



The distortion of distributed metric social choice [☆]

Elliot Anshelevich ^a, Aris Filos-Ratsikas ^b, Alexandros A. Voudouris ^{c,*}

^a Department of Computer Science, Rensselaer Polytechnic Institute, USA

^b Department of Computer Science, University of Liverpool, UK

^c School of Computer Science and Electronic Engineering, University of Essex, UK



ARTICLE INFO

Article history:

Received 12 November 2021

Received in revised form 21 March 2022

Accepted 24 March 2022

Available online 26 March 2022

Keywords:

Social choice

Distortion

Mechanism design

ABSTRACT

We consider a social choice setting with agents that are partitioned into disjoint groups, and have metric preferences over a set of alternatives. Our goal is to choose a single alternative aiming to optimize various objectives that are functions of the distances between agents and alternatives in the metric space, under the constraint that this choice must be made in a distributed way: The preferences of the agents within each group are first aggregated into a representative alternative for the group, and then these group representatives are aggregated into the final winner. Deciding the winner in such a way naturally leads to loss of efficiency, even when complete information about the metric space is available. We provide a series of (mostly tight) bounds on the distortion of distributed mechanisms for variations of well-known objectives, such as the (average) total cost and the maximum cost, and also for new objectives that are particularly appropriate for this distributed setting and have not been studied before.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The main goal of social choice theory [33] is to come up with outcomes that accurately reflect the collective opinions of individuals within a society. A prominent example is that of elections, where the preferences of voters over different candidates are aggregated into a single winner, or a set of winners in the case of committee elections. Besides elections, the abstract social choice theory setting, where a set of *agents* express preferences over a set of possible *alternatives*, captures very broad decision-making application domains, such as choosing public policies, allocations of resources, or the most appropriate position to locate a public facility.

In the field of computational social choice, Procaccia and Rosenschein [32] defined the notion of *distortion* to measure the loss in an aggregate cardinal objective (typically the utilitarian social welfare), due to making decisions whilst having access to only limited (*ordinal*, in particular) information about the preferences of the agents, rather than their true cardinal values (or costs). Following their work, a lot of effort has been put forward to bound the distortion of social choice rules, with Anshelevich et al. [7] and Anshelevich and Postl [5] being the first to consider settings with *metric* preferences. In such settings, agents and alternatives are points in a metric space, and thus the distances between them (which define the costs of the agents) satisfy the triangle inequality. The metric space can be high-dimensional, and can be thought of as evaluating

[☆] This paper unifies and partially extends earlier versions that appear in *Proceedings of the 14th International Symposium on Algorithmic Game Theory* [21] and *Proceedings of the 17th Conference on Web and Internet Economics* [9].

* Corresponding author.

E-mail address: alexandros.voudouris@essex.ac.uk (A.A. Voudouris).

the proximity between agents and alternatives for different political issues or ideological axes (e.g., liberal to conservative, or libertarian to authoritarian). The distortion in metric social choice has received significant attention, with many variants of the main setting being considered over the recent years.

In contrast to the *centralized* decision-making settings considered in the papers mentioned above, there are cases where it is logistically too difficult to aggregate the preferences of the agents directly, or different groups of agents play inherently different roles in the process. In such scenarios, the collective decisions have to be carried out in a *distributed* manner, as follows. The agents are partitioned into groups (such as electoral districts, focus groups, or sub-committees), and the members of each group locally decide a single alternative that is representative of their preferences, without taking into account the agents of different groups. Then, the final outcome is decided based on properties of the group representatives, and not on the underlying agents within the groups; for example, the representatives act as agents themselves and choose an outcome according to their own preferences. However, since the representatives cannot perfectly capture all the information about the preferences of the agents (even when it is available in the group level), it is not surprising that choosing the final outcome this way may lead to loss of efficiency.

Motivated by this, Filos-Ratsikas et al. [23] initiated the study of the deterioration of the social welfare in general normalized distributed social choice settings. They extended the notion of distortion to account for the information about the preferences of the agents that is lost after the local decision step, and showed bounds on the distortion of max-weight mechanisms when the number of groups is given. In this paper, we consider the distributed distortion problem with metric preferences and bound the distortion of (cardinal and ordinal) mechanisms based on the following axes: We show bounds for (a) general metrics (including refined bounds for the line metric), (b) four different objectives (including the average total cost, the maximum cost, as well as two new objectives that are tailor-made and clearly motivated by the distributed nature of the setting), and (c) groups of agents that could vary in size. We paint an almost complete picture of the distortion landscape of distributed mechanisms when the agents have metric preferences.

1.1. Our contributions

We consider a distributed social choice setting with a set of agents and a set of alternatives, all of whom are located in a metric space. The preferences of the agents for the alternatives are given by their distances in the metric space, and as such, they satisfy the triangle inequality. Furthermore, the agents are partitioned into a given set of *districts* of possibly different sizes. A distributed mechanism selects an alternative based on the preferences of the agents in two steps: first, each district selects a *representative* alternative using some local aggregation rule, and then the mechanism uses *only* information about the representatives to select the final winning alternative.

The goal is to choose the alternative that optimizes some aggregate objective that is a function of the distances between agents and alternatives. In the main part of the paper we consider the following four cost minimization objectives, which can be defined as compositions of objectives applied over and within the districts:

- The *average of the average agent distance in each district*¹ (denoted by $\text{AVG} \circ \text{AVG}$);
- The *average of the maximum agent distance in each district* (denoted by $\text{AVG} \circ \text{MAX}$);
- The *maximum agent distance in any district* (denoted by $\text{MAX} \circ \text{MAX}$);
- The *maximum of the average agent distance in each district* (denote by $\text{MAX} \circ \text{AVG}$).

While $\text{AVG} \circ \text{AVG}$ and $\text{MAX} \circ \text{MAX}$ are adaptations of objectives that have been considered in the centralized setting, $\text{AVG} \circ \text{MAX}$ and $\text{MAX} \circ \text{AVG}$ are only meaningful in the context of distributed social choice. In particular, $\text{MAX} \circ \text{AVG}$ can be thought of as a fairness-inspired objective guaranteeing that no district has a very large cost, where the cost of a district is the average cost of its members. Similarly, $\text{AVG} \circ \text{MAX}$ guarantees that the average district cost is small, where the cost of a district is now defined as the egalitarian (maximum) cost of any of its members. We consider the introduction and study of these objectives as one of the major contributions of our work.

We measure the performance of a distributed mechanism by its *distortion*, defined as the worst-case ratio (over all instances of the problem) between the objective value of the alternative chosen by the mechanism and the minimum possible objective value achieved over all alternatives. The distortion essentially measures the deterioration of the objective due to the fact that the mechanism must make a decision via a distributed two-step process, on top of other possible informational limitations related to the preferences of the agents. We consider *deterministic* mechanisms that are either *cardinal* (in which case they have access to the exact distances between agents and alternatives) or *ordinal* (in which case they have access only to the rankings that are induced by the distances). Table 1 gives an overview of our bounds on the distortion of distributed mechanisms for the four objectives defined above. We provide bounds that hold for general metric spaces, and then more refined bounds for the fundamental special case where the metric is a line.

¹ Note that this objective is not exactly equivalent to the well-known (average) *social cost* objective, defined as the (average) total agent distance over all districts. All of our results extend for this objective as well, by adapting our mechanisms to weigh the representatives proportionally to the district sizes. When all districts have the same size, $\text{AVG} \circ \text{AVG}$ coincides with the average social cost.

Table 1

An overview of the distortion bounds for the various settings studied in this paper. Each entry consists of an interval showing a lower bound on the distortion of all deterministic distributed mechanisms for the corresponding setting, and an upper bound that is achieved by some mechanism; when a single number is presented, the bound is tight.

	General metric		Line metric	
	Cardinal	Ordinal	Cardinal	Ordinal
AVG $\circ$ AVG	3	[7, 11]	3	7
AVG $\circ$ MAX	3	$[2 + \sqrt{5}, 11]$	3	$[2 + \sqrt{5}, 5]$
MAX $\circ$ MAX	$[1 + \sqrt{2}, 3]$	[3, 5]	$1 + \sqrt{2}$	3
MAX $\circ$ AVG	$[1 + \sqrt{2}, 3]$	$[2 + \sqrt{5}, 5]$	$1 + \sqrt{2}$	$[2 + \sqrt{5}, 5]$

1.2. Our techniques

Several of our bounds for general metric spaces follow from a novel composition technique for designing distributed mechanisms. In particular, we prove a rather general *composition theorem*, which appears in many versions throughout our paper, depending on the objective at hand. Roughly speaking, the theorem relates the distortion of a distributed mechanism to the distortion of the centralized voting rules it uses for the local (in-district) and global (over-districts) aggregation steps. In particular, for two such voting rules with distortion bounds α and β , the distortion of the composed mechanism is at most $\alpha + \beta + \alpha\beta$. This effectively enables us to plug in voting rules with known distortion bounds, and obtain distributed mechanisms with low distortion. The theorem is also robust in the sense that the AVG and MAX objectives can be substituted with more general objectives satisfying specific properties, such as monotonicity and subadditivity; we provide more details on that in Section 6.2.

To demonstrate the strength of the composition theorem, consider the objective $\text{AVG} \circ \text{AVG}$. The upper bound of 3 for cardinal mechanisms and general metrics is obtained by using optimal centralized voting rules (with distortion 1) for both aggregation steps. Similarly, the upper bound of 11 for ordinal mechanisms follows by using the ordinal PLURALITYMATCHING rule of Gkatzelis et al. [24] in both steps of the distributed mechanism; this rule is known to have distortion at most 3 for general instances, and at most 2 when all agents are at distance 0 from their most-preferred alternative (which is the case when the representatives are thought of as agents in the second step of the mechanism).

Even though the composition theorem is evidently a very powerful tool, it comes short of providing tight bounds in some cases. To this end, we design explicit distributed mechanisms with improved distortion guarantees, both for general metrics as well as the fundamental special case where the metric is a line. A compelling highlight of our work is a novel mechanism for objectives of the form $\text{MAX} \circ G$, to which we refer as λ -ACCEPTABLE-RIGHTMOST-LEFTMOST (λ -ARL). While this mechanism has the counter-intuitive property of *not* being *unanimous* (i.e., there are cases where all agents agree on the best alternative, but the mechanism does not choose this alternative as the winner), it achieves the best possible distortion of $1 + \sqrt{2}$ among all distributed mechanisms on the line. In contrast, we prove that unanimous mechanisms cannot achieve distortion better than 3. To the best of our knowledge, this is the first time that not satisfying unanimity turns out to be a necessary ingredient for achieving the best possible distortion in the metric social choice literature.

For the lower bounds, we employ the following main idea. We construct instances of the problem for which any mechanism with low distortion (depending on the bound we are aiming for) must choose some specific alternative for some particular district. Then, because of the distributed nature of the mechanism, we can exploit the fact that this alternative must represent this district in any instance that contains it, and use such instances to either argue about the distortion of the mechanism, or to impose constraints on the representatives of other districts. In some of our more involved constructions, we use the idea highlighted above inductively, together with arguments involving ordinality when necessary (for example, see Lemma 4.1 and Lemma 5.2).

1.3. Related work

The distortion of centralized social choice voting rules has been studied extensively for many different settings. For a comprehensive introduction to the distortion literature, we refer the interested reader to the recent survey of Anshelevich et al. [8].

After the work of Procaccia and Rosenschein [32], a series of papers adopted their normalized setting, where the agents have unit-sum values for the alternatives, and proved asymptotically tight bounds on the distortion of ordinal single-winner rules [14,15], multi-winner rules [16], rules that choose rankings of alternatives [11], and strategyproof rules [12]. Recent papers considered more general questions related to how the distortion is affected by the amount of available information about the values of the agents [2,29,30]. The normalized distortion has also been investigated in other related problems, such as participatory budgeting [10] and one-sided matching [3,22].

The metric distortion setting was first considered by Anshelevich et al. [7] who, among many results, showed a lower bound of 3 on the distortion of deterministic single-winner ordinal rules for the social cost, and an upper bound of 5, achieved by the Copeland rule. Following their work, many papers were devoted to bridging this gap (e.g., see [28,31])

until, finally, Gkatzelis et al. [24] designed the PLURALITYMATCHING rule that achieves an upper bound of 3; in fact, this bound holds for the more general *fairness ratio* [25] (which captures various different objectives, including the social cost and the maximum cost). Besides the main setting, many other works have shown bounds on the metric distortion for randomized rules [5,20], rules that use less than ordinal information about the preferences of the agents [4,19,27], committee elections [17,18,26], primary elections [13], and for many other problems [1,6].

2. Preliminaries

An instance of our problem is defined as a tuple $I = (N, A, D, \delta)$, where

- N is a set of n agents.
- A is a set of m alternatives.
- D is a collection of k districts, which define a partition of N (i.e., each agent belongs to a single district). Let N_d be the set of agents that belong to district $d \in D$, and denote by $n_d = |N_d|$ the size of d .
- δ is a metric space that contains points representing the agents and the alternatives. In particular, δ defines a distance $\delta(i, j)$ between any $i, j \in N \cup A$, such that the triangle inequality is satisfied, i.e., $\delta(i, j) \leq \delta(i, x) + \delta(x, j)$ for every $i, j, x \in N \cup A$.

A distributed mechanism takes as input information about the metric space, which can be of cardinal or ordinal nature. In particular, a *cardinal* mechanism is given access to the distances between agents and alternatives, while an *ordinal* mechanism is given as input *strict linear orderings* (or, *rankings*) that are induced by the distances (an agent i ranks alternative x higher than alternative y in case $\delta(i, x) \leq \delta(i, y)$). The output of the mechanism is a single winner alternative $w \in A$ that is computed by implementing the following two steps:

- Step 1: For every district $d \in D$, the agents therein decide a *representative* alternative $y_d \in A$.
- Step 2: Given the district representatives, the output is an alternative $w \in A$.

In both steps, the decisions are made by using *direct voting rules*, which map the preferences of a given subset of agents to an alternative. To be more specific, in the first step, an *in-district* direct voting rule is applied for each district $d \in D$ with input the preferences of the agents in the district (set N_d) to decide its representative $y_d \in A$. Then, in the second step, the district representatives can be thought of as pseudo-agents, and an *over-districts* direct voting rule is applied with input their preferences to decide the final winner $w \in A$.² In the special case of instances consisting of a single district, the process is not distributed, and thus the two steps collapse into one: The final winner is the alternative chosen to be the district's representative.

Observe that, since a distributed mechanism uses specific in-district and over-districts voting rules, it satisfies the following useful properties: The representative y_d of some district d is the same over all instances where the same amount of information about d is available. In addition, the final winner is the same over all instances where the same alternatives have been selected as district representatives and the same amount of information about them is available. We extensively exploit these properties in our lower bound constructions.

2.1. Objectives

We consider standard minimization objectives that have been studied in the related literature and also propose new ones that are appropriate in the context distributed setting. Each objective assigns a value to every alternative as a cost function composition $F \circ G$ of an objective function F that is applied over the districts and an objective function G that is applied within the districts. Our main four objectives are defined by considering all possible combinations of $F, G \in \{\text{AVG}, \text{MAX}\}$, where the functions AVG and MAX define an average and a max over districts or agents within a district, respectively. In particular, we have:

- The *average of average cost* of an alternative $j \in A$ is defined as

$$(\text{AVG} \circ \text{AVG})(j|I) = \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, j) \right).$$

- The *max of max cost* of an alternative $j \in A$ is defined as

$$(\text{MAX} \circ \text{MAX})(j|I) = \max_{d \in D} \max_{i \in N_d} \delta(i, j) = \max_{i \in N} \delta(i, j).$$

² Even though we consider over-districts rules that can in principle choose the final winner from the set of all alternatives, most of our mechanisms actually use rules that select the final winner from the set of representatives only; see the discussion in Section 6.1.

- The *average of max cost* of an alternative $j \in A$ is defined as

$$(\text{AVG} \circ \text{MAX})(j|I) = \frac{1}{k} \sum_{d \in D} \max_{i \in N_d} \delta(i, j).$$

- The *max of average cost* of an alternative $j \in A$ is defined as

$$(\text{MAX} \circ \text{AVG})(j|I) = \max_{d \in D} \left\{ \frac{1}{n_d} \sum_{i \in N_d} \delta(i, j) \right\}.$$

The $\text{AVG} \circ \text{AVG}$ objective is similar to the well-known utilitarian *average social cost* objective measuring the average total distance between all agents and alternative j ; actually, $\text{AVG} \circ \text{AVG}$ coincides with the average social cost when the districts are symmetric (i.e., have the same size), but not in general. The $\text{MAX} \circ \text{MAX}$ objective coincides with the egalitarian *max cost* measuring the maximum distance from j among all agents. The new objectives $\text{AVG} \circ \text{MAX}$ and $\text{MAX} \circ \text{AVG}$ make sense in the context of distributed voting, and can be thought of as measures of fairness between districts. For example, minimizing the $\text{MAX} \circ \text{AVG}$ objective corresponds to making sure that the final choice treats each district fairly so that the average social cost of each district is almost equal to that of any other district. Of course, besides combinations of AVG and MAX , one can define many more objectives; we consider such generalizations in Section 6.2.

2.2. Distortion of voting rules and distributed mechanisms

Direct voting rules can be suboptimal, especially when they have limited access to the metric space (for example, when ordinal information is known about the preferences of the agents over the alternatives). This inefficiency is typically captured in the related literature by the notion of distortion, which is the worst-case ratio between the objective value of the optimal alternative over the objective value of the alternative chosen by the rules. Formally, given a minimization cost objective $F \in \{\text{AVG}, \text{MAX}\}$, the F -distortion of a voting rule V is

$$\text{dist}_F(V) = \sup_{I=(N,A,\delta)} \frac{F(V(I)|I)}{\min_{j \in A} F(j|I)},$$

where $V(I)$ denotes the alternative chosen by the voting rule when given as input the (single-district) instance I consisting of a set of agents N , a set of alternatives A , and a metric space δ .

The notion of distortion can be naturally extended for the case of distributed mechanisms. Given a composition objective $F \circ G$, the $(F \circ G)$ -distortion of a distributed mechanism M is

$$\text{dist}_{F \circ G}(M) = \sup_{I=(N,A,D,\delta)} \frac{(F \circ G)(M(I)|I)}{\min_{j \in A} (F \circ G)(j|I)},$$

where $M(I)$ is the alternative chosen by the mechanism when given as input an instance I consisting of a set of agents N , a set of alternatives A , a set D of districts, and a metric space δ . Clearly, the distortion of any mechanism is at least 1. In case the denominator of the ratio (or tends to) 0 while the numerator is positive, we will say that the distortion is unbounded. Our goal is to bound the distortion of distributed mechanisms for the different objectives we consider. To this end, we will either show how known results from the literature about the distortion of direct voting rules can be composed to yield distortion bounds for distributed mechanisms or design explicit mechanisms with low distortion.

Before we proceed to our technical results, we argue that it is without loss of generality to consider mechanisms satisfying particular unanimity properties, which will be used extensively in all our lower bound constructions in the upcoming sections. We say that a distributed mechanism satisfies *unanimity* (is *unanimous*) if it chooses the representative of a district d to be an alternative a whenever all agents in N_d prefer a over all other alternatives. We say that a distributed mechanism satisfies *cardinal unanimity* (is *cardinally unanimous*) if it chooses the representative of a district d to be an alternative a whenever all agents in N_d are positioned at the same point as a in the metric space and all alternatives are positioned somewhere else. Observe that unanimity implies cardinal unanimity, but not vice versa. The proofs of the following lemmas are similar.

Lemma 2.1. *For any $F, G \in \{\text{AVG}, \text{MAX}\}$, every ordinal distributed mechanism with finite $(F \circ G)$ -distortion must satisfy unanimity.*

Proof. Suppose towards a contradiction that there is an ordinal distributed mechanism M with finite $(F \circ G)$ -distortion that is not unanimous. Then, there must exist a district d consisting of agents all of whom prefer an alternative a over all other alternatives, such that M chooses the representative of d to be some alternative $b \neq a$. Now, consider the instance I that includes only this district d , which means that M chooses b to be the final winner as well. A metric space that is consistent to the available ordinal information is the one where alternative a and all agents are positioned at the same point, while all other alternatives (including b) are positioned at different points. For this metric space, we have that $(F \circ G)(a|I) = G(a|I) = 0$ and $(F \circ G)(b|I) = G(b|I) > 0$, yielding that the distortion of M is unbounded, a contradiction. $\square$

Lemma 2.2. For any $F, G \in \{\text{AVG}, \text{MAX}\}$, every cardinal distributed mechanism with finite $(F \circ G)$ -distortion must satisfy cardinal unanimity.

Proof. Suppose towards a contradiction that there is a cardinal distributed mechanism M with finite $(F \circ G)$ -distortion that does not satisfy cardinal unanimity. Then, there must exist a district d consisting of agents all of whom are positioned at the same point as an alternative a while all other alternatives are positioned somewhere else, such that M chooses the representative of d to be some alternative $b \neq a$. Now, consider the instance I that includes only this district d . Since there is a single district in I , b is chosen to be the final winner as well. However, since b is not located at the same point as a in the metric space, we have that $(F \circ G)(a|I) = G(a|I) = 0$ and $(F \circ G)(b|I) = G(b|I) > 0$, yielding that the distortion of M is unbounded, a contradiction. $\square$

3. Composition results for general metrics

In this section, we consider general metric spaces, and show how known distortion bounds for direct voting rules can be composed to yield distortion bounds for distributed mechanisms that rely on those voting rules. Given an objective $F \circ G$, we say that a distributed mechanism is α -in- β -over if it works as follows.

α -in- β -over mechanisms

1. For each district $d \in D$, choose its representative using an in-district voting rule with G -distortion at most α .
2. Choose the final winner using an over-districts voting rule with F -distortion at most β .

Our first technical result is an upper bound on the $(F \circ G)$ -distortion of α -in- β -over mechanisms, for any $F, G \in \{\text{AVG}, \text{MAX}\}$. The proof of the following theorem also follows from the more general Theorem 6.5 in Section 6.2.1, which considers objectives that are compositions of functions satisfying particular properties, such as monotonicity and subadditivity.

Theorem 3.1. For any $F, G \in \{\text{AVG}, \text{MAX}\}$, the $(F \circ G)$ -distortion of any α -in- β -over mechanism is at most $\alpha + \beta + \alpha\beta$.

Proof. Here, we present a proof only for the $\text{AVG} \circ \text{AVG}$ objective; the proof for the other objectives is similar. Consider an arbitrary α -in- β -over mechanism M and an arbitrary instance $I = (N, A, D, \delta)$. Let w be the alternative that M outputs as the final winner when given I as input, and denote by o an optimal alternative. By the definition of M , we have the following two properties:

$$\forall j \in A, d \in D : \sum_{i \in N_d} \delta(i, y_d) \leq \alpha \sum_{i \in N_d} \delta(i, j) \quad (1)$$

and

$$\forall j \in A : \sum_{d \in D} \delta(y_d, w) \leq \beta \sum_{d \in D} \delta(y_d, j) \quad (2)$$

By the triangle inequality, we have $\delta(i, w) \leq \delta(y_d, w) + \delta(i, y_d)$ for any agent $i \in N$. Using this, we obtain

$$\begin{aligned} (\text{AVG} \circ \text{AVG})(w|I) &= \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, w) \right) \\ &\leq \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(y_d, w) \right) + \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, y_d) \right) \\ &= \frac{1}{k} \sum_{d \in D} \delta(y_d, w) + \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, y_d) \right). \end{aligned}$$

By (1) and (2) for $j = o$, we obtain

$$(\text{AVG} \circ \text{AVG})(w|I) \leq \beta \cdot \frac{1}{k} \sum_{d \in D} \delta(y_d, o) + \alpha \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, o) \right)$$

$$= \beta \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(y_d, o) \right) + \alpha \cdot (\text{AVG} \circ \text{AVG})(o|I).$$

By the triangle inequality, we have $\delta(y_d, o) \leq \delta(i, y_d) + \delta(i, o)$ for any agent $i \in N$. Using this and (1) for $j = o$, we can upper bound the first term of the last expression above as follows:

$$\begin{aligned} \beta \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(y_d, o) \right) &\leq \beta \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, y_d) \right) + \beta \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, o) \right) \\ &\leq (\beta + \alpha\beta) \cdot \frac{1}{k} \sum_{d \in D} \left(\frac{1}{n_d} \sum_{i \in N_d} \delta(i, o) \right) \\ &= (\beta + \alpha\beta) \cdot (\text{AVG} \circ \text{AVG})(o|I). \end{aligned}$$

Putting everything together, we obtain a distortion upper bound of $\alpha + \beta + \alpha\beta$. $\square$

By applying Theorem 3.1 using known results, we can derive bounds on the $(F \circ G)$ -distortion of distributed mechanisms for any $F, G \in \{\text{AVG}, \text{MAX}\}$. In particular, due to the structure of F and G , if the whole metric space is known (that is, we have access to the exact distances between agents and alternatives), then we can easily compute the alternative that optimizes F and G . In other words, there exist direct voting rules with F - and G -distortion 1, which can be used to obtain a 1-in-1-over distributed mechanism, and the following statement.

Corollary 3.2. *For any $F, G \in \{\text{AVG}, \text{MAX}\}$, there exists a cardinal distributed mechanism with $(F \circ G)$ -distortion at most 3.*

If only ordinal information is available about the distances between agents and alternatives, then we can employ the PLURALITYMATCHING voting rule of Gkatzelis et al. [24] both within and over the districts. This rule is known to achieve the best possible distortion of 3 among all ordinal rules, for any $F, G \in \{\text{AVG}, \text{MAX}\}$. In fact, this rule achieves a distortion bound of 2 when all agents are at distance 0 from their top alternative; this is the case when the agents are a subset of the alternatives as in the second step of a distributed mechanism.³ Hence, we have a 3-in-2-over mechanism, and Theorem 3.1 yields the following statement.

Corollary 3.3. *For any $F, G \in \{\text{AVG}, \text{MAX}\}$, there exists an ordinal distributed mechanism with $(F \circ G)$ -distortion at most 11.*

Corollaries 3.2 and 3.3 demonstrate the power that Theorem 3.1 gives us in designing distributed mechanisms with constant distortion upper bounds, by using known results from the literature as black boxes. However, this method does not always lead to the best possible distributed mechanisms. In particular, let us consider the objectives $\text{MAX} \circ G$, for $G \in \{\text{AVG}, \text{MAX}\}$ and the class of ordinal mechanisms. We can improve upon the upper bound of 11 due to Corollary 3.3 using the following, much simpler mechanisms.

α -in-arbitrary-over mechanisms

1. For each district $d \in D$, choose its representative using an ordinal in-district voting rule with G -distortion at most α .
2. Output an arbitrary representative as the final winner.

Theorem 3.4. *For any $G \in \{\text{AVG}, \text{MAX}\}$, the $(\text{MAX} \circ G)$ -distortion of any α -in-arbitrary-over mechanism is at most $2 + \alpha$.*

Proof. We will present a proof only for the $\text{MAX} \circ \text{AVG}$ objective; the proof for $\text{MAX} \circ \text{MAX}$ follows by similar, even simpler arguments. Consider any α -in-arbitrary-over mechanism M and any instance $I = (N, A, D, \delta)$. Let w be the alternative chosen by M when given as input the ordinal information of I , and denote the optimal alternative by o . In addition, let $d^* \in \arg \max_{d \in D} \frac{1}{n_d} \sum_{i \in N_d} \delta(i, w)$ be the district that gives the max cost for w . By the triangle inequality, and since $(\text{MAX} \circ \text{AVG})(o|I) \geq \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, o)$ for every $d \in D$, we have

³ See Theorem 1 and Proposition 6 in the arxiv version of the paper of Gkatzelis et al. [24].

$$\begin{aligned}
(\text{MAX} \circ \text{AVG})(w|I) &= \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, w) \\
&\leq \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \left(\delta(i, o) + \delta(w, o) \right) \\
&\leq (\text{MAX} \circ \text{AVG})(o|I) + \delta(w, o)
\end{aligned}$$

Now, let d_w be a district whose representative is w . By the triangle inequality and the fact that the in-district voting rule used to choose w as the representative of d_w has AVG-distortion at most α , we obtain

$$\begin{aligned}
\delta(w, o) &\leq \frac{1}{n_{d_w}} \sum_{i \in N_{d_w}} \left(\delta(i, o) + \delta(i, w) \right) \\
&\leq (1 + \alpha) \cdot \frac{1}{n_{d_w}} \sum_{i \in N_{d_w}} \delta(i, o) \\
&\leq (1 + \alpha) \cdot (\text{MAX} \circ \text{AVG})(o|I).
\end{aligned}$$

Putting everything together, we obtain a distortion upper bound of $2 + \alpha$. $\square$

Now, using again the voting rule of Gkatzelis et al. [24] with AVG- and MAX-distortion at most 3 as an in-district rule, we obtain a 3-in-arbitrary-over distributed mechanism, and Theorem 3.4 implies the following result.

Corollary 3.5. *For any $G \in \{\text{AVG}, \text{MAX}\}$, there exists an ordinal distributed mechanism with $(\text{MAX} \circ G)$ -distortion at most 5.*

4. Line metric: cardinal mechanisms

In this section, as well as in the next one, we focus on the line metric, where agents and alternatives are points on the line of real numbers. Exploiting this structure, there are classes of mechanisms for which we can obtain significantly improved bounds compared to those implied by the general composition Theorem 3.1 and Theorem 3.4.

We start with distributed mechanisms that have access to the line metric, and are thus aware of the distances between agents and alternatives. Recall that for such mechanisms, Corollary 3.2 implies a distortion bound of 3 for all the objectives we have considered so far. For objectives of the form $\text{AVG} \circ G$, we will show that the bound of 3 is tight, even when the metric is a line. This lower bound will follow for both objectives, by considering instances with districts consisting of a single agent (or, more generally, the same number of agents located at the same position), so that $\text{AVG} = \text{MAX}$ within each district. To do that, we will need a structural lemma to argue about the representative of particular districts that all distributed mechanisms must choose to achieve distortion smaller than 3.

Lemma 4.1. *For any $G \in \{\text{AVG}, \text{MAX}\}$, let M be a mechanism with $(\text{AVG} \circ G)$ -distortion strictly less than 3. Consider any instance with two alternatives a and b that are located at 0 and 1, respectively. Then, for every integer $\mu \geq 1$, M must choose alternative a as the representative of the district in which all agents are positioned at $\frac{2\mu+1}{4(\mu+1)}$.*

Proof. We prove the statement by induction on μ .

Base case for $\mu = 1$:

We will first argue that for any instance I with three districts, such that a is the representative of one district, and b is the representative of two districts, M must choose b as the overall winner. Suppose that this is not true, and let I include the following three districts:

- The first district consists of an agent positioned at $1/4$. The representative of this district must be a , as otherwise the distortion of the mechanism for an instance consisting only of (copies of) this district would be 3.
- Each of the other two districts consist of an agent positioned at 1. Due to cardinal unanimity, the representative of this district must be b .

Since

$$(\text{AVG} \circ G)(a|I) = \frac{1}{3} \left(\frac{1}{4} + 2 \right) = \frac{1}{3} \cdot \frac{9}{4}$$

and

$$(\text{AVG} \circ G)(b|I) = \frac{1}{3} \cdot \frac{3}{4},$$

the distortion of M is at least 3, a contradiction.

Now, assume towards a contradiction that the representative of the district in which all agents are positioned at $\frac{2\mu+1}{4(\mu+1)} = 3/8$ is b . Consider the following instance J with three districts:

- In the first district, there is an agent positioned at 0. Due to cardinal unanimity, the representative of this district must be a .
- In each of the other two districts, there is an agent positioned at $3/8$. By our assumption, the representative of these two districts is b .

Since a is the representative of one district and b is the representative of two districts, by our discussion above about instance I , M will choose b as the overall winner for J . However, since

$$(\text{AVG} \circ G)(a|J) = \frac{1}{3} \cdot 2 \cdot \frac{3}{8} = \frac{1}{3} \cdot \frac{3}{4}$$

and

$$(\text{AVG} \circ G)(b|J) = \frac{1}{3} \left(1 + 2 \cdot \frac{5}{8} \right) = \frac{1}{3} \cdot \frac{9}{4},$$

we have that the distortion of M is at least 3, a contradiction.

Induction step:

We assume the statement is true for $\mu = \ell - 1$, and will show that it is also true for $\mu = \ell$.

We will first argue that for any instance I with $2\ell + 1$ districts, such that a is the representative of ℓ districts, and b is the representative of the remaining $\ell + 1$ districts, M must choose b as the overall winner. Consider the following $2\ell + 1$ districts:

- In each of the first ℓ districts, there is an agent positioned at $\frac{2\ell-1}{4\ell} = \frac{2(\ell-1)+1}{4((\ell-1)+1)}$. By the induction hypothesis, the representative of all these districts is a .
- In each of the remaining $\ell + 1$ districts, there is an agent positioned at 1. Due to cardinal unanimity, the representative of these districts is b .

We now have that

$$(\text{AVG} \circ G)(a|I) = \frac{1}{2\ell+1} \left(\ell \cdot \frac{(2\ell-1)}{4\ell} + (\ell+1) \right) = \frac{1}{2\ell+1} \cdot \frac{3(2\ell+1)}{4}$$

and

$$(\text{AVG} \circ G)(b|I) = \frac{1}{2\ell+1} \cdot \ell \cdot \frac{(2\ell+1)}{4\ell} = \frac{1}{2\ell+1} \cdot \frac{(2\ell+1)}{4}.$$

If M chooses a as the overall winner, then its distortion would be at least 3, leading to a contradiction. So, the overall winner must be b .

Now, assume towards a contradiction that the representative of the district in which all agents are positioned at $\frac{2\ell+1}{4(\ell+1)}$ is b . Consider the following instance J with $2\ell + 1$ districts:

- In each of the first ℓ districts, there is an agent positioned at 0. Due to cardinal unanimity, the representative of these districts is a .
- In each of the remaining $\ell + 1$ districts, there is an agent positioned at $\frac{2\ell+1}{4(\ell+1)}$. By our assumption, the representative of these districts is b .

Since a is the representative of ℓ districts and b is the representative of $\ell + 1$ districts, by our discussion above about instance I , M will choose b as the overall winner for J . However, since

$$(\text{AVG} \circ G)(a|J) = \frac{1}{2\ell+1} \cdot (\ell+1) \cdot \frac{(2\ell+1)}{4(\ell+1)} = \frac{1}{2\ell+1} \cdot \frac{\lambda(2\ell+1)}{4}$$

and

$$(\text{AVG} \circ G)(b|J) = \frac{1}{2\ell+1} \left(\ell + (\ell+1) \cdot \frac{(2\ell+3)}{4(\ell+1)} \right) = \frac{1}{2\ell+1} \cdot \frac{3(2\ell+1)}{4},$$

the distortion of M is at least 3, a contradiction.

This concludes the proof of the lemma. $\square$

We are now ready to prove the lower bound.

Theorem 4.2. *For any $G \in \{\text{AVG}, \text{MAX}\}$, the $(\text{AVG} \circ G)$ -distortion of any distributed mechanism is at least $3 - \varepsilon$, for any $\varepsilon > 0$, even when the metric is a line.*

Proof. Let $\mu \geq 1$ be some integer. Consider any distributed mechanism M , which is given as input an instance I with two alternatives a and b located at 0 and 1, respectively, and the following two districts:

- In the first district, there is an agent that is positioned at $\frac{2\mu+1}{4(\mu+1)}$. By Lemma 4.1, the representative of this district is a .
- In the second district, there is an agent that is positioned at 1. Due to cardinal unanimity, the representative of this district is b .

So, we have an instance in which a is the representative of one district, and b is the representative of the other district. Recall that in such a case it is without loss of generality to assume that M chooses the left-most district representative as the overall winner, that is, alternative a (located at 0). Since

$$(\text{AVG} \circ G)(a|I) = \frac{1}{2} \left(1 + \frac{2\mu}{4(\mu+1)} \right) = \frac{1}{2} \cdot \frac{6\mu+5}{4(\mu+1)}$$

and

$$(\text{AVG} \circ G)(b|I) = \frac{1}{2} \left(1 - \frac{2\mu}{4(\mu+1)} \right) = \frac{1}{2} \cdot \frac{2\mu+3}{4(\mu+1)},$$

this distortion leads to a distortion of at least $\frac{6\mu+5}{2\mu+3} = 3 - \frac{4}{2\mu+3}$. By choosing a sufficiently large value for μ , we obtain the desired lower bound. $\square$

For objectives of the form $\text{MAX} \circ G$, we design a novel distributed mechanism that is tailor-made for the line metric and achieves a distortion of at most $1 + \sqrt{2} \leq 2.42$. This mechanism is particularly interesting. While it satisfies cardinal unanimity, it does *not* satisfy unanimity: Even when all agents in a district prefer an alternative a to every other alternative (i.e., a is the closest alternative to all agents), the mechanism may end up choosing a different alternative as the representative. In fact, to break the distortion barrier of 3, our mechanism *has* to be non-unanimous, as we can prove that any unanimous distributed mechanism cannot achieve a $(\text{MAX} \circ G)$ -distortion better than 3; this is obtained by Theorem 5.7 in Section 5, by setting $x = 1$.

For a given $\lambda \geq 1$, we say that an alternative is λ -acceptable for a district $d \in D$ if her G -value for the agents in N_d is at most λ times the G -value of any other alternative for the agents in N_d . Given an objective G , we define a class of distributed mechanisms parameterized by λ that work as follows:

λ -ACCEPTABLE-RIGHTMOST-LEFTMOST (λ -ARL)

1. For each district d , choose its representative to be the *rightmost* λ -acceptable alternative for the district.
2. Output the leftmost district representative as the final winner.

Theorem 4.3. *For any $G \in \{\text{AVG}, \text{MAX}\}$, the $(\text{MAX} \circ G)$ -distortion of λ -ARL is at most $\max \left\{ 2 + \frac{1}{\lambda}, \lambda \right\}$.*

Proof. We will present the proof only for the $\text{MAX} \circ \text{AVG}$ objective; the proof for $\text{MAX} \circ \text{MAX}$ follows by similar, even simpler arguments. Consider an arbitrary instance $I = (N, A, D, \delta)$, where δ is a line metric. Let w be the alternative selected by λ -ARL when given as input I , and denote by o the optimal alternative. We consider the following two cases depending on the relative positions of w and o .

Case 1: w is to the left of o . Let $d^* \in \arg \max_{d \in D} \sum_{i \in N_d} \delta(i, w)$ be the district that gives the max cost for w . By the triangle inequality, and since $(\text{MAX} \circ \text{AVG})(o|I) \geq \frac{1}{n_d} \sum_{i \in N_d} \delta(i, o)$ for every $d \in D$, we have

$$\begin{aligned}
(\text{MAX} \circ \text{AVG})(w|I) &= \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, w) \\
&\leq \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \left(\delta(i, o) + \delta(w, o) \right) \\
&\leq (\text{MAX} \circ \text{AVG})(o|I) + \delta(w, o).
\end{aligned}$$

Now, let d_w be a district whose representative is w . In other words, w is the rightmost λ -acceptable alternative for d_w . Let o_w be the alternative that minimizes the total distance of the agents in d_w . We make the following two observations:

- Since o_w is trivially λ -acceptable, it must be the case that she is located to the left of w , and thus $\delta(w, o) \leq \delta(o_w, o)$.
- Since o is to the right of w and is not the representative of d_w , it must be the case that she is not λ -acceptable for d_w , that is, $\sum_{i \in N_{d_w}} \delta(i, o) > \lambda \sum_{i \in N_{d_w}} \delta(i, o_w)$.

Combining these two observations together with the triangle inequality, we obtain

$$\begin{aligned}
\delta(w, o) &\leq \delta(o, o_w) = \frac{1}{n_{d_w}} \sum_{i \in N_{d_w}} \delta(o, o_w) \\
&\leq \frac{1}{n_{d_w}} \sum_{i \in N_{d_w}} \left(\delta(i, o) + \delta(i, o_w) \right) \\
&\leq \left(1 + \frac{1}{\lambda} \right) \cdot \frac{1}{n_{d_w}} \sum_{i \in N_{d_w}} \delta(i, o) \\
&\leq \left(1 + \frac{1}{\lambda} \right) \cdot (\text{MAX} \circ \text{AVG})(o|I).
\end{aligned}$$

Putting everything together, we obtain a distortion of at most $2 + \frac{1}{\lambda}$.

Case 2: w is to the right of o . As in the previous case, let d^* be the district with maximum cost for w . Denote by y^* the representative of d^* , and by o^* the optimal alternative for d^* that minimizes the (average) total distance of the agents in d^* . We consider the following two subcases:

- **w is not λ -acceptable for d^* .** To argue about this case, we will need the following folklore technical lemma for instances on the line:

Lemma 4.4. *Let S be a set of agents, and denote by o_S the optimal alternative for the agents in S that minimizes their total distance. Then, for every two alternatives x and y , such that $y < x < o_S$, or $o_S < x < y$, it holds that $\sum_{i \in S} \delta(i, x) \leq \sum_{i \in S} \delta(i, y)$.*

The fact that w is the leftmost district representative implies that y^* is to the right of w . In case $o^* < w < y$, Lemma 4.4 would yield that

$$\sum_{i \in N_{d^*}} \delta(i, w) \leq \sum_{i \in N_{d^*}} \delta(i, y^*) \leq \alpha \sum_{i \in N_{d^*}} \delta(i, o^*),$$

thus contradicting that w is not λ -acceptable for d^* . Hence, $o < w < o^*$, and Lemma 4.4 implies that

$$(\text{MAX} \circ \text{AVG})(w|I) = \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, w) \leq \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, o) \leq (\text{MAX} \circ \text{AVG})(o|I),$$

meaning that w is no worse than the optimal alternative in this case.

- **w is λ -acceptable for d^* .** In this case, we clearly have that

$$\begin{aligned}
(\text{MAX} \circ \text{AVG})(w|I) &= \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, w) \\
&\leq \lambda \cdot \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, o^*) \\
&\leq \lambda \cdot \frac{1}{n_{d^*}} \sum_{i \in N_{d^*}} \delta(i, o)
\end{aligned}$$

$$\leq \lambda \cdot (\text{MAX} \circ \text{AVG})(o|I),$$

and thus the distortion is at most λ in this case.

Putting everything together, the distortion of the mechanism is at most $\max\{2 + \frac{1}{\lambda}, \lambda\}$. $\square$

By optimizing the distortion bound achieved by the class of λ -ARL mechanisms over the parameter λ , we see that the best such mechanism achieves a distortion of at most $1 + \sqrt{2}$.

Corollary 4.5. *For any $G \in \{\text{AVG}, \text{MAX}\}$, the $(\text{MAX} \circ G)$ -distortion of $(1 + \sqrt{2})$ -ARL is at most $1 + \sqrt{2}$.*

We conclude this section by presenting a lower bound of $1 + \sqrt{2}$ on the $(\text{MAX} \circ G)$ -distortion of distributed mechanisms, which holds even when the metric is a line, thus showing that $(1 + \sqrt{2})$ -ARL is the best possible on the line.

Theorem 4.6. *For any $G \in \{\text{AVG}, \text{MAX}\}$, the $(\text{MAX} \circ G)$ -distortion of any distributed mechanism is at least $1 + \sqrt{2}$, even when the metric is a line.*

Proof. Suppose towards a contradiction that there exists a distributed mechanism M with $(\text{MAX} \circ G)$ -distortion strictly smaller than $1 + \sqrt{2}$. We will consider instances with two alternatives a and b that are located at 0 and 2, respectively. In addition, our instances will consist of single-agent districts so that $\text{AVG} = \text{MAX}$ therein, and thus the distortion bounds hold for any $G \in \{\text{AVG}, \text{MAX}\}$. We can assume without loss of generality that M outputs a as the winner whenever it is given as input an instance with two districts of the same size, such that both alternatives are representative of some district.

First, consider an instance I_1 with a single district consisting of one agent that is located at $2 - \sqrt{2}$. Clearly, we have that $(\text{MAX} \circ G)(a|I_1) = 2 - \sqrt{2}$ and $(\text{MAX} \circ G)(b|I_1) = \sqrt{2}$. Therefore, M must choose a as the district representative, and thus the overall winner; otherwise its distortion would be at least $\frac{\sqrt{2}}{2 - \sqrt{2}} = 1 + \sqrt{2}$.

Next, consider an instance I_2 with a single district consisting of one agent that is located at $2 + \sqrt{2}$. Here, we have that $(\text{MAX} \circ G)(a|I_2) = 2 + \sqrt{2}$ and $(\text{MAX} \circ G)(b|I_2) = \sqrt{2}$. Therefore, M must choose b as the district representative; otherwise its distortion would be at least $\frac{2 + \sqrt{2}}{\sqrt{2}} = 1 + \sqrt{2}$.

Finally, consider an instance I_3 with the following two districts:

- The first district is similar to the one in I_1 and consists of an agent that is located at $2 - \sqrt{2}$. By the above discussion, the representative of this district is a .
- The second district is similar to the one in I_2 and consists of an agent that is located at $2 + \sqrt{2}$. By the above discussion, the representative of this district is b .

Since both alternatives are representative of some district and the districts have the same unit size, M outputs a as the final winner. However, it holds that $(\text{MAX} \circ G)(a|I_3) = 2 + \sqrt{2}$ (realized by the agent in the second district) and $(\text{MAX} \circ G)(b|I_3) = \sqrt{2}$ (realized by both agents), and thus the distortion of M is at least $\frac{2 + \sqrt{2}}{\sqrt{2}} = 1 + \sqrt{2}$, a contradiction. $\square$

5. Line metric: ordinal mechanisms

We now turn our attention to ordinal distributed mechanisms and the two objectives $\text{AVG} \circ G$ for $G \in \{\text{AVG}, \text{MAX}\}$. Recall that Corollary 3.3 implies a distortion bound of at most 11 for these objectives. However, when the metric is a line, we can do much better by observing that there is an ordinal direct over-districts voting rule with AVG -distortion of 1. In particular, we can identify the median district representative (since the input rankings are strict and the metric space is a line) and choose it as the final winner. Using the rule of Gkatzelis et al. [24] as the direct in-district voting rule, we obtain a distributed 3-in-1-over mechanism with distortion at most 7 due to Theorem 3.1.

Corollary 5.1. *When the metric is a line, there exists an ordinal distributed mechanism with $(\text{AVG} \circ G)$ -distortion at most 7, for any $G \in \{\text{AVG}, \text{MAX}\}$.*

For $\text{AVG} \circ \text{AVG}$, we will show that the bound of 7 is tight, even when the metric is a line. We first argue about the representative of particular districts that all ordinal mechanisms would have to choose to achieve distortion smaller than 7. This structural property of ordinal mechanisms will then be used to derive a lower bound on their distortion that is arbitrarily close to 7. The following lemma is similar in spirit to Lemma 4.1 in Section 4.

Lemma 5.2. *Let M be an ordinal distributed mechanism with distortion less than 7. Consider instances with two alternatives a and b , and districts of size λ . Then, for every integer $\mu \geq 1$, M must choose a as the representative of the district that consists of $\lambda \left(\frac{1}{2} + \frac{1}{4\mu}\right)$ agents who prefer a over b , and $\lambda \left(\frac{1}{2} - \frac{1}{4\mu}\right)$ agents who prefer b over a .*

Proof. We prove the statement by induction on μ . The induction depends on showing the following property (P): When given as input instances with $k = 2\mu + 1$ districts, such that a is the representative of μ districts, and b is the representative of $\mu + 1$ districts, M must choose b as the overall winner. This property will be used in the induction step to show the statement of the lemma, which in turn will be used to show property (P). To simplify our calculations and discussion, we will consider the objective SUM (the total cost of the agents in all districts) which is equivalent to $\text{AVG} \circ \text{AVG}$ in terms of distortion when the districts are symmetric (this will be the case in all instances considered in the proof).

Base case for $\mu = 1$:

First, consider an instance I_1 consisting of (multiple copies of) a single district with λ agents in total, such that $3\lambda/4$ agents prefer a over b , while the remaining $\lambda/4$ agents prefer b over a . Suppose that M choose b as the representative of this district, and thus the overall winner for I_1 . Consider the following line metric, which is consistent to the ordinal preferences of the agents:

- Alternative a is located at 0 and alternative b is located at 1.
- The $3\lambda/4$ agents that prefer a over b are positioned at 0.
- The $\lambda/4$ agents that prefer b over a are positioned at $1/2$.

Since

$$\text{SUM}(a|I_1) = \frac{\lambda}{4} \cdot \frac{1}{2} = \frac{\lambda}{8}$$

and

$$\text{SUM}(b|I_1) = \frac{\lambda}{4} \cdot \frac{1}{2} + \frac{3\lambda}{4} = \frac{7\lambda}{8},$$

the distortion of M is at least 7, a contradiction. Hence, it must be the case that a is the representative of this district (with the particular set of ordinal preferences).

We now argue about the base case of property (P), that is, when given as input instances with three districts, such that a is the representative of one district, and b is the representative of two districts, M must choose b as the overall winner. In particular, consider the instance I_2 with the following districts:

- The first district consists of $3\lambda/4$ agents that prefer a over b , and $\lambda/4$ agents that prefer b over a . By the above discussion, the representative of this district is a .
- In each of the other two districts, all λ agents prefer b over a . Due to unanimity, the representative of this district is b .
- In the other two districts, all agents are positioned at 1. Due to unanimity, the representative of both districts is 1.

Suppose towards a contradiction that M chooses a as the overall winner when given as input this instance. Consider the following line metric:

- Alternative a is located at 0 and alternative b is located at 1.
- In the first district, the $3\lambda/4$ agents that prefer a are positioned at $1/2$, while the $\lambda/4$ agents that prefer b are positioned at 1.
- In the other two districts, all agents are positioned at 1.

Since

$$\text{SUM}(a|I_2) = \frac{3\lambda}{4} \cdot \frac{1}{2} + \frac{\lambda}{4} + 2\lambda = \frac{21\lambda}{8}$$

and

$$\text{SUM}(b|I_2) = \frac{3\lambda}{4} \cdot \frac{1}{2} = \frac{3\lambda}{8},$$

the distortion of M is at least 7, a contradiction. Hence, b is chosen as the overall winner.

Induction step:

We assume that property (P) is true for $\mu = \ell - 1$. We will use this to show that the statement of the lemma is true for $\mu = \ell$, which will in turn be used to show that property (P) is true for $\mu = \ell$, thus completing the induction.

Consider an instance $I_1^{(\ell)}$ with the following $2\ell - 1$ districts:

- Each of the first ℓ districts is a copy of the district defined in the statement: There are $\lambda(\frac{1}{2} + \frac{1}{4\ell})$ agents that prefer a over b , and $\lambda(\frac{1}{2} - \frac{1}{4\ell})$ agents that prefer b over a .

- In each of the remaining $\ell - 1$ districts, there are λ agents that prefer a over b . Due to unanimity, the representative of all these districts is a .

Suppose that b is the representative of each of the first ℓ districts. Combined with the fact that a is the representative of $\ell - 1$ districts, by the induction hypothesis about property (P), we have that the overall winner must be b . Consider the following line, which is consistent to the ordinal preferences of all agents:

- Alternative a is located at 0 and alternative b is located at 1.
- In each of the first ℓ districts, the $\lambda(\frac{1}{2} + \frac{1}{4\ell})$ agents that prefer a are positioned at 0, while the $\lambda(\frac{1}{2} - \frac{1}{4\ell})$ agents that prefer b are positioned at $1/2$.
- In each of the last $\ell - 1$ districts, all agents are positioned at 0.

Since

$$\text{SUM}(a|I_1^{(\ell)}) = \ell \cdot \lambda \left(\frac{1}{2} - \frac{1}{4\ell} \right) \cdot \frac{1}{2} = \frac{(2\ell - 1)\lambda}{8}$$

and

$$\text{SUM}(b|I_1^{(\ell)}) = \ell \left(\frac{\ell\lambda}{2} + \frac{\lambda}{4\ell} \right) + \ell \cdot \left(\frac{\lambda}{2} - \frac{\lambda}{4\ell} \right) \cdot \frac{1}{2} + (\ell - 1) \cdot \lambda = \frac{7(2\ell - 1)\lambda}{8},$$

choosing b at the overall winner leads to a distortion of at least 7, a contradiction. Hence, it must be the case that the representative of each of the first ℓ districts is a (with the particular set of ordinal preferences).

To complete the proof, we need to also show that property (P) holds (so that we can use it in the above argument about instance $I_1^{(\ell)}$). Consider the instance $I_2^{(\ell)}$ with the following $2\ell + 1$ districts:

- Each of the first ℓ districts is a copy of the district defined in the statement: There are $\lambda(\frac{1}{2} + \frac{1}{4\ell})$ agents that prefer a over b , and $\lambda(\frac{1}{2} - \frac{1}{4\ell})$ agents that prefer b over a . By the above discussion (about instance $I_1^{(\ell)}$), the representative of these districts is a .
- In each of the remaining $\ell + 1$ districts, there are λ agents that prefer b over a . Due to unanimity, the representative of these districts is b .

Suppose towards a contradiction that M chooses a as the overall winner, and consider the following line metric, which is consistent to the ordinal preferences of the agents:

- Alternative a is located at 0 and alternative b is located at 1.
- In each of the first ℓ districts, the $\lambda(\frac{1}{2} + \frac{1}{4\ell})$ agents that prefer a are positioned at $1/2$, while the $\lambda(\frac{1}{2} - \frac{1}{4\ell})$ agents that prefer b are positioned at 1.
- In each of the last $\ell + 1$ districts, all agents are positioned at 1.

Since

$$\text{SUM}(a|I_2^{(\ell)}) = \ell \left(\frac{\lambda}{2} + \frac{\lambda}{4\ell} \right) \cdot \frac{1}{2} + \ell \cdot \left(\frac{\lambda}{2} - \frac{\lambda}{4\ell} \right) + (\ell + 1) \cdot \lambda = \frac{7(2\ell + 1)\lambda}{8},$$

and

$$\text{SUM}(b|I_2^{(\ell)}) = \ell \cdot \left(\frac{\lambda}{2} + \frac{\lambda}{4\ell} \right) \cdot \frac{1}{2} = \frac{(2\ell + 1)\lambda}{8},$$

choosing a as the overall winner leads to a distortion of 7, a contradiction. So, the overall winner must be b , and the proof is now complete. $\square$

We are now ready to show the lower bound on the $\text{AVG} \circ \text{AVG}$ -distortion of ordinal mechanisms.

Theorem 5.3. *The $(\text{AVG} \circ \text{AVG})$ -distortion of any ordinal distributed mechanism is at least $7 - \varepsilon$, for any $\varepsilon > 0$, even when the metric is a line.*

Proof. Let $\mu \geq 1$ be some integers. Consider any ordinal distributed mechanism M , which is given as input an instance I with two alternatives a and b , and the following two districts:

- In the first district, there are $\lambda(\frac{1}{2} + \frac{1}{4\mu})$ agents that prefer a over b , and $\lambda(\frac{1}{2} - \frac{1}{4\mu})$ agents that prefer b over a . By Lemma 5.2, the representative of this district is a .

- In the second district, all λ agents prefer b over a . Due to unanimity, the representative of this district is b .

So, we have an instance with a as the representative of one district, and b as the representative of the other district. In this case, recall that it is without loss of generality to assume that M chooses a as the overall winner. Consider the following line metric, which is consistent to the ordinal preferences of the agents:

- In the first district, the $\lambda \left(\frac{1}{2} + \frac{1}{4\mu} \right)$ agents that prefer a are positioned at $1/2$, while the $\lambda \left(\frac{1}{2} - \frac{1}{4\mu} \right)$ agents that prefer b are positioned at 1 .
- In the second district, all agents are positioned at 1 .

Since

$$\text{SUM}(a|I) = \frac{1}{2} \cdot \lambda \left(\frac{1}{2} + \frac{1}{4\mu} \right) + \lambda \left(\frac{1}{2} - \frac{1}{4\mu} \right) + \lambda = \frac{\lambda}{4} \left(7 - \frac{1}{2\mu} \right)$$

and

$$\text{SUM}(b|I) = \frac{1}{2} \cdot \lambda \left(\frac{1}{2} + \frac{1}{4\mu} \right) = \frac{\lambda}{4} \left(1 + \frac{1}{2\mu} \right),$$

choosing a as the overall winner leads to a distortion of at least

$$\frac{7 - 1/(2\mu)}{1 + 1/(2\mu)} = 7 - \frac{4}{\mu + 2}.$$

The lower bound of $7 - \varepsilon$ follows by choosing a sufficiently large value for μ . $\square$

For $\text{AVG} \circ \text{MAX}$, the bound of 7 implied by Corollary 5.1 is a first improvement, but we can do even better with the following distributed mechanism.

ARBITRARY-MEDIAN

1. For every district $d \in D$, choose its representative to be the favorite alternative of an arbitrary agent $j_d \in N_d$.
2. Output the median representative as the final winner.

Theorem 5.4. *When the metric is a line, ARBITRARY-MEDIAN has $(\text{AVG} \circ \text{MAX})$ -distortion at most 5.*

Proof. Consider any instance $I = (N, A, D, \delta)$, in which δ is a line metric. Let w be the alternative chosen by ARBITRARY-MEDIAN when given as input (the ordinal information of) I , and denote by o the optimal alternative. For every district $d \in D$, let $i_d = \arg \max_{i \in N_d} \delta(i, w)$ be the agent in d that gives the max cost for w . By the triangle inequality, and since $(\text{AVG} \circ \text{MAX})(o|I) \geq \frac{1}{k} \sum_{d \in D} \delta(i_d, o)$, we have

$$\begin{aligned} (\text{AVG} \circ \text{MAX})(w|I) &= \frac{1}{k} \sum_{d \in D} \delta(i_d, w) \\ &\leq \frac{1}{k} \sum_{d \in D} \left(\delta(i_d, o) + \delta(w, o) \right) \\ &\leq (\text{AVG} \circ \text{MAX})(o|I) + \delta(w, o). \end{aligned}$$

Without loss of generality, we can assume that w is to the left of o on the line. Since w is the median representative, there is a set S of at least $k/2$ districts whose representatives are to the left of (or coincide with) w . As a result, for every district $d \in S$, agent j_d (whose favorite alternative becomes the representative of d) prefers w over o , and thus $d(j_d, o) \geq \delta(w, o)/2$. Using this, we obtain

$$\begin{aligned} (\text{AVG} \circ \text{MAX})(o|I) &= \frac{1}{k} \sum_{d \in D} \max_{i \in N_d} \delta(i, o) \geq \frac{1}{k} \sum_{d \in S} \delta(j_d, o) \geq \frac{1}{2} \cdot \frac{\delta(w, o)}{2} \\ &\Leftrightarrow \delta(w, o) \leq 4 \cdot (\text{AVG} \circ \text{MAX})(o|I). \end{aligned}$$

Putting everything together, we obtain a distortion upper bound of 5. $\square$

Our next theorem shows an almost matching lower bound of approximately 4.23.

Theorem 5.5. *The $(\text{AVG} \circ \text{MAX})$ -distortion of any ordinal distributed mechanism is at least $2 + \sqrt{5} - \varepsilon$, for any $\varepsilon > 0$, even when the metric is a line.*

Proof. Suppose towards a contradiction that there is an ordinal distributed mechanism M with distortion strictly smaller than $2 + \sqrt{5} - \varepsilon$, for any $\varepsilon > 0$. We will reach a contradiction by defining instances with two alternatives a and b , and districts consisting of the same size. Without loss of generality, we assume that M chooses alternative a as the final winner when given as input any instance with only two districts, such that both alternatives are representative of some district. Let x and y be two integers such that $\phi > y/x \geq \phi - \varepsilon/2$, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio.

First, consider an instance I_1 consisting of the following two districts:

- The first district consists of two agents, such that one of them prefers alternative a and the other prefers alternative b .
- The second district consists of two agents, such that both of them prefer alternative b . Due to unanimity, the representative of this district must be b .

Suppose that M chooses a as the representative of the first district, in which case both alternatives are representative of some district, and thus M chooses a as the final winner. Consider the following metric:

- Alternative a is positioned at 0, and alternative b is positioned at 1.
- In the first district, the agent that prefers alternative a is positioned at $1/2$, whereas the agent that prefers alternative b is positioned at $3/2$.
- In the second district, both agents are positioned at 1.

Then, we have that $(\text{AVG} \circ \text{MAX})(a|I_1) = \frac{1}{2} \left(\frac{3}{2} + 1 \right) = 5/4$ and $(\text{AVG} \circ \text{MAX})(b|I_1) = \frac{1}{2} \left(\frac{1}{2} + 0 \right) = 1/4$, leading to a distortion of 5. Consequently, M must choose b as the representative of the first district (where one agent prefers a and one prefers b).

Next, we argue that for instances with $x + y$ districts such that a is the representative of x districts and b is the representative of y districts, M must choose b as the final winner. Assume otherwise that M chooses a in such a situation, and consider an instance I_2 consisting of the following $x + y$ districts:

- Each of the first x districts consists of a single agent that prefers alternative a ; thus, a is the representative of all these districts.
- Each of the next y districts consists of a single agent that prefers alternative b ; thus, b is the representative of all these districts.

By assumption, M chooses a as the final winner. Consider the following metric:

- Alternative a is positioned at 0 and alternative b is positioned at 1.
- In each of the first x districts, the agent therein is positioned at $1/2$.
- In each of the next y districts, the agent therein is positioned at 1.

Since

$$(\text{AVG} \circ \text{MAX})(a|I_2) = \frac{1}{x+y} \left(\frac{x}{2} + y \right) = \frac{1}{x+y} \cdot \frac{x+2y}{2}$$

and

$$(\text{AVG} \circ \text{MAX})(b|I_2) = \frac{1}{x+y} \left(\frac{x}{2} + 0 \right) = \frac{1}{x+y} \cdot \frac{x}{2},$$

the distortion is

$$\frac{x+2y}{x} = 1 + 2\frac{y}{x} \geq 1 + 2\phi - \varepsilon = 2 + \sqrt{5} - \varepsilon.$$

Consequently, M must choose b , whenever there are $x + y$ districts such that a is the representative of x districts and b is the representative of the remaining y districts.

Finally, consider an instance I_3 with the following $x + y$ districts:

- Each of the first x districts consists of two agents that prefer alternative a . Due to unanimity, a must be the representative of all these districts.
- Each of the next y districts consists of two agents, such that one of them prefers alternative a , while the other prefers alternative b . By the discussion above (about instance I_1), b must be the representative of these districts.

Since a is the representative of x districts and b is the representative of y districts, by the discussion above (about instance I_2), M chooses b as the final winner. Consider the following metric:

- Alternative a is positioned at 0 and alternative b is positioned at 1.
- In each of the first x districts, the two agents therein are positioned at 0.
- In each of the next y districts, the agent that prefers a is positioned at $-1/2$, whereas the agent that prefers b is positioned at $1/2$.

Since

$$(\text{AVG} \circ \text{MAX})(a|I_3) = \frac{1}{x+y} \left(0 + \frac{y}{2} \right) = \frac{1}{x+y} \cdot \frac{y}{2}$$

and

$$(\text{AVG} \circ \text{MAX})(b|I_3) = \frac{1}{x+y} \left(x + \frac{3y}{2} \right) = \frac{1}{x+y} \cdot \frac{2x+3y}{2},$$

the distortion is

$$\frac{2x+3y}{y} = 3 + \frac{2x}{y} > 3 + \frac{2}{\phi} = 1 + 2\phi = 2 + \sqrt{5}.$$

This contradicts our assumption that M has distortion strictly smaller than $2 + \sqrt{5} - \varepsilon$, thus completing the proof. $\square$

Next, we consider the objectives $\text{MAX} \circ G$ for $G \in \{\text{AVG}, \text{MAX}\}$. Corollary 3.5 implies a distortion bound of at most 5 for these objectives. When $G = \text{MAX}$ and the metric is a line, we get an improved upper bound of 3 using a rather simple distributed mechanism, which essentially outputs the favorite alternative of an arbitrary agent.

ARBITRARY-DICTATOR

1. For each district $d \in D$, choose its representative to be the favorite alternative of an arbitrary agent in N_d .
2. Output an arbitrary district representative as the final winner.

Theorem 5.6. *When the metric is a line, ARBITRARY-DICTATOR has $(\text{MAX} \circ \text{MAX})$ -distortion at most 3.*

Proof. Let $I = (N, A, D, \delta)$ be an arbitrary instance, with δ a line metric. Without loss of generality, we assume that the alternative w chosen by ARBITRARY-DICTATOR is positioned to the right of the dictator agent i^* . In addition, we denote by o the optimal alternative, by ℓ the leftmost agent, and by r the rightmost agent. With some abuse of notation, in the following we will also use w , i^* , o , ℓ and r to denote the positions of the corresponding agents and alternatives on the line.

First observe that the distortion is at most 2 when $w \in [\ell, r]$. In particular, since there is at least one alternative in-between the agents, we have that $(\text{MAX} \circ \text{MAX})(o|I) \geq \frac{\delta(\ell, r)}{2}$ and $(\text{MAX} \circ \text{MAX})(w|I) \leq \delta(\ell, r)$. Hence, we can assume that $w > r$ (since $w \geq i^*$). Furthermore, it cannot be the case that $i^* \leq o \leq w$ (since then i^* would prefer o over w), neither that $r \leq w \leq o$ (since then o would not be the optimal alternative). Consequently, $o < i^* \leq r < w$. Since w is the favorite alternative of i^* , w is also the favorite alternative of r . Hence, $\delta(r, w) \leq \delta(r, o)$. We distinguish between the following two subcases:

- If $o \in [\ell, r]$, we clearly have that $(\text{MAX} \circ \text{MAX})(o|I) \geq \frac{\delta(\ell, r)}{2}$, and

$$(\text{MAX} \circ \text{MAX})(w|I) = \delta(\ell, w) = \delta(\ell, r) + \delta(r, w) \leq 3 \cdot (\text{MAX} \circ \text{MAX})(o|I).$$

- If $o < \ell$, we have that

$$(\text{MAX} \circ \text{MAX})(o|I) = \delta(r, o) = \delta(\ell, o) + \delta(\ell, r) \geq \delta(\ell, r)$$

and

$$(\text{MAX} \circ \text{MAX})(w|I) = \delta(\ell, w) = \delta(r, w) + \delta(\ell, r) \leq \delta(r, o) + \delta(\ell, r) \leq 2 \cdot (\text{MAX} \circ \text{MAX})(o|I).$$

This completes the proof. $\square$

The following theorem shows that the bound of 3 is the best possible we can hope for the $\text{MAX} \circ \text{MAX}$ objective using a unanimous distributed mechanism, even when the metric is a line. This lower bound directly extends to ordinal mechanisms, as any such mechanism with finite distortion has to be unanimous (Lemma 2.1).

Theorem 5.7. *The $(\text{MAX} \circ \text{MAX})$ -distortion of any unanimous distributed mechanism is at least 3, even when the metric is a line.*

Proof. Consider an arbitrary unanimous distributed mechanism M , and the following instance I with two alternatives a and b . There are two districts with $x \geq 1$ agents each: All agents in the first district prefer alternative a , whereas all agents in the second district prefer b . Due to unanimity, a must be the representative of the first district, and b the representative of the second district. The final winner can be any of the two, say b without loss of generality. Now, consider the following metric:

- Alternative a is located at 1 and alternative b is located at 3.
- All agents in the first district (who prefer a) are located at 0.
- All agents in the second district (who prefer b) are located at 2.

Consequently, we have that $(\text{MAX} \circ \text{MAX})(a|I) = 1$ and $(\text{MAX} \circ \text{MAX})(b|I) = 3$, leading to a distortion of 3. $\square$

Finally, let us focus on the objective $\text{MAX} \circ \text{AVG}$, for which we show a lower bound of approximately 4.23 on the distortion of all ordinal distributed mechanisms, thus almost matching the upper bound of 5 implied by Corollary 3.5.

Theorem 5.8. *The $(\text{MAX} \circ \text{AVG})$ -distortion of any ordinal distributed mechanism is at least $2 + \sqrt{5} - \varepsilon$, for any $\varepsilon > 0$, even when the metric is a line.*

Proof. Suppose towards a contradiction that there is an ordinal distributed mechanism M with $(\text{MAX} \circ \text{AVG})$ -distortion strictly smaller than $2 + \sqrt{5} - \varepsilon$, for any $\varepsilon > 0$. We will consider instances with two alternatives a and b . We can assume without loss of generality that M outputs a as the winner whenever it is given as input an instance with two districts of the same size, such that both alternatives are representative of some district. However, we do not know how M decides the representative of a district in case there is a tie between the two alternatives therein. Let $\theta = \frac{\sqrt{5}-1}{2} = \phi - 1 \approx 0.618$, and denote by x a sufficiently large integer, such that θx and $(1 - \theta)x$ are both integers.⁴ Observe that θ is such that $\frac{1+\theta}{1-\theta} = \frac{2+\theta}{\theta} = 1 + 2\phi = 2 + \sqrt{5}$.

First, consider an instance I_1 with a single district consisting of x agents, such that θx agents prefer a , and $(1 - \theta)x$ agents prefer b . Suppose that M outputs b (as the district representative and the final winner) when given these ordinal preferences of the agents in the district as input, and consider the following metric:

- Alternative a is located at 0 and alternative b is located at 1.
- The θx agents that prefer a are located at 0.
- The $(1 - \theta)x$ agents that prefer b are located at $1/2$.

Since there is only one district, we clearly have that

$$(\text{MAX} \circ \text{AVG})(a|I_1) = \frac{1}{x} \left(0 + \frac{(1 - \theta)x}{2} \right) = \frac{1}{x} \cdot \frac{(1 - \theta)x}{2}$$

and

$$(\text{MAX} \circ \text{AVG})(b|I_1) = \frac{1}{x} \left(\theta x + \frac{(1 - \theta)x}{2} \right) = \frac{1}{x} \cdot \frac{(1 + \theta)x}{2}.$$

Hence, the distortion in this case is $\frac{1+\theta}{1-\theta} = 2 + \sqrt{5}$. As this contradicts our assumption that M has distortion strictly smaller than $2 + \sqrt{5}$, M must choose a as the winner in such a single-district instance.

Now, consider an instance I_2 with the following two districts:

- The first district is similar to the one in instance I_1 , i.e., it consists of θx agents that prefer a , and $(1 - \theta)x$ agents that prefer b . By the above discussion, the representative of this district must be alternative a .
- The second district consists of x agents, all of whom prefer b . Due to unanimity (Lemma 2.1), the representative of this district must be alternative b .

Since both a and b are representative of some district, M outputs a as the final winner. Now, consider the following metric:

- Alternative a is located at 0 and alternative b is located at 1.

⁴ Here, we assume the existence of such an x to simplify the presentation of the proof. Since θ is irrational, to be exact, we can choose x to be a sufficiently large integer so that θx and $(1 - \theta)x$ are approximately equal to their floors, which would lead to the distortion lower bound being $2 + \sqrt{5} - \varepsilon$, for any $\varepsilon > 0$.

- In the first district, the θx agents that prefer a are located at $1/2$, and the remaining $(1 - \theta)x$ agents that prefer b are located at 1 . The total distance of the agents in the district is $\frac{\theta x}{2} + (1 - \theta)x = \frac{(2 - \theta)x}{2}$ from a , and $\frac{\theta x}{2}$ from b .
- In the second district, all x agents are located at $1 + \frac{\theta}{2}$. The total distance of the agents in the district is $\frac{(2 + \theta)x}{2}$ from a , and $\frac{\theta x}{2}$ from b .

Consequently, from the second district, we have that

$$(\text{MAX} \circ \text{AVG})(a|I_2) = \frac{1}{x} \cdot \frac{(2 + \theta)x}{2}$$

and, from both districts, we have that

$$(\text{MAX} \circ \text{AVG})(b|I_2) = \frac{1}{x} \cdot \frac{\theta x}{2}.$$

Thus, the distortion of the mechanism for this instance is $\frac{2 + \theta}{\theta} = 2 + \sqrt{5}$, a contradiction. $\square$

6. Extensions and generalizations

6.1. Mechanisms that select from the set of representatives

In the previous sections we looked at distributed mechanisms, which can choose any alternative as the final winner by essentially considering the district representatives as proxies for the preferences of their districts. We now focus on the case where the final winner can only be chosen from among the district representatives. To make the distinction between general mechanisms and those that select from the pool of district representatives clear, we will use the term *representative-selecting* to refer to the latter.

It is not hard to see that, with the exception of the bounds implied by Theorem 3.1 and its corollaries for general metric spaces (Corollaries 3.2 and 3.3), the rest of our results follow from representative-selecting mechanisms. In particular:

- Every α -in-arbitrary-over mechanism, as well as ARBITRARY-DICTATOR, choose some arbitrary representative (Theorem 3.4, Corollary 3.5, and Theorem 5.6);
- The 1-in-1-over cardinal mechanism and the 3-in-1-over ordinal mechanism for $\text{AVG} \circ G$ in the line metric, as well as ARBITRARY-MEDIAN, choose the median representative (Corollary 3.2 for line metric, Corollary 5.1, and Theorem 5.4);
- Every λ -ARL mechanism chooses the leftmost representative (Theorem 4.3 and Corollary 4.5).

It is also not hard to observe that all our lower bounds (Theorems 4.6, 5.5, 5.7, and 5.8) also extend for the class of representative-selecting mechanisms: some representative is always chosen as the final winner in all instances used in the constructions. Based on the above discussion, we have the following corollary, which collects the best distortion bounds for the different objectives we consider.

Theorem 6.1. *We can form ordinal representative-selecting mechanisms with distortion at most 5 for $\text{MAX} \circ G$ and general metric spaces. When the metric is a line, the distortion of cardinal representative-selecting mechanisms is exactly 3 for $\text{AVG} \circ G$, and exactly $1 + \sqrt{2}$ for $\text{MAX} \circ G$. When the metric space is a line, the worst-case distortion of ordinal mechanisms is exactly 7 for $\text{AVG} \circ \text{AVG}$, between $2 + \sqrt{5}$ and 5 for $\text{AVG} \circ \text{MAX}$, exactly 3 for $\text{MAX} \circ \text{MAX}$, and at least $2 + \sqrt{5}$ for $\text{MAX} \circ \text{AVG}$.*

Now, let us see how choosing only from the district representatives affects the bounds implied by Theorem 3.1 for general metric spaces. For clarity, we focus on objectives of the form $\text{AVG} \circ G$; our discussion can easily be adapted for objectives of the form $\text{MAX} \circ G$. Let M be a representative-selecting mechanism that uses some in-district and over-districts direct voting rules. Given an instance $I = (N, M, D, \delta)$, let $R = R_M(I)$ be the set of district representatives chosen by M , and denote by $w = M(I) \in R$ the final winner. Clearly, Theorem 3.1 would hold without any modifications if, for any instance I , w satisfies inequality (2) in the proof of the theorem:

$$\forall j \in A : \sum_{i \in R} \delta(i, w) \leq \beta \cdot \sum_{i \in R} \delta(i, j).$$

However, the distortion guarantees of direct voting rules used by M in and over the districts are usually only with respect to the set of alternatives from which they are allowed to choose. So, if M uses an over-districts rule that has AVG -distortion at most γ , we have that, for any instance I , w is such that

$$\forall j \in R : \sum_{i \in R} \delta(i, w) \leq \gamma \cdot \sum_{i \in R} \delta(i, j). \quad (3)$$

Inequality (3) cannot directly substitute inequality (2) in the proof of Theorem 3.1 as it may be the case that the optimal alternative is not included in the set of representatives. So, we need to understand the relation between β and γ , and then use it to obtain a distortion bound for M .

Lemma 6.2. *For any mechanism M , it holds that $\beta \leq 2\gamma$.*

Proof. By definition, β is the worst-case ratio when comparing the total distance of the representatives from the outcome of M to their total distance from the optimal *alternative*, whereas γ is the worst-case ratio when comparing the total distance of the representatives from the outcome of M to their total distance from the optimal *representative*. So, to prove the statement, it suffices to show that

$$\min_{j \in R} \sum_{i \in R} \delta(i, j) \leq 2 \cdot \min_{j \in A} \sum_{i \in R} \delta(i, j),$$

for every instance that results in the set of district representatives R . This inequality holds trivially when there is only one district, so we assume that $k \geq 2$. To simplify notation, denote by $r \in R$ the optimal representative and by $o \in A$ the optimal alternative. For every $j \in R$, we have $\sum_{i \in R} \delta(i, r) \leq \sum_{i \in R} \delta(i, j)$. So, summing over all representatives, and using the triangle inequality, we obtain

$$\begin{aligned} \sum_{i \in R} \delta(i, r) &\leq \frac{1}{k} \sum_{j \in R} \sum_{i \in R} \delta(i, j) \\ &= \frac{2}{k} \sum_{i, j \in R} \delta(i, j) \\ &\leq \frac{2}{k} \sum_{i, j \in R} (\delta(i, o) + \delta(j, o)) \\ &= \frac{2(k-1)}{k} \sum_{i \in R} \delta(i, o). \end{aligned}$$

Thus, we have that $\beta \leq \frac{2(k-1)}{k} \gamma \leq 2\gamma$, and the proof is complete. $\square$

Due to Lemma 6.2, Theorem 3.1 implies the following distortion bounds for general metric spaces and α -in- γ -over representative-selecting mechanisms.

Theorem 6.3. *For general metric spaces and any $F, G \in \{\text{AVG}, \text{MAX}\}$, the $(F \circ G)$ -distortion of any α -in- γ -over representative-selecting mechanism is at most $\alpha + 2\gamma + 2\alpha\gamma$.*

Using appropriate direct voting rules in and over the districts, we can now again obtain concrete upper bounds on the distortion of cardinal and ordinal representative-selecting mechanisms. In particular, for every objective $F \circ G$, similarly to the case of general mechanisms in Section 3, there is a 1-in-1-over cardinal mechanism and a 3-in-2-over ordinal mechanism. Therefore, by Theorem 6.3, we obtain the following corollary; note that Theorem 6.1 implies better ordinal bounds for objectives of the form $\text{MAX} \circ G$. This corollary shows that some loss in distortion may be experienced by forcing mechanisms to select a winner from only among the set of representatives, but that loss is not too large.

Corollary 6.4. *For general metric spaces and any $F, G \in \{\text{AVG}, \text{MAX}\}$, there is a representative-selecting mechanism with $(F \circ G)$ -distortion at most 5, and an ordinal representative-selecting mechanism with $(F \circ G)$ -distortion at most 19.*

6.2. More general objectives

We now consider again mechanisms that can choose the final winner from the set of all alternatives, and discuss how some of our results (in particular, Theorems 3.1 and 4.3) can be extended for objectives $F \circ G$ beyond the cases where $F, G \in \{\text{AVG}, \text{MAX}\}$.

6.2.1. Generalizing Theorem 3.1

We previously showed in Theorem 3.1 that the $(F \circ G)$ -distortion of distributed mechanisms can be bounded in terms of the F - and G -distortion of the voting rules used in and over the districts. Here, we show that this theorem still holds for a much more general class of functions. To define this properly, we should think of F and G as functions that take as

input vectors of distances. More precisely, given an instance $I = (N, A, D, \delta)$, let f and g be functions so that the cost of any alternative $j \in A$ is

$$(F \circ G)(j|I) = f\left(g\left(\vec{\delta}_1(j)\right), \dots, g\left(\vec{\delta}_k(j)\right)\right),$$

where $\vec{\delta}_d(j)$ is the vector consisting of the distances $\delta(i, j)$ between every agent $i \in N_d$ and alternative j . To give a few examples, $g\left(\vec{\delta}_d(j)\right) = \frac{1}{n_d} \sum_{i \in N_d} \delta(i, j)$ if $G = \text{AVG}$, and $g\left(\vec{\delta}_d(j)\right) = \max_{i \in N_d} \delta(i, j)$ if $G = \text{MAX}$. More generally, we consider functions f and g which satisfy the following properties:

- **Monotonicity:** A function f is *monotone* if $f(\vec{v}) \leq f(\vec{u})$, for any two vectors $\vec{v}$ and $\vec{u}$ such that $v_\ell \leq u_\ell$ for every index ℓ .
- **Subadditivity:** A function f is *subadditive* if $f(\vec{v} + \vec{u}) \leq f(\vec{v}) + f(\vec{u})$, for any two vectors $\vec{v}$ and $\vec{u}$. Moreover, for any scalar $c \geq 1$, it must be that $f(c \cdot \vec{v}) \leq c \cdot f(\vec{v})$.⁵
- **Consistency:** A function f is *consistent* if $f(\vec{v}) = c$, for any vector $\vec{v}$ such that $v_\ell = c$ for every index ℓ .

Note that both AVG and MAX , as well as many other functions, obey all of the above properties. To justify the necessity of these properties for obtaining reasonable bounds on the distortion, note that non-monotone functions would not make sense for our setting, since choosing an outcome that is worse from the agents' perspective would improve the objective function. Subadditivity is needed to enable us to apply the triangle inequality in the analysis. In fact, it is not difficult to see that there are simple instances for which non-subadditive objective functions could result in unbounded distortion. Finally, consistency is needed to normalize the objective function values; it can be relaxed somewhat (e.g., to include the SUM objective instead of AVG) with minor changes to the proofs.

Theorem 6.5. *The distortion of any α -in- β -over mechanism is at most $\alpha + \beta + \alpha\beta$, for any objective $F \circ G$ defined by functions f and g which are monotone, subadditive, and consistent.*

Proof. Consider an arbitrary α -in- β -over mechanism M and an arbitrary instance $I = (N, A, D, \delta)$. Denote by y_d the representative of each district $d \in D$, by w the final winner, and by o the optimal alternative. In addition, for each $d \in D$, let $\vec{t}_d(j)$ be a vector of size n_d with every element being equal to the distance $\delta(y_d, j)$. Then, the cost of w is

$$\begin{aligned} (F \circ G)(w|I) &= f\left(g\left(\vec{\delta}_1(w)\right), \dots, g\left(\vec{\delta}_k(w)\right)\right) \\ &\leq f\left(g\left(\vec{\delta}_1(y_1) + \vec{t}_1(w)\right), \dots, g\left(\vec{\delta}_k(y_k) + \vec{t}_k(w)\right)\right) \\ &\leq f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) + f\left(g\left(\vec{t}_1(w)\right), \dots, g\left(\vec{t}_k(w)\right)\right) \\ &= f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) + f\left(\delta(y_1, w), \dots, \delta(y_k, w)\right), \end{aligned} \quad (4)$$

where the first inequality follows by the monotonicity of f and g together with the triangle inequality, the second inequality is due to subadditivity, and the last equality is due to the consistency of g .

Since M is α -in- β -over, we have that the cost of w if the representatives were pseudo-agents is at most a factor β worse than the cost of the optimal alternative o . That is,

$$\begin{aligned} f\left(\delta(y_1, w), \dots, \delta(y_k, w)\right) &\leq \beta \cdot f\left(\delta(y_1, o), \dots, \delta(y_k, o)\right) \\ &= \beta \cdot f\left(g\left(\vec{t}_1(o)\right), \dots, g\left(\vec{t}_k(o)\right)\right). \end{aligned}$$

Since $\delta(y_d, o) \leq \delta(y_d, i) + \delta(i, o)$ for each agent $i \in N_d$, and the fact that the functions are monotone and subadditive, we obtain

$$\begin{aligned} f\left(g\left(\vec{t}_1(o)\right), \dots, g\left(\vec{t}_k(o)\right)\right) &\leq f\left(g\left(\vec{\delta}_1(y_1) + \vec{\delta}_1(o)\right), \dots, g\left(\vec{\delta}_k(y_k) + \vec{\delta}_k(o)\right)\right) \\ &\leq f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) + f\left(g\left(\vec{\delta}_1(o)\right), \dots, g\left(\vec{\delta}_k(o)\right)\right). \end{aligned}$$

⁵ The latter condition, sometimes known as sub-homogeneity, is not usually included in the standard definition of subadditive functions. It is easily implied by the first subadditivity condition when c is an integer.

Using the above inequalities, (4) yields

$$\begin{aligned} (F \circ G)(w|I) &\leq (1 + \beta) \cdot f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) + \beta \cdot f\left(g\left(\vec{\delta}_1(o)\right), \dots, g\left(\vec{\delta}_k(o)\right)\right) \\ &= (1 + \beta) \cdot f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) + \beta \cdot (F \circ G)(o|I). \end{aligned} \quad (5)$$

Since the in-district voting rule used by M has G -distortion at most α , we have that, for each district d , the cost of y_d is at most α times the cost of o for the agents in d , that is, $g\left(\vec{\delta}_d(y_d)\right) \leq \alpha \cdot g\left(\vec{\delta}_d(o)\right)$. Combining this with the second subadditivity property of f , we obtain

$$f\left(g\left(\vec{\delta}_1(y_1)\right), \dots, g\left(\vec{\delta}_k(y_k)\right)\right) \leq \alpha \cdot f\left(g\left(\vec{\delta}_1(o)\right), \dots, g\left(\vec{\delta}_k(o)\right)\right) = \alpha \cdot (F \circ G)(o|I).$$

By this, inequality (5) now gives the desired upper bound of $\alpha + \beta + \alpha\beta$ on the distortion of M . $\square$

6.2.2. Generalizing Theorem 4.3

In this section, we will show that, when the metric space is a line, Theorem 4.3 holds for more general objectives of the form $\text{MAX} \circ G$. In particular, we are aiming to minimize the maximum cost of any district, which for an alternative $j \in A$ is given by some function $g\left(\vec{\delta}_d(j)\right)$ of the vector $\vec{\delta}_d(j)$ containing the distances between the members of d and j . As in Section 6.2.1, we will assume that g is monotone, subadditive, and consistent. In addition, we require that g is *single-peaked*, that is, for any district d , there is a unique alternative j that minimizes $g\left(\vec{\delta}_d(j)\right)$, and g increases monotonically as we move further away from the location of j (to the left or the right). It is easy to see that many functions, including AVG and MAX , obey these properties.

As in Section 4, the upper bound on the distortion is due to the λ -ARL mechanism, which chooses the representative of each district to be the rightmost λ -acceptable alternative for the district, and then outputs the leftmost representative as the final winner. Recall that the set of λ -acceptable alternatives for a district d contains all the alternatives x such that $g\left(\vec{\delta}_d(x)\right) \leq \lambda \cdot \min_{j \in A} g\left(\vec{\delta}_d(j)\right)$.

Theorem 6.6. *The distortion of λ -ARL is at most $\max\left\{2 + \frac{1}{\lambda}, \lambda\right\}$ for any objective of the form $\text{MAX} \circ G$, where G is implemented by a monotone, subadditive, consistent, and single-peaked function g .*

Proof. Consider an arbitrary instance $I = (N, A, D, \delta)$, where δ is a line metric. Let w be the alternative chosen by λ -ARL when given I as input, and denote by o the optimal alternative. We consider the following two cases depending on the relative positions of w and o .

Case 1: w is to the left of o . Let $d^* \in \arg \max_{d \in D} g\left(\vec{\delta}_d(w)\right)$ be the district that gives the max cost for w . Also, for every $d \in D$, let $\vec{\delta}_d(w, o)$ be a vector of size n_d with all its entries equal to $\delta(w, o)$. By the triangle inequality, the monotonicity, subadditivity and consistency of g , as well as the fact that $(\text{MAX} \circ G)(o|I) \geq g\left(\vec{\delta}_d(o)\right)$ for any district d , we have

$$\begin{aligned} (\text{MAX} \circ G)(w|I) &= g\left(\vec{\delta}_{d^*}(w)\right) \leq g\left(\vec{\delta}_{d^*}(o) + \vec{\delta}_{d^*}(w, o)\right) \\ &\leq g\left(\vec{\delta}_{d^*}(o)\right) + g\left(\vec{\delta}_{d^*}(w, o)\right) \\ &\leq (\text{MAX} \circ G)(o|I) + \delta(w, o). \end{aligned}$$

Now, let d_w be a district whose representative is w . In other words, w is the rightmost λ -acceptable alternative for d_w . Let o_w be the alternative with minimum cost according to g for the agents in d_w . We make the following two observations:

- Since o_w is trivially λ -acceptable, it must be the case that she is located to the left of w , and thus $\delta(w, o) \leq \delta(o_w, o)$.
- Since o is to the right of w and is not the representative of d_w , it must be the case that she is not λ -acceptable for d_w , that is, $g\left(\vec{\delta}_{d_w}(o)\right) > \lambda \cdot g\left(\vec{\delta}_{d_w}(o_w)\right)$.

Combining these two observations together with the triangle inequality and the properties of g (monotonicity, subadditivity, and consistency), we obtain

$$\begin{aligned} \delta(w, o) &\leq \delta(o, o_w) = g\left(\vec{\delta}_{d_w}(o, o_w)\right) \\ &\leq g\left(\vec{\delta}_{d_w}(o) + \vec{\delta}_{d_w}(o_w)\right) \end{aligned}$$

$$\begin{aligned}
&\leq g(\vec{\delta}_{d_w}(o)) + g(\vec{\delta}_{d_w}(o_w)) \\
&\leq \left(1 + \frac{1}{\lambda}\right) \cdot g(\vec{\delta}_{d_w}(o)) \\
&\leq \left(1 + \frac{1}{\lambda}\right) \cdot (\text{MAX} \circ G)(o|I).
\end{aligned}$$

Putting everything together, we obtain a distortion of at most $2 + \frac{1}{\lambda}$.

Case 2: w is to the right of o . As in the previous case, let d^* be the district with maximum cost for w . Denote by y^* the representative of d^* , and by o^* the optimal alternative for d^* that minimizes the cost of the agents in d^* according to g . We consider the following two subcases:

- **w is not λ -acceptable for d^* .** The fact that w is the leftmost district representative implies that y^* is to the right of w . If w is between o^* and y^* , then the fact that g is single-peaked would yield that

$$g(\vec{\delta}_{d^*}(w)) \leq g(\vec{\delta}_{d^*}(y^*)) \leq \lambda \cdot g(\vec{\delta}_{d^*}(o^*)),$$

thus contradicting that w is not λ -acceptable for d^* in this subcase. Hence, it must be the case that w is between o and o^* , and since g is single-peaked, we have that

$$(\text{MAX} \circ G)(w|I) = g(\vec{\delta}_{d^*}(w)) \leq g(\vec{\delta}_{d^*}(o)) \leq (\text{MAX} \circ G)(o|I).$$

- **w is λ -acceptable for d^* .** In this case, we clearly have that

$$\begin{aligned}
(\text{MAX} \circ G)(w|I) &= g(\vec{\delta}_{d^*}(w)) \\
&\leq \lambda \cdot g(\vec{\delta}_{d^*}(o^*)) \\
&\leq \lambda \cdot g(\vec{\delta}_{d^*}(o)) \\
&\leq \lambda \cdot (\text{MAX} \circ G)(o|I).
\end{aligned}$$

Putting everything together, we obtain an upper bound of $\max\{2 + \frac{1}{\lambda}, \lambda\}$. $\square$

By optimizing over the parameter λ , we obtain the following generalization of Corollary 4.5.

Corollary 6.7. *The distortion of $(1 + \sqrt{2})$ -ARL is at most $1 + \sqrt{2}$ for any objective of the form $\text{MAX} \circ G$, where G is defined by a monotone, subadditive, consistent, and single-peaked function g .*

7. Open problems

In this paper, we showed bounds on the distortion of single-winner distributed mechanisms for many different objectives, some of which are novel and make sense only in this particular setting. Still, there are several challenging open questions, as well as new directions for future research. Starting with our results, it would be interesting to close the gaps between the lower and upper bounds presented in Table 1 for the various scenarios we considered. For cases where our bounds for general metrics and the line differ significantly, such as for ordinal mechanisms and the $\text{AVG} \circ \text{MAX}$ objective, one could focus on other well-structured metrics, like the Euclidean space or generalizations of it.

Since we focused exclusively on deterministic mechanisms, a possible direction could be to consider randomized mechanisms and investigate whether better distortion bounds are possible. Note that our composition theorem (Theorem 3.1 and its variants) already provides randomized bounds by plugging in appropriate randomized in-district and over-districts direct voting rules. However, these bounds seem extremely loose, and different techniques are required to obtain tight bounds. Going beyond the single-winner setting, one could study the distortion of distributed mechanisms that output committees of a given number of alternatives, or rankings of all alternatives. Finally, another interesting direction would be to study what happens when agents act strategically, and either understand how this behavior affects given distributed mechanisms, or aim to design strategyproof mechanisms that are resilient to manipulation and at the same time achieve low distortion.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] Ben Abramowitz, Elliot Anshelevich, Utilitarians without utilities: maximizing social welfare for graph problems using only ordinal preferences, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI*, 2018, pp. 894–901.
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros A. Voudouris, Peeking behind the ordinal curtain: improving distortion via cardinal queries, *Artif. Intell.* 296 (2021) 103488.
- [3] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros A. Voudouris, A few queries go a long way: information-distortion tradeoffs in matching, in: *Proceedings of the 35th AAAI Conference on Artificial Intelligence, AAAI*, 2021, pp. 5078–5085.
- [4] Ioannis Anagnostides, Dimitris Fotakis, Panagiotis Patsilinakos, Metric-distortion bounds under limited information, in: *Proceedings of the 14th International Symposium on Algorithmic Game Theory, SAGT*, 2021, pp. 299–313.
- [5] Elliot Anshelevich, John Postl, Randomized social choice functions under metric preferences, *J. Artif. Intell. Res.* 58 (2017) 797–827.
- [6] Elliot Anshelevich, Wennan Zhu, Ordinal approximation for social choice, matching, and facility location problems given candidate positions, in: *Proceedings of the 14th International Conference on Web and Internet Economics, WINE*, 2018, pp. 3–20.
- [7] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, Piotr Skowron, Approximating optimal social choice under metric preferences, *Artif. Intell.* 264 (2018) 27–51.
- [8] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, Alexandros A. Voudouris, Distortion in social choice problems: the first 15 years and beyond, in: *Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI*, 2021, pp. 4294–4301.
- [9] Elliot Anshelevich, Aris Filos-Ratsikas, Alexandros A. Voudouris, The distortion of distributed metric social choice, in: *Proceedings of the 17th Conference on Web and Internet Economics, WINE*, 2021.
- [10] Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, Nisarg Shah, Preference elicitation for participatory budgeting, in: *Proceedings of the 31st AAAI Conference on Artificial Intelligence, AAAI*, 2017, pp. 376–382.
- [11] Gerdus Benadè, Ariel D. Procaccia, Mingda Qiao, Low-distortion social welfare functions, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI*, 2019, pp. 1788–1795.
- [12] Umang Bhaskar, Varsha Dani, Abheek Ghosh, Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI*, 2018, pp. 925–932.
- [13] Allan Borodin, Omer Lev, Nisarg Shah, Tyrone Strangway, Primarily about primaries, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI*, 2019, pp. 1804–1811.
- [14] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, Or Sheffet, Optimal social choice functions: a utilitarian view, *Artif. Intell.* 227 (2015) 190–213.
- [15] Ioannis Caragiannis, Ariel D. Procaccia, Voting almost maximizes social welfare despite limited communication, *Artif. Intell.* 175 (9–10) (2011) 1655–1671.
- [16] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, Nisarg Shah, Subset selection via implicit utilitarian voting, *J. Artif. Intell. Res.* 58 (2017) 123–152.
- [17] Ioannis Caragiannis, Nisarg Shah, Alexandros A. Voudouris, The metric distortion of multiwinner voting, in: *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI*, 2022.
- [18] Xujin Chen, Minming Li, Chenhao Wang, Favorite-candidate voting for eliminating the least popular candidate in a metric space, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI*, 2020, pp. 1894–1901.
- [19] Brandon Fain, Ashish Goel, Kamesh Munagala, Nina Prabh, Random dictators with a random referee: constant sample complexity mechanisms for social choice, in: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI*, 2019, pp. 1893–1900.
- [20] Michal Feldman, Amos Fiat, Iddan Golomb, On voting and facility location, in: *Proceedings of the 2016 ACM Conference on Economics and Computation, EC*, 2016, pp. 269–286.
- [21] Aris Filos-Ratsikas, Alexandros A. Voudouris, Approximate mechanism design for distributed facility location, in: *Proceedings of the 14th International Symposium on Algorithmic Game Theory, SAGT*, 2021, pp. 49–63.
- [22] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Jie Zhang, Social welfare in one-sided matchings: random priority and beyond, in: *Proceedings of the 7th Symposium on Algorithmic Game Theory, SAGT*, 2014, pp. 1–12.
- [23] Aris Filos-Ratsikas, Evi Micha, Alexandros A. Voudouris, The distortion of distributed voting, *Artif. Intell.* 286 (2020) 103343.
- [24] Vasilis Gkatzelis, Daniel Halpern, Nisarg Shah, Resolving the optimal metric distortion conjecture, in: *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, 2020, pp. 1427–1438.
- [25] Ashish Goel, Anilesh K. Krishnaswamy, Kamesh Munagala, Metric distortion of social choice rules: lower bounds and fairness properties, in: *Proceedings of the 2017 ACM Conference on Economics and Computation, EC*, 2017, pp. 287–304.
- [26] Michal Jaworski, Piotr Skowron, Evaluating committees for representative democracies: the distortion and beyond, in: *Proceedings of the 29th International Joint Conference on Artificial Intelligence, IJCAI*, 2020, pp. 196–202.
- [27] David Kempe, Communication, distortion, and randomness in metric voting, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI*, 2020, pp. 2087–2094.
- [28] David Kempe, An analysis framework for metric voting based on LP duality, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI*, 2020, pp. 2079–2086.
- [29] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, David P. Woodruff, Efficient and thrifty voting by any means necessary, in: *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems, NeurIPS*, 2019, pp. 7178–7189.
- [30] Debmalya Mandal, Nisarg Shah, David P. Woodruff, Optimal communication-distortion tradeoff in voting, in: *Proceedings of the 21st ACM Conference on Economics and Computation, EC*, 2020, pp. 795–813.
- [31] Kamesh Munagala, Kangning Wang, Improved metric distortion for deterministic social choice rules, in: *Proceedings of the 2019 ACM Conference on Economics and Computation, EC*, 2019, pp. 245–262.
- [32] Ariel D. Procaccia, Jeffrey S. Rosenschein, The distortion of cardinal preferences in voting, in: *International Workshop on Cooperative Information Agents, CIA*, 2006, pp. 317–331.
- [33] Amartya Sen, Social choice theory, in: *Handbook of Mathematical Economics*, vol. 3, 1986, pp. 1073–1181.