)$ instead of $\text{top}_{i,q}(S)$. We also denote by $T_i = \text{top}_i(A)$ the set of q alternatives ranked highest by agent i .

Unbounded Distortion With $q \leq k/3$

We begin with a strong impossibility result for the case where $q \leq k/3$. In particular, we show that every multiwinner voting rule has unbounded distortion, even if it is allowed to use randomization.

Theorem 1. *For $q \leq k/3$, the (k, q) -distortion of every (even randomized) multiwinner voting rule is unbounded.*

Proof. Fix k and q such that $q \leq k/3$, and a multiwinner voting rule f . Let $L = \lfloor k/q \rfloor + 1 \geq 4$. We consider an instance with $n = L$ agents, partitioned into two sets: $V = \{v_1, \dots, v_{\lfloor L/2 \rfloor}\}$ and $U = \{u_1, \dots, u_{\lfloor L/2 \rfloor}\}$. There are $m = Lq$ alternatives, partitioned into L sets $X_1, \dots, X_{\lfloor L/2 \rfloor}, Y_1, \dots, Y_{\lfloor L/2 \rfloor}$ of size q each.³ Let $X = \bigcup_{\ell=1}^{\lfloor L/2 \rfloor} X_\ell$ and $Y = \bigcup_{\ell=1}^{\lfloor L/2 \rfloor} Y_\ell$.


Consider any preference profile such that:

- Every agent in V ranks the alternatives in X higher than those in Y .
- Every agent in U ranks the alternatives in Y higher than those in X .
- For $\ell \in [\lfloor L/2 \rfloor]$, every agent ranks the alternatives in X_ℓ as well as those in Y_ℓ arbitrarily.
- For $\ell \in [\lfloor L/2 \rfloor]$, agent $v_\ell \in V$ ranks the alternatives in X_i higher than those in X_j whenever $|\ell - i| < |\ell - j|$, for $i, j \in [\lfloor L/2 \rfloor]$; agent v_ℓ ranks the sets of Y in the order $Y_1, \dots, Y_{\lfloor L/2 \rfloor}$ from highest to lowest.
- For $\ell \in [\lfloor L/2 \rfloor]$, agent $u_\ell \in U$ ranks the alternatives in Y_i higher than those in Y_j whenever $|\ell - i| < |\ell - j|$, for $i, j \in [\lfloor L/2 \rfloor]$; agent u_ℓ ranks the sets of X in the order $X_{\lfloor L/2 \rfloor}, \dots, X_1$ from highest to lowest.


Table 1 presents an example preference profile for $k = 8$ and $q = 2$; Figure 1 depicts the possible underlying metric spaces used in the two cases below for this example instance.

Because $m = Lq = (\lfloor k/q \rfloor + 1) \cdot q > k$, not all alternatives can be included in the committee. We switch between the following two cases.

³We use an instance with $n \leq m$. Note, however, that the lower bound continues to hold even if one assumes $n \geq m$. We can simply create t copies of each agent, for an appropriately large $t \in \mathbb{N}$.



(a) The metric space in Case 1.



(b) The metric space in Case 2.

Figure 1: The two metrics used in the proof of Theorem 1 when $k = 8$ and $q = 2$. Both are consistent with the ordinal preferences presented in Table 1.

Case 1: Some alternative in Y is not included in the committee with a positive probability.

Suppose that an alternative in Y_{ℓ^*} is not included in the committee with a positive probability, for some $\ell^* \in \lceil \lceil L/2 \rceil \rceil$. Consider the following one-dimensional Euclidean metric, which is consistent with the ordinal preferences of the agents defined above:

- All agents in V and all alternatives in X are located at 0.
- For every $\ell \in \lceil \lceil L/2 \rceil \rceil$, agent $u_\ell \in U$ and the alternatives in Y_ℓ are located at $\lceil L/2 \rceil + \ell$.

Since some alternative in Y_{ℓ^*} is not included in the committee with a positive probability, the expected q -cost of agent u_{ℓ^*} is positive, and thus the expected q -social cost under f is also positive. However, it is possible to achieve q -social cost 0 by including in the committee all the alternatives of Y and an arbitrary subset of $k - \lceil L/2 \rceil q \geq q$ alternatives from X ; here, the inequality follows because $q \leq k/3$. So, the distortion of f is unbounded in this case.

Case 2: Every alternative in Y is included in the committee with probability 1.

Consider the following metric, which is consistent with the ordinal preferences of the agents defined above:

- For $\ell \in \lceil \lceil L/2 \rceil \rceil$, agent v_ℓ and the alternative in X_ℓ are located at $-L + \ell$.
- All agents of U and all alternatives in Y are located at 0.

Since f chooses a committee that includes all of Y with probability 1, it excludes at least one alternative of X with probability 1. Whenever an alternative in X_ℓ is excluded, note that agent v_ℓ has q -cost at least 1. Hence, the expected q -social cost is at least 1. However, it is possible to achieve q -social cost 0 by choosing the committee containing all the alternatives of X and an arbitrary subset of $k - \lceil L/2 \rceil q \geq q$ alternatives from Y . So, the distortion of f is unbounded in this case as well. \square

Algorithm 1: POLAROPPOSITES

- 1 Choose an arbitrary agent $i \in N$;
 - 2 Choose an agent $j \in \arg \max_{\ell \in N \setminus \{i\}} c_i(T_\ell)$;
 - 3 Output an arbitrary committee $W \supseteq T_i \cup T_j$;
-

Linear Distortion With $k/3 < q \leq k/2$

We now turn our attention to the case of $q \in (k/3, k/2]$. In this case, the best possible distortion turns out to be bounded, but linear in the number of agents, which could be very large.

We propose a novel deterministic multiwinner voting rule, called POLAROPPOSITES (see Algorithm 1), which runs in polynomial time and achieves a distortion of $O(n)$. Recall that, for a fixed value of q , $\text{top}_i(S)$ is the set of the q most favorite alternatives of agent i in S , and $T_i = \text{top}_i(A)$. POLAROPPOSITES is relatively straightforward: we choose an agent i arbitrarily, and another agent j whose T_j has the highest cost⁴ for agent i ; then we output any committee that includes $T_i \cup T_j$. However, the analysis of the distortion upper bound of our rule is intricate.

Before we proceed with bounding the distortion of our rule, we present a structural lemma, which holds for all possible values of k and q , and will be extremely useful in the proof of the bound.

Lemma 1. *Consider any instance $I = (N, A, d, k, q)$ and let $O \in \arg \min_{C: |C|=k} \text{SC}(C)$ be an optimal committee. There exists a subset of agents S with $|S| \leq \lfloor k/q \rfloor$ such that for every agent $i \in N$ there exists an agent $j \in S$ with $\text{top}_i(O) \cap \text{top}_j(O) \neq \emptyset$ and $c_j(O) \leq c_i(O)$. In addition, for every agent $i \in N$ and committee $C \supseteq \bigcup_{j \in S} \text{top}_j(A)$, we have $c_i(C) \leq 3 \cdot c_i(O)$.*

We are now ready to bound the distortion of the POLAROPPOSITES rule.

Theorem 2. *For any $q \in (k/3, k/2]$, the (k, q) -distortion of POLAROPPOSITES is $O(n)$.*

Proof. Let $I = (N, A, d, k, q)$ be an instance with $q \in (k/3, k/2]$. Let i and j be the agents chosen by POLAROPPOSITES on I , $W \supseteq T_i \cup T_j$ be the committee returned by it, and $O \in \arg \min_{C: |C|=k} \text{SC}(C)$ be an optimal committee for I . We will show that $c_\ell(W) \leq c_\ell(O) + 4 \cdot \text{SC}(O)$ for every agent $\ell \in N$. Then, by summing over all agents, we will obtain that $\text{SC}(W) \leq (4n+1) \cdot \text{SC}(O)$, thus implying an upper bound of $4n+1$ on the distortion of POLAROPPOSITES. We distinguish between the following two cases.

Case 1: $\text{top}_i(O) \cap \text{top}_j(O) \neq \emptyset$. Let $x \in \text{top}_i(O) \cap \text{top}_j(O)$ be an alternative that both i and j consider to be among their top q alternatives in the optimal committee O . For any agent $\ell \in N$, let

- $y_\ell = \arg \max_{z \in T_\ell} d(\ell, z)$ be the q -th overall favorite alternative of ℓ , and

⁴Here, we use the fact that in our model, we can compare two sets of alternatives in terms of their cost to agent i using only agent i 's preference ranking over individual alternatives.

- $y_{\ell_j} = \arg \max_{z \in T_j} d(\ell, z)$ be the q -th favorite alternative of ℓ among those in T_j .

Since $T_j \subseteq W$, by the definition of y_{ℓ_j} , we have $c_\ell(W) \leq c_\ell(T_j) = d(\ell, y_{\ell_j})$. Combining this with the triangle inequality, we obtain

$$\begin{aligned} c_\ell(W) &\leq d(\ell, y_{\ell_j}) \\ &\leq d(\ell, y_\ell) + d(i, y_\ell) + d(i, x) + d(j, x) + d(j, y_{\ell_j}). \end{aligned}$$

For any agent v , note that $c_v(T_v) \leq c_v(O)$, and for any agent u and alternative $y \in T_v$, we also have $d(u, y) \leq c_u(T_v)$. Combining these with the definitions of j, y_ℓ, x , and y_{ℓ_j} , we can now bound each of the terms in the expression above as follows:

- $d(\ell, y_\ell) = c_\ell(T_\ell) \leq c_\ell(O)$;
- $d(i, y_\ell) \leq c_i(T_\ell) \leq c_i(T_j)$;
- $d(i, x) \leq c_i(O)$;
- $d(j, x) \leq c_j(O)$;
- $d(j, y_{\ell_j}) \leq c_j(T_j) \leq c_j(O)$.

Substituting these, we get

$$c_\ell(W) \leq c_\ell(O) + c_i(T_j) + c_i(O) + 2c_j(O).$$

We can bound the term $c_i(T_j)$ using the definition of alternative y_{ij} , the triangle inequality, and some of the above observations, as follows:

$$\begin{aligned} c_i(T_j) &= d(i, y_{ij}) \\ &\leq d(i, x) + d(j, x) + d(j, y_{ij}) \\ &\leq c_i(O) + c_j(O) + c_j(T_j) \\ &\leq c_i(O) + 2c_j(O). \end{aligned}$$

So, combining everything, and using the fact that $\text{SC}(O) \geq c_i(O) + c_j(O)$, we have that


$$\begin{aligned} c_\ell(W) &\leq c_\ell(O) + 2c_i(O) + 4c_j(O) \\ &\leq c_\ell(O) + 4 \cdot \text{SC}(O), \end{aligned}$$

as desired.


Case 2: $\text{top}_i(O) \cap \text{top}_j(O) = \emptyset$. Consider the set S that is guaranteed to exist by Lemma 1. Since $q \in (k/3, k/2]$, we have that $\lfloor k/q \rfloor = 2$, and hence $|S| \leq 2$.

If $|S| = 1$, then every agent $\ell \in N$ is mapped to the single agent $v \in S$, and it holds that $\text{top}_\ell(O) \cap \text{top}_v(O) \neq \emptyset$. If $|S| = 2$, we claim that there is a function $g : N \rightarrow S$ such that for every agent $\ell \in N$, it holds that $\text{top}_\ell(O) \cap \text{top}_{g(\ell)}(O) \neq \emptyset$ and $S = \{g(i), g(j)\}$. Lemma 1 already guarantees a function g meeting the first condition. The only case in which the second condition cannot be met is if there exists $v \in S$ such that $\text{top}_i(O) \cap \text{top}_v(O) = \text{top}_j(O) \cap \text{top}_v(O) = \emptyset$. However, this, along with $\text{top}_i(O) \cap \text{top}_j(O) = \emptyset$ implies that $\text{top}_i(O)$, $\text{top}_j(O)$, and $\text{top}_v(O)$ are disjoint subsets of O of size q each. This is a contradiction since $|O| = k < 3q$.

Now, consider any agent $\ell \in N$, and suppose that $g(\ell) = g(j) = v \in S$; the case $g(\ell) = g(i)$ if $|S| = 2$ is similar. By the properties of S , there exist alternatives $x \in \text{top}_\ell(O) \cap \text{top}_v(O)$ and $z \in \text{top}_j(O) \cap \text{top}_v(O)$. As in case 1, let $y_{\ell_j} \in$



(a) The metric space in Case 1.



(b) The metric space in Case 2.

Figure 2: The metric spaces used in the proof Theorem 3. Each of the sets X, Y, Z has size q . Both metrics are consistent with the ordinal profile according to which the preference of all x agents in V is $X \succ Y \succ Z$, the preference of the x agents in U is $Y \succ X \succ Z$, and the preference of agent w is $Z \succ Y \succ X$. If the q alternatives of Z are not all included in the committee with positive probability (Case 1), then agent w has positive expected q -cost in the first metric space, leading to unbounded distortion as the committee that includes Z and q alternatives from $X \cup Y$ has q -social cost 0. If all alternatives of Z are included in the committee (Case 2), then the x agents of V (in case not all alternatives of X are included in the committee) or the x agents in U (in case not all alternatives of Y are included in the committee) have q -cost 1 in the second metric space, leading to a distortion of $x = \Omega(n)$ as the committee that includes all the alternatives of $X \cup Y$ and some alternatives of Z has q -social cost 1. In any case, the distortion of the voting rule is $\Omega(n)$.

T_j be the q -th favorite alternative of ℓ in T_j ; so, $c_\ell(T_j) = d(\ell, y_{\ell_j})$. By the fact that $T_j \subseteq W$, and using the triangle inequality, we obtain

$$\begin{aligned} c_\ell(W) &\leq c_\ell(T_j) = d(\ell, y_{\ell_j}) \\ &\leq d(\ell, x) + d(v, x) + d(v, z) + d(j, z) + d(j, y_{\ell_j}). \end{aligned}$$

By the definitions of x, z and y_{ℓ_j} , we have $d(\ell, x) \leq c_\ell(O)$, $d(v, x) \leq c_v(O)$, $d(v, z) \leq c_v(O)$, $d(j, z) \leq c_j(O)$, and $d(j, y_{\ell_j}) \leq c_j(T_j) \leq c_j(O)$. Combined with the fact that $\text{SC}(O) \geq c_v(O)$ and $\text{SC}(O) \geq c_j(O)$, we get

$$\begin{aligned} c_\ell(W) &\leq c_\ell(O) + 2c_v(O) + 2c_j(O) \\ &\leq c_\ell(O) + 4 \cdot \text{SC}(O), \end{aligned}$$

as desired. \square

We conclude this section by showing a matching lower bound of $\Omega(n)$ on the distortion of any (even randomized) multiwinner voting rule. Hence, when $q \in (k/3, k/2]$, POLAROPPOSITES is the asymptotically best possible rule in terms of distortion, even among randomized rules. See Figure 2 and its caption for a sketch of the proof.

Theorem 3. For $q \in (k/3, k/2]$, the (k, q) -distortion of every (even randomized) multiwinner voting rule is $\Omega(n)$.

Constant Distortion With $q > k/2$

We now turn our attention to the case $q > k/2$, where it is possible to achieve constant distortion. A crucial observation, which we exploit in this section, is that the q -costs of the agents form a new metric over the agents and all possible committees of alternatives.

Lemma 2. *For any instance $I = (N, A, d, k, q)$ with $q > k/2$, the q -costs $c_i(C)$ of agents $i \in N$ for k -sized committees of alternatives $C \subseteq A$ form a pseudometric.*

Proof. To prove the statement, we need to show that the q -costs satisfy the triangle inequality, i.e., $c_i(X) \leq c_i(Y) + c_j(Y) + c_j(X)$, for any two agents $i, j \in N$ and two k -sized committees X, Y .

For any $q \leq k$, there exists $x \in \text{top}_j(X)$ that is among the $k - q + 1$ least favorite alternatives (ranked in some position from q to k) of agent i in X ; thus, x is such that $c_i(X) \leq d(i, x)$ and $c_j(X) \geq d(j, x)$. Since $q > k/2$, there also exists $y \in \text{top}_i(Y) \cap \text{top}_j(Y)$, that is, y is such that $c_i(Y) \geq d(i, y)$ and $c_j(Y) \geq d(j, y)$. Combining these with the triangle inequality for the distances between agents and alternatives, we have that

$$\begin{aligned} c_i(X) &\leq d(i, x) \leq d(i, y) + d(j, y) + d(j, x) \\ &\leq c_i(Y) + c_j(Y) + c_j(X), \end{aligned}$$

as desired. \square

Since the q -costs form a metric, we can transform the original profile in which the agents rank the alternatives to one in which the agents rank all possible k -sized committees. In particular, to decide whether agent i prefers a committee X over another committee Y , it suffices to compare her q -th favorite alternatives in X and Y ; we can break ties arbitrarily. Given the rankings of the agents over the committees, we can then employ any single-winner rule to decide the final committee. Specifically, using the PLURALITY-MATCHING rule of Gkatzelis, Halpern, and Shah (2020), we obtain a best-possible distortion bound of 3.

Corollary 1. *For $q > k/2$, there exists a multiwinner voting rule with (k, q) -distortion at most 3.*

Unfortunately, the above approach requires us to apply a single-winner voting rule to a profile consisting of an exponential number of alternatives (the set of all k -sized committees). This is an inherent obstacle in our attempts to get constant distortion bounds by naively applying known deterministic single-winner voting rules.

Interestingly, it is still easy to compute the favorite k -sized committee of an agent, which consists of the agent's k most preferred alternatives; we refer to this as the *top- k* committee of the agent. Consequently, randomized dictatorship, the single-winner voting rule which selects an agent uniformly at random and returns her most favorite alternative, can be efficiently implemented in our setting. Using its distortion analysis by Anshelevich and Postl (2017), we obtain the following.

Corollary 2. *For any $q > k/2$, the (k, q) -distortion of randomized dictatorship, which can be implemented efficiently, is at most $3 - 2/n$.*

In the following, we restrict our attention to deterministic polynomial-time multiwinner voting rules. Our main result provides such a rule with distortion at most 9 by exploiting the following lemma.

Lemma 3. *Consider any instance $I = (N, A, d, k, q)$ with $q > k/2$, and let O be an optimal committee for I . There exists a committee C that is top- k for some agent, and such that $SC(C) \leq 3 \cdot SC(O)$.*

Proof. By Lemma 2, the q -costs of the agents form a metric space. Let j be the closest agent to the optimal committee O according to her q -cost, i.e., $c_j(O) \leq c_i(O)$ for every $i \in N$. Let C be the top- k committee of j , and thus $c_j(C) \leq c_j(O)$. Combining the above with the triangle inequality for the q -cost metric, for any agent $i \in N$, we have

$$\begin{aligned} c_i(C) &\leq c_i(O) + c_j(O) + c_j(C) \\ &\leq c_i(O) + 2c_j(O) \leq 3c_i(O), \end{aligned}$$

The lemma follows by summing over all agents. \square

Essentially, Lemma 3 suggests that the best top- k committee must 3-approximate the optimal committee in terms of social cost. So, by considering the n top- k committees (one per agent) as alternatives, we can deploy the single-winner rule of Gkatzelis, Halpern, and Shah (2020) to obtain in polynomial time a committee that is within a factor of 3 away of the best top- k committee in terms of social cost, and thus has a (k, q) -distortion of at most 9.

Corollary 3. *For any $q > k/2$, there is a polynomial-time multiwinner voting rule with (k, q) -distortion at most 9.*

Actually, the approach used to obtain Corollary 3 can be used as a template that could potentially lead to even lower distortion bounds by deterministic polynomial-time voting rules. All we need is an algorithm that takes as input the ranking profile and decides — in polynomial time — on a set P of k -sized committees, which, for every distance metric consistent with the ranking profile, contains a committee that approximates the optimal social cost within a factor of ρ (e.g., Lemma 3 suggests that this is possible for $\rho = 3$). Then, applying the voting rule of Gkatzelis, Halpern, and Shah (2020) on the reduced ranking profile with only the committees in P as alternatives, we obtain a deterministic polynomial-time multiwinner voting rule with (k, q) -distortion at most 3ρ .

We present two results related to this template. The first one is positive and shows that a guarantee of $\rho = 1$ is possible when the number of agents is constant, making the optimal distortion bound of 3 achievable in polynomial time.

Theorem 4. *For any $q > k/2$ and constant number of agents, there is a deterministic polynomial-time multiwinner voting rule with (k, q) -distortion at most 3.*

Proof. We prove the theorem by defining a voting rule that follows our template with $\rho = 1$. In particular, given a ranking profile, our rule first identifies a set P of at most m^n committees such that one of minimum social cost is always included, no matter which distance function d , consistent with the preference profile, is used in the definition of

the social cost. Then, by invoking the single-winner rule of Gkatzelis, Halpern, and Shah (2020) using the committees in P as alternatives, we get the desired multiwinner voting rule with (k, q) -distortion of at most 3.

To identify the committees in P , we enumerate over all possibilities for the alternatives in a k -sized committee that the agents can have as their q -th closest ones. In particular, for each of the m^n possible vectors of alternatives $\ell = \langle \ell_1, \ell_2, \dots, \ell_n \rangle$, we include in P a committee $C \subseteq A$ with $|C| = k$, $\bigcup_{i=1}^n \{\ell_i\} \subseteq C$, such that ℓ_i is the alternative in C that is the q -th closest to agent i , for $i \in [n]$, if such a committee exists. We refer to each committee having these properties as a (k, q) -completion of the vector of alternatives ℓ . Notice that including in P just one (k, q) -completion for each vector of alternatives is enough for our purposes since any two committees C, C' that are (k, q) -completions for the same vector ℓ of alternatives have the same social cost for a given distance function d : $SC_q(C|d) = SC_q(C'|d) = \sum_{i=1}^n d(i, \ell_i)$. Clearly, since n is a constant, P has polynomial size. To prove that its whole construction takes polynomial time, it suffices to show how a (k, q) -completion C for a given vector of alternatives ℓ can be identified (if it exists) in polynomial time. Essentially, we need to decide which alternatives in addition to the ones in ℓ form a k -sized committee and ensure that alternative ℓ_i is the q -th closest to agent i among those in C . To do so, we define the following classification of the alternatives into types from the set

$$\mathcal{T} = \{\langle t_1, t_2, \dots, t_n \rangle : t_i \in \{+1, 0, -1\}, i \in N\}.$$

An alternative $a \in A$ belongs to type $t \in \mathcal{T}$ if and only if

- $a \succ_i \ell_i$ for each $i \in N$ with $t_i = +1$,
- $a = \ell_i$ for each $i \in N$ with $t_i = 0$, and
- $\ell_i \succ_i a$ for each $i \in N$ with $t_i = -1$.

Notice that the classification is such that replacing an alternative in a (k, q) -completion by another alternative from the same class results in another (k, q) -completion. Hence, to identify a (k, q) -completion C for a given vector of alternatives ℓ , we need to identify the number of alternatives from each type of \mathcal{T} that C should have.

For each type $t \in \mathcal{T}$, denote by $H(t)$ the number of alternatives of type t . Also, set $L(t) = 1$ if t is the type of some alternative in ℓ and $L(t) = 0$, otherwise. The quantities $L(t)$ and $H(t)$ are lower and upper bounds on the number $h(t)$ of alternatives of type t that can be included in a (k, q) -completion of ℓ . In particular, notice that each alternative in ℓ is the unique alternative in its type t , i.e., $H(t) = 1$. Setting $L(t) = 1$ for this type guarantees that this alternative will be included in the (k, q) -completion. Now, the existence of a (k, q) -completion is equivalent to the feasibility of the following integer linear program:

$$\begin{aligned} \sum_{t \in \mathcal{T}} h(t) &= k \\ \sum_{t \in \mathcal{T}: t_i \geq 0} h(t) &= q, \forall i \in N \\ L(t) &\leq h(t) \leq H(t), \forall t \in \mathcal{T} \\ h(t) &\in \mathbb{N}_{\geq 0}, \forall t \in \mathcal{T} \end{aligned}$$

Notice that we can naively check the feasibility of the above ILP by enumerating all possible values for the variables $h(t)$. As there are 3^n types, there are 3^n such variables (i.e., constantly many since n is a constant), each taking at most $m + 1$ values. Even this naive solution takes only polynomial time. Finally, once we have a feasible solution for the above ILP, we can easily define the (k, q) -completion by just including $h(t)$ alternatives of type t in it, for $t \in \mathcal{T}$. \square

Our last result is negative and reveals a limitation of our template. It shows a lower bound of $1 + 2/e$ on ρ (unless $P = NP$) and, consequently, a lower bound of approximately 5.207 on the distortion that can be achieved via this template. A slightly weaker lower bound of 4.5 can be shown using purely information-theoretic arguments that do not rely on any complexity assumption.

Theorem 5. *Let A be an algorithm which on input a ranking profile with n agents and m alternatives, and integers k and q with $1 \leq k/2 < q \leq k < m$, runs in time polynomial in m and n and returns a set P of (polynomially many) k -sized committees of alternatives with the following property: For every distance function d , there is a committee $C \in P$ such that $SC_q(C|d) < (1 + 2/e + \varepsilon) \cdot SC_q(O|d)$, where O is a committee of minimum social cost according to d , and ε is a positive constant. Then, $P = NP$.*

Extensions and Open Problems

In this work, we extended the metric distortion framework to multiwinner voting. By modeling the cost of an agent for a k -sized committee as her distance from her q -favorite alternative therein, we revealed a quite surprising trichotomy on the distortion of multiwinner rules in terms of q and k , providing asymptotically tight bounds. The main question that our work leaves open is to identify the best possible distortion for when $q > k/2$ that can be achieved by an efficient deterministic rule.

We exclusively focused on the social cost and did not consider other objectives, such as the maximum agent cost. It is not hard to observe that our methods provide bounds in terms of this objective as well, albeit not necessarily tight. When $q \leq k/3$, the distortion remains unbounded (in the instances used in the proof of Theorem 1, the optimal committee guarantees cost 0 to all agents, whereas any voting rule yields a positive cost to some agent). When $q \in (k/3, k/2]$, by carefully inspecting the proof of the upper bound of POLAROPPOSITES in Theorem 2, we can show that it achieves constant distortion in terms of the maximum cost. Finally, when $q > k/2$, since the q -costs form a metric space, known results from single-winner voting extend to multiwinner voting; that is, there exists a deterministic multiwinner rule with distortion at most 3, which is based on the rule of Gkatzelis, Halpern, and Shah (2020); this rule is known to achieve a *fairness ratio* of at most 3, which implies 3-approximation of the maximum cost as well. The efficient upper bound of 9 via our template also extends to the maximum cost, as the bound of 3 in Lemma 3 actually holds per agent. In the future, it would be interesting to go beyond objectives such as the social and maximum costs, and consider others that make sense in multiwinner voting.

Acknowledgements

We thank Elliot Anshelevich for fruitful discussions in early stages of this work, and Piotr Skowron for suggesting an example demonstrating unbounded distortion for $q = 1$, which has been generalized to $q \leq k/3$ in Theorem 1.

References

- Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 777–783.
- Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 4294–4301.
- Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58: 797–827.
- Arrow, K. J. 1951. *Social Choice and Individual Values*. Wiley.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2): 461–485.
- Benadè, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 376–382.
- Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227: 190–213.
- Bredereck, R.; Faliszewski, P.; Igarashi, A.; Lackner, M.; and Skowron, P. 2018. Multiwinner Elections With Diversity Constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 933–940.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58: 123–152.
- Caragiannis, I.; and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9-10): 1655–1671.
- Chamberlin, J. R.; and Courant, P. N. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review*, 77(3): 718–733.
- Chen, X.; Li, M.; and Wang, C. 2020. Favorite-Candidate Voting for Eliminating the Least Popular Candidate in a Metric Space. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1894–1901.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3): 599–632.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74(2017): 27–47.
- Gkatzelis, V.; Halpern, D.; and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 1427–1438.
- Goel, A.; Hulett, R.; and Krishnaswamy, A. K. 2018. Relating Metric Distortion and Fairness of Social Choice Rules. arXiv:1810.01092.
- Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, 287–304.
- Jaworski, M.; and Skowron, P. 2020. Evaluating Committees for Representative Democracies: the Distortion and Beyond. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 196–202.
- Kempe, D. 2020a. An analysis framework for metric voting based on LP duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2079–2086.
- Kempe, D. 2020b. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2087–2094.
- Lu, T.; and Boutilier, C. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Meir, R.; Sandomirskiy, F.; and Tennenholtz, M. 2021. Representative committees of peers. *Journal of Artificial Intelligence Research*, 71: 401–429.
- Monroe, B. L. 1995. Fully Proportional Representation. *American Political Science Review*, 89(4): 925–940.
- Munagala, K.; and Wang, K. 2019. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 245–262.
- Peters, D.; Pierczynski, G.; Shah, N.; and Skowron, P. 2021. Market-Based Explanations of Collective Decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5656–5663.
- Peters, D.; and Skowron, P. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 793–794.
- Procaccia, A. D.; and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents (CIA)*, 317–331.
- Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare*, 30(3): 353–362.
- Skowron, P. K.; and Elkind, E. 2017. Social choice under metric preferences: Scoring rules and STV. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, 706–712.