Contents lists available at ScienceDirect





journal homepage: www.elsevier.com/locate/artint



Improved metric distortion via threshold approvals Elliot Anshelevich^{a,[®]}, Aris Filos-Ratsikas^{b,[®]}, Christopher Jerrett^{a,[®]}, Alexandros A. Voudouris^{c,[®]},*

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

^b School of Informatics, University of Edinburgh, UK

° School of Computer Science and Electronic Engineering, University of Essex, UK

ABSTRACT

We consider a social choice setting in which agents and alternatives are represented by points in a metric space, and the cost of an agent for an alternative is the distance between the corresponding points in the space. The goal is to choose a single alternative to (approximately) minimize the social cost (cost of all agents) or the maximum cost of any agent, when only limited information about the preferences of the agents is given. Previous work has shown that the best possible distortion one can hope to achieve is 3 when access to the ordinal preferences of the agents is given, even when the distances between alternatives in the metric space are known. We improve upon this bound of 3 by designing deterministic mechanisms that exploit a bit of cardinal information. We show that it is possible to achieve distortion $1 + \sqrt{2}$ by using the ordinal preferences of the agents, the distances between alternatives, and a threshold approval set per agent that contains all alternatives that are at distance from the agent within an appropriately chosen factor of the minimum distance of the agents from any alternative. We show that this bound is the best possible for any deterministic mechanism in general metric spaces, and also provide improved bounds for the fundamental case of a line metric.

1. Introduction

Social choice theory is concerned with aggregating the heterogeneous preferences of individuals over a set of outcomes into a single decision [17]. Besides its many applications in traditional domains, such as political elections or voting for policy issues, social choice theory has also been at the epicenter of areas such as *multi-agent systems* [29], *recommendation systems* [37], *search engines* [26], and *crowdsourcing* applications [48]. Aggregation methods inspired from social choice theory have also been used in relation to machine learning, such as for *regression* and *estimation* tasks [18,24,25,41], as well as *virtual democracy* [40,50,51].

A central theme underpinning many of these applications is the interplay between the amount of available information and the efficiency of the implemented decisions [5,16,46,56,57]. Indeed, it is often the case that access to the preferences of the participants (or *agents*) is restricted due to lack of enough statistical data, or due to computational or cognitive limitations in the elicitation process. This leads to the natural question of whether we can design aggregation rules (or *mechanisms*) that achieve high efficiency even when presented with informational restrictions.

This is the main question studied in the literature of *distortion* [12,53], which measures the deterioration of an aggregate social objective due to limited information about the preferences of the agents. In its original setup (e.g., see [16]), limited information is interpreted as having access only to the *ordinal preference rankings* over the possible outcomes, instead of a complete *cardinal utility* structure that fully captures the intensity of the preferences. Over the years, the notion of distortion has been refined to capture different kinds of informational limitations, including restricted ordinal information [39,42], communication complexity [46,47], and query complexity [3,5].

https://doi.org/10.1016/j.artint.2025.104295

Available online 30 January 2025

^{*} Corresponding author. *E-mail address:* alexandros.voudouris@essex.ac.uk (A.A. Voudouris).

Received 21 January 2024; Received in revised form 25 January 2025; Accepted 28 January 2025

^{0004-3702/© 2025} 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/).

Table 1

An overview of our distortion bounds for deterministic mechanisms that use different types of information. ORD is used to refer to the ordinal preferences of the agents, DIST is used the distances between alternatives, and TAS is used for α -threshold approval sets.

		$TAS \wedge ORD \wedge DIST$	$TAS\wedgeORD$	$\text{TAS} \land \text{DIST}$	TAS
Social cost	General Line	$[2, 1 + \sqrt{2}]$ $2\sqrt{2} - 1$	$[1 + \sqrt{2}, 3]$ $2\sqrt{2} - 1$	$\frac{1+\sqrt{2}}{2}$	$\Theta(n)$
Max cost	General Line	$1 + \sqrt{2}$ $1 + \sqrt{2}$	$[1 + \sqrt{2}, 3]$		3

The literature on *metric distortion* considers scenarios where the agents and the outcomes exist in a (possibly high-dimensional) metric space, and the costs (rather than utilities) are given by distances that satisfy the triangle inequality. The economic interpretation of these distances is typically in terms of the proximity for political or ideological issues along different axes (e.g., liberal to conservative, or libertarian to authoritarian). The metric distortion of social choice mechanisms, firstly studied in the works of Anshelevich et al. [11] and Anshelevich and Postl [6], is one of the most well-studied settings in the context of this literature, with many variants of the main setting being considered in recent years (see Section 3 of the survey of Anshelevich et al. [12] for an overview). Since its inception in 2015 (the conference version of [11]), the "holy grail" of this literature was a mechanism with a metric distortion of 3 for the *social cost objective* (i.e., the sum of costs of all agents), which would match the corresponding lower bound shown by Anshelevich et al. [11]. The upper bound was finally established by Gkatzelis et al. [38] via the PLURALITY MATCHING mechanism, and then later also by Kizilkaya and Kempe [43] via the much simpler PLURALITY VETO mechanism; see also [44,52]. For the other most natural objective, that of minimizing the *maximum cost* of any agent, the best possible bound can also easily be seen to be 3; this bound can be achieved by the aforementioned mechanisms, and many other simpler ones (see Section 4).

Crucially, these tight bounds of 3 are achieved by mechanisms without access to any information of a cardinal nature, i.e., they work by using only the ordinal preference rankings of the agents. It can also be observed (e.g., see [10]) that 3 is still the best bound one can hope for even if one also assumes access to the distances between the outcomes in the metric space, a natural assumption when, for example, the outcomes correspond to alternative locations where a public facility can be built or to alternative opinions (across multiple dimensions) expressed by election candidates. At the same time, a growing recent literature advocates to take the natural next step, and study the distortion when a small amount of cardinal information about the preferences is also available, which, in many cases, is reasonable to elicit [2–5,14,15]. This motivates the following question:

Is it possible to beat the barrier of 3 on the metric distortion if we also have access to limited information of a cardinal nature?

1.1. Our contribution

We consider a setting in which a set N of agents and a set A of alternatives are positioned on a metric space, and the preferences of the agents are captured by the distances between their positions and that of the different alternatives. We are interested in the distortion of deterministic mechanisms in terms of the social cost and the maximum cost objectives, i.e., the ratio of the cost of the alternative chosen by the mechanism over the smallest possible cost over all alternatives, taken over all possible instances. Note that with full cardinal information, achieving an optimal distortion of 1 is trivial.

We answer the main question posed above in the affirmative, by assuming access to the ordinal preferences, the distances between the alternatives in the metric space, and some additional limited cardinal information about the costs of the agents. In particular, for each agent $i \in N$, we have access to an α -threshold approval set (α -TAS) A_i , which contains all the alternatives for which the agent has cost at most a factor α times the cost for her¹ most-preferred alternative, with $\alpha \ge 1$ being a value that can in general be chosen by the mechanism designer. Such threshold approval information has been used before in the distortion literature for other problems (see the related work discussion below for more details) and can be thought of as asking the agents to partition the alternatives into "buckets" of different importance (that is determined by their distance to them relative to the minimum distance over all alternatives). For example, information of this type is implicitly elicited in conference bidding systems and in other systems for rating hotels or restaurants.

For general metric spaces, we design mechanisms that achieve a distortion of $1 + \sqrt{2}$ in terms of the social cost and the maximum cost objectives, thus beating the aforementioned barrier of 3. We complement these results with lower bounds on the distortion that *any* deterministic mechanism can achieve when using this amount of information. For the fundamental special case of a line metric, we provide refined tight bounds on the distortion for both objectives, showing an even larger separation from the bound of 3 that holds in the absence of any cardinal information and even when the metric is a line. More generally, we fine-tune our analysis to the presence/absence of all three types of information (i.e., only α -TAS, α -TAS + ordinal preferences, α -TAS + known alternative distances) and show that, in most cases, the best possible distortion bounds can be obtained even when only using two types of information. Our results are summarized in Table 1.

¹ For clarity, we will use female pronouns for agents and male pronouns for alternatives whenever necessary.

1.2. Related work and discussion

The distortion literature is rather extensive, focusing primarily on the utilitarian normalized setting (e.g., see [14-16,19,21,23, 27,32,34,46,47]) or the metric setting that we consider in this paper (e.g., see [1,2,6-11,13,20,22,30,31,33,35,36,38,43,49,55]). For more details, we refer the reader to the recent survey of Anshelevich et al. [12]. Below we discuss the works that are mostly relevant to the setting that we study here.

The term "approval threshold" (in fact, "approval threshold vote") was introduced in the conference version of the paper of Benade et al. [14] was used in the context of participatory budgeting to refer to an elicitation device which returns a set of alternatives that an agent values higher than a given threshold; it was later also used by Bhaskar et al. [15] for truthful multiwinner voting, and by Latifian and Voudouris [45] for one-sided matching. Amanatidis et al. [3] introduced a setting where limited cardinal information is elicited via a set of *value* or *comparison queries*; an agent is asked to provide her value for a given alternative, or indicate whether she prefers one alternative over another by more than a given factor. An approval threshold in their setting can be constructed by a single value query (for the agent's most-preferred alternative), and then a number of comparison queries that is logarithmic in the number of alternatives.

The α -TAS that we consider are closely related to the comparison query model of Amanatidis et al. [3], in that they do not encode information about any absolute value for the cost of the agents, but rather only information that is *relative* to their cost for their most-preferred alternatives. In this sense, these sets are *relative*, and can be viewed as the outcomes of a single relative threshold query, or, equivalently, the outcome of logarithmically-many comparison queries. From a cognitive standpoint, this is an even more conceivable elicitation device since the agents are only asked to perform (cardinal) comparisons rather than to come up with absolute cost numbers. Such comparisons are motivated by and rooted in the ideas of the Von Neumann-Morgenstern utility theory [54]; see also [3].

A common characteristic of all aforementioned works is that they do not consider the distortion in the metric setting. Abramowitz et al. [2] studied the metric distortion of mechanisms that use limited cardinal information, namely information about whether the ratio between the costs of an agent for any two candidates is larger or smaller than a chosen threshold τ . Their elicitation method is different from ours, and requires information about the relative costs for all pairs of alternatives. In contrast, our α -TAS only require aggregate information about the set A_i of alternatives. The two settings coincide only in the case where there are just 2 alternatives overall, which allows us to obtain a lower bound of $2\sqrt{2} - 1 \approx 1.83$ for the social cost for the line metric from their work, for which we also show a matching upper bound.

2. Preliminaries

We consider a single-winner social choice setting with a set N of n agents and a set A of m alternatives. Agents and alternatives are represented by points in a metric space. For any $x, y \in N \cup A$, let d(x, y) denote the distance between the points representing x and y. The distances satisfy the standard conditions for metric spaces, namely that d(x, y) = d(y, x), d(x, y) = 0 if x = y, and $d(x, y) \leq d(x, z) + d(z, y)$; the last condition is called the *triangle inequality*. We will use the tuple I = (N, A, d) to denote an *instance* of our setting, and will use I to denote the set of all instances. When the instance is clear from context in our proofs, we will simplify our notation by omitting it.

2.1. Input information

A (deterministic) *mechanism* M takes as input some information related to the distances between agents and alternatives, and outputs a single alternative as the *winner*. We will consider combinations of the following three types of information:

- ORD: The *ordinal preferences* $(\succ_i)_{i \in N}$ of the agents which are induced by the distances, where \succ_i is a complete ordering over the elements of *A*, and $x \succ_i y$ implies that $d(i, x) \le d(i, y)$ for any $i \in N$ and $x, y \in A$.
- DIST: The distances d(x, y) between any pair of alternatives $x, y \in A$.
- TAS: A set of α -threshold approval sets (α -TAS) $(A_i)_{i \in N}$ such that $d(i, x) \le \alpha \cdot d(i, o_i)$ for any $x \in A_i$, where $\alpha \ge 1$ and o_i is the most-preferred alternative of agent *i* (i.e., $o_i \succ_i y$ for any $y \in A$). We will say that agent *i* approves alternative *x* in case $x \in A_i$.

Clearly, all types of information can be inferred from an instance *I*. We will use the term Inf(I) to denote the information that is available to a mechanism for instance *I*, which can be any combination of ORD, DIST and TAS. We will use the logical AND (\land) to indicate the types of information that a mechanism has access to.

As we already explained in the Introduction, the bulk of the metric distortion literature considers mechanisms that have access to ORD, which results in a best-possible distortion of 3 [11,38,43] (see below for the formal definition). For mechanisms with access to ORD \land DIST, the best possible distortion is still 3 [10]. Our main results are for the case of having access to all types of information (ORD \land DIST \land TAS), but we also provide results for ORD \land TAS and DIST \land TAS independently, as well as for TAS in isolation.

2.2. Objectives and distortion

We consider two well-known minimization objectives, the *social cost* (SC) and the *maximum cost* (MC). Given an instance I, the social cost of an alternative x is

Mechanism 1: α-MINISUM-TAS-DISTANCE.

Input: Distances between alternatives, *a*-TAS;
Output: Winner *w*;
$$w \in \arg \min_{x \in A} \left\{ \sum_{i \in N} \min_{j \in A_i} d(j, x) \right\};$$

$$\mathrm{SC}(x|I) = \sum_{i \in \mathbb{N}} d(i, x),$$

and the maximum cost of *x* is

$$\mathrm{MC}(x|I) = \max_{i \in \mathbb{N}} d(i, x).$$

Given an objective $F \in \{SC, MC\}$, the *distortion* of a mechanism M for F is defined as

$$\sup_{I \in \mathcal{I}} \frac{\mathrm{F}(M(\mathrm{Inf}(I))|I)}{\min_{j \in A} \mathrm{F}(j|I)},$$

where M(Inf(I)) denotes the alternative chosen by the mechanism when given as input the information Inf(I) for instance I. Clearly, the distortion is at least 1 for any mechanism, and our goal is to determine the best possible distortion when given combinations of Section 2.1.

3. Social cost

In this section, we show bounds on the distortion for the social cost objective for mechanisms that have access to ORD \land DIST \land TAS. Our main result is a mechanism, coined $(1 + \sqrt{2})$ -MINISUM-TAS-DISTANCE, which achieves a distortion of $1 + \sqrt{2} \approx 2.42$, thus breaking the barrier of 3, which is the best possible without access to the α -TAS. In fact, this mechanism does not require access to the ordinal preferences of the agents, it only uses DIST \land TAS. We complement this result with a lower bound of 2 on the distortion of any deterministic mechanism with access to ORD \land DIST \land TAS.

For mechanisms that use ORD \wedge TAS and DIST \wedge TAS, we provide stronger lower bounds of $1 + \sqrt{2}$; this establishes that $(1 + \sqrt{2})$ -MINISUM-TAS-DISTANCE is the best possible among mechanisms with access to DIST \wedge TAS. We also prove that the distortion of any mechanism with access to TAS is bound to be bad, namely $\Theta(n)$. For the special case of the line metric, we provide more refined *tight* bounds depending on the amount of information available in the input.

3.1. Results for general metric spaces

We start by showing an upper bound of $1 + \sqrt{2}$ for general metric spaces using a mechanism to which we refer as α -MINISUM-TAS-DISTANCE. This mechanism is quite intuitive and does not require any information about the ordinal preferences. The mechanism chooses as the winner an alternative that minimizes the sum of distances from the α -TAS of the agents, where the distance between an alternative *x* and a set A_i is defined as the distance of *x* from its closest alternative in A_i . See Mechanism 1 for a description using pseudocode.

Theorem 3.1. In general metric spaces, the distortion of α -MINISUM-TAS-DISTANCE for the social cost objective is at most max $\left\{\alpha, 2+\frac{1}{\alpha}\right\}$.

Proof. Recall that o_i is the most-preferred alternative of agent *i*; hence, $d(i, j) \le \alpha \cdot d(i, o_i)$ for any $j \in A_i$. Let *o* be the optimal alternative (i.e., the alternative with the smallest social cost) and let *w* be the alternative chosen by the mechanism. For any agent *i*, denote by $x_i \in A_i$ the alternative that is closest to *w* among those approved by *i*, and by $y_i \in A_i$ the alternative that is closest to *o*. By the definition of *w*, we have that $\sum_{i \in N} d(x_i, w) \le \sum_{i \in N} d(y_i, o)$. Using the triangle inequality, we now have:

$$\begin{split} \mathsf{SC}(w) &= \sum_{i \in N} d(i, w) \leq \sum_{i \in N} d(i, x_i) + \sum_{i \in N} d(x_i, w) \\ &\leq \sum_{i \in N} d(i, x_i) + \sum_{i \in N} d(y_i, o). \end{split}$$

We make the following observations:

- For every agent *i* such that $o \in A_i$, $y_i = o$, and thus $d(y_i, o) = 0$, as well as $d(i, o_i) \le d(i, o)$.
- For every agent *i* such that $o \notin A_i$, $d(y_i, o) \le d(o_i, o)$ and $d(i, o_i) < \frac{1}{\alpha} \cdot d(i, o)$.

Combining the above with the fact that $d(i, x_i) \le \alpha \cdot d(i, o_i)$ and using the triangle inequality, we obtain

$$SC(w) \le \sum_{i \in N} d(i, x_i) + \sum_{i \in N} d(y_i, o)$$



Fig. 1. An example of the metric space used in the proof of Lemma 3.3 with i = 1 to show a lower bound of α , assuming that $\alpha \leq 3$. Agents are represented by rectangles, and alternatives by circles.

$$\begin{split} &= \sum_{i:o \in A_i} \left(d(i,x_i) + d(y_i,o) \right) + \sum_{i:o \notin A_i} \left(d(i,x_i) + d(y_i,o) \right) \\ &\leq \sum_{i:o \in A_i} d(i,x_i) + \sum_{i:o \notin A_i} \left(d(i,x_i) + d(o_i,o) \right) \\ &\leq \alpha \cdot \sum_{i:o \in A_i} d(i,o_i) + \sum_{i:o \notin A_i} \left(\alpha \cdot d(i,o_i) + d(i,o_i) + d(i,o) \right) \\ &\leq \alpha \cdot \sum_{i:o \notin A_i} d(i,o) + \left(2 + \frac{1}{\alpha} \right) \cdot \sum_{i:o \notin A_i} d(i,o) \\ &\leq \max \left\{ \alpha, 2 + \frac{1}{\alpha} \right\} \cdot \operatorname{SC}(o). \end{split}$$

Consequently, the distortion is at most $\max \left\{ \alpha, 2 + \frac{1}{\alpha} \right\}$. \Box

By optimizing over α , we obtain the following corollary.

Corollary 3.2. In general metric spaces, the distortion of $(1 + \sqrt{2})$ -MINISUM-TAS-DISTANCE for the social cost objective is at most $1 + \sqrt{2}$.

We complement the upper bound of $1 + \sqrt{2}$ by a nearly tight lower bound of 2 for mechanisms that use all three types of information. We first show a lower bound of min{3, α } on the distortion of such mechanisms; the following lemma will be used in several lower bound proofs, and also suggests that to beat the bound of 3, a mechanism must use $\alpha < 3$.

Lemma 3.3. For the social cost objective, the distortion of any mechanism M with access to $ORD \wedge DIST \wedge TAS$ is at least min{3, α }, for any $\alpha \ge 1$.

Proof. We consider an instance with *n* agents and *n* alternatives $A = \{a_1, ..., a_n\}$. The distance between any two alternatives is 2. The agents have cyclic ordinal preferences over the alternatives such that for agent *i*:

$$a_i \succ_i \ldots \succ_i a_n \succ_i a_1 \succ \ldots \succ_i a_{i-1}.$$

In addition, all agents approve all alternatives, that is, $a \in A_i$ for every $a \in A$ and every agent *i*. Suppose that the mechanism chooses alternative a_i for some $i \in [n]$ as the winner when given as input this information about the instance.

Observe that, besides agent *i*, all other n - 1 agents prefer a_{i-1} to a_i , interpreting a_0 as a_n . Given this, we aim to define a metric space that is consistent with the distances between the alternatives, as well as with the ordinal preferences of the agents, so as to maximize the social cost of a_i while minimizing the social cost of a_{i-1} . In particular, we have the following metric space (see Fig. 1 for an example with i = 1):

• Agent *i* is at distance 1 from all alternatives.

- Agent j < i is at distance 1 from all alternatives in $\{a_j \dots, a_{i-1}\}$ and at distance min $\{3, \alpha\}$ from the remaining alternatives in $\{a_i, \dots, a_n, a_1, \dots, a_{i-1}\}$.
- Agent j > i is at distance 1 from the alternatives in $\{a_j, \ldots, a_n\}$ and at distance min $\{3, \alpha\}$ from the remaining alternatives in $\{a_1, \ldots, a_{i-1}\}$; note that a_i belong to the latter set of alternatives in this case.
- The distance between two alternatives is determined by the length of the minimum path that connects them; hence, it is exactly 2 (due to the distance 1 of agent *i* from all alternatives).

Since all agents, beside *i*, are at distance 1 from a_{i-1} and at distance $min\{3, \alpha\}$ from a_i , we have

$$\frac{\operatorname{SC}(a_i)}{\operatorname{SC}(a_{i-1})} = \frac{1 + (n-1) \cdot \min\{3, \alpha\}}{n}.$$

As *n* becomes arbitrarily large, we obtain a lower bound of min $\{3, \alpha\}$ on the distortion.

We are now ready to show the lower bound of 2 on the distortion of any mechanism that uses all three types of information.

Theorem 3.4. For the social cost objective, the distortion of any mechanism M with access to $ORD \land DIST \land TAS$ is at least 2.

Proof. By Lemma 3.3, the distortion of any deterministic mechanism is at least min $\{3, \alpha\}$, even when the ordinal preferences of the agents are given and the distances between alternatives are fixed. This implies that the distortion is at least 2 if $\alpha \ge 2$.

For the case $\alpha < 2$, let $\delta > 0$ be an infinitesimal and consider the following instance with *n* agents and 2*n* alternatives $A = \{a_1, \ldots, a_n, b_1, \ldots, b_n\}$. Any pair of *a*-type alternatives are at distance $1 + 2\alpha$, any pair of *b*-type alternatives are at distance 2α , and any pair of alternatives consisting of one *a*-type and one *b*-type are at distance $1 + \alpha$. Let [X] be an arbitrary permutation of the elements of *X*. Each agent $i \in [n]$ has the ordinal preference

 $a_i \succ_i b_i \succ_i \ldots \succ_i b_n \succ_i b_1 \succ \ldots \succ_i b_{i-1} \succ [\{a_1, \ldots, a_n\} \setminus \{a_i\}].$

In addition, agent *i* reports that she only approves alternative a_i , that is, $A_i = \{a_i\}$. Now, suppose that, for some $i \in [n]$, the mechanism chooses a_i or b_i as the winner. We consider the following metric space:

- Agent *i* is at distance 1 from alternative a_i , at distance $\alpha + \delta$ from each alternative in $\{b_1, \dots, b_n\}$, and at distance $1 + 2\alpha + 2\delta$ from the remaining alternatives in $\{a_1, \dots, a_n\} \setminus \{a_i\}$.
- Agent j < i is at distance 1 from alternative a_j , at distance $\alpha + \delta$ from every alternative in $\{b_j, \dots, b_{i-1}\}$, and at distance $2 + \alpha$ from every alternative in $\{b_i, \dots, b_n, \dots, b_{j-1}, a_1, \dots, a_n\} \setminus \{a_j\}$.
- Agent j > i is at distance 1 from alternative a_j , at distance $\alpha + \delta$ from every alternative in $\{b_j, \dots, b_n\}$, and at distance $2 + \alpha$ from every alternative in $\{b_1, \dots, b_{i-1}, a_1, \dots, a_n\} \setminus \{a_i\}$.

Clearly, the specified distances are consistent with the ordinal preferences of the agents (with ties broken as required). It is also not hard to verify that these distances, together with the fixed distances between alternatives, satisfy the triangle inequality; for an example of the metric space for the case i = 1 see Fig. 2. By interpreting b_0 as b_n , observe that every agent $j \neq i$ is at distance $\alpha + \delta$ from b_{i-1} , and distance $1 + 2\alpha$ from b_i and a_i , and thus

$$SC(a_i) = 1 + (n-1)(2 + \alpha),$$

$$SC(b_i) = \alpha + \delta + (n-1)(2+\alpha)$$

and

$$SC(b_{i-1}) = n(\alpha + \delta).$$

As *n* becomes arbitrarily large and δ becomes arbitrarily small, since $\alpha < 2$, the distortion is at least $1 + 2/\alpha > 2$.

3.2. Refined lower bounds for general metric spaces

As we mentioned earlier, α -MINISUM-TAS-DISTANCE achieves a distortion of $1 + \sqrt{2}$ without having access to the ordinal preferences of the agents. This begs the natural question of whether this is the best possible among all mechanisms with this information available in their inputs. We prove that this is indeed the case in the following theorem.

Theorem 3.5. For the social cost objective, the distortion of any mechanism M with access to DIST \wedge TAS is at least $1 + \sqrt{2}$.

Proof. By Lemma 3.3, any mechanism has distortion at least min{3, α }, even when the distances between alternatives are known. This implies that the distortion is at least $1 + \sqrt{2}$ if $\alpha \ge 1 + \sqrt{2}$.



Fig. 2. An example of the metric space used in the proof of Theorem 3.4 to show a lower bound of $1 + 2/\alpha$ for i = 1. Distances between alternatives are not shown, but are assumed to be fixed. Agents are represented by rectangles, and alternatives by circles. The notation α^+ is used to represent a value slight larger than α (instead of $\alpha + \delta$). It is not hard to verify that the specified distances satisfy the triangle inequality; for example, the distance of the edge between agent *j* and alternative b_1 is $2 + \alpha$, whereas the length of a path from *j* to b_1 via any other *b*-type alternative is $2\alpha + \alpha^+$, and the length of the path via a_j is $1 + 1 + \alpha = 2 + \alpha$.



Fig. 3. The metric space used in the proof of Theorem 3.5. The notation α^+ is used to represent a value slight larger than α (essentially α plus an infinitesimal). The red and green squares are potential agent positions, depending on the decision made by the mechanism when given that the α -TAS of each agent *i* is $A_i = \{x_i, z_i\}$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

For the case $\alpha < 1 + \sqrt{2}$, we consider instances with *n* agents and m = 2n + 2 alternatives $\{x_1, ..., x_n, z_1, ..., z_n, y, w\}$. The α -TAS of agent *i* is $A_i = \{x_i, z_i\}$. The metric space is presented in Fig. 3; in the figure, α^+ is used to denote α plus some infinitesimal. Without loss of generality, we assume that the mechanism chooses one of $\{x_1, ..., x_n, w\}$; the case where the mechanism chooses one of $\{z_1, ..., z_n, y\}$ is symmetric.

We place the agents at the red nodes in Fig. 3; for the case where the mechanism chooses the winner from the set $\{z_1, ..., z_n, y\}$, we would place the agents at the green nodes in Fig. 3. Agent *i* is at distance 1 from z_i , at distance $1 + \alpha - 1 = \alpha$ from x_i , and at distance strictly larger than α from any other alternative; so, this placement of the agents is consistent with the α -TAS. Since the distance of agent *i* is essentially α from *y*, at least $1 + 2\alpha$ from *w*, and 3α from any alternative x_j with $j \neq i$, we have

$$SC(x_i) = \alpha + 3(n-1)\alpha = (3n-2)\alpha$$
$$SC(w) = n(1+2\alpha),$$

and

 $SC(y) = n\alpha$.

So, the distortion is at least 3 - 2/n if any alternative x_i is chosen, and at least $2 + 1/\alpha$ if w is chosen. Since $\alpha < 1 + \sqrt{2}$, this leads to a lower bound of $2 + 1/\alpha > 1 + \sqrt{2}$.

We also show the same distortion lower bound of $1 + \sqrt{2}$ for mechanisms that have access to the α -TAS and the ordinal preferences of the agents, but not the distances between alternatives.



Fig. 4. An example of the metric space used in the proof of Theorem 3.6 to show a lower bound of $2 + 1/\alpha$ for i = 1. Agents are represented by rectangles, and alternatives by circles. The notation α^+ is used to represent a value slightly larger than α (instead of $\alpha + \delta$).

Theorem 3.6. For the social cost objective, the distortion of any mechanism M with access to $ORD \land TAS$ is at least $1 + \sqrt{2}$.

Proof. By Lemma 3.3, the distortion of any deterministic mechanism is at least $\min\{3, \alpha\}$, even when the ordinal preferences of the agents are given. This implies that the distortion is at least $1 + \sqrt{2}$ if $\alpha \ge 1 + \sqrt{2}$.

For the case $\alpha < 1 + \sqrt{2}$, we consider an instance similar to the one used in the proof of Theorem 3.4 without the requirement to have fixed distances between alternatives (as this information is not used by the mechanism). There are *n* agents and 2*n* alternatives $A = \{a_1, \dots, a_n, b_1, \dots, b_n\}$. Let [X] be an arbitrary permutation of the elements of *X*. Each agent $i \in [n]$ has the ordinal preference

$$a_i \succ_i b_i \succ_i \ldots \succ_i b_n \succ_i b_1 \succ \ldots \succ_i b_{i-1} \succ [\{a_1, \ldots, a_n\} \setminus \{a_i\}].$$

In addition, agent *i* reports that she only approves of alternative a_i , that is, $A_i = \{a_i\}$. Let $\delta > 0$ be an infinitesimal. Now, suppose that, for some $i \in [n]$, the mechanism chooses a_i or b_i as the winner. We consider the following metric space:

- Agent *i* is at distance 1 from alternative a_i , at distance $\alpha + \delta$ from each alternative in $\{b_1, \dots, b_n\}$, and at distance $1 + 2(\alpha + \delta)$ from the remaining alternatives in $\{a_1, \dots, a_n\} \setminus \{a_i\}$.
- Agent j < i is at distance 1 from alternative a_j , at distance $\alpha + \delta$ from every alternative in $\{b_j, \dots, b_{i-1}\}$, at distance $1 + 2\alpha$ from every alternative in $\{b_i, \dots, b_n, \dots, b_{j-1}\}$, and at distance $1 + 2(\alpha + \delta)$ from every alternative in $\{a_1, \dots, a_n\} \setminus \{a_j\}$.
- Agent j > i is at distance 1 from alternative a_j , at distance $\alpha + \delta$ from every alternative in $\{b_j, \dots, b_n\}$, at distance $1 + 2\alpha$ from every alternative in $\{b_1, \dots, b_{j-1}\}$, and at distance $1 + 2(\alpha + \delta)$ from every alternative in $a_1, \dots, a_n\} \setminus \{a_j\}$.

An example of the metric space for the case i = 1 is given in Fig. 4. Interpreting b_0 as b_n , we can observe that every agent $j \neq i$ has value $\alpha + \delta$ for b_{i-1} , and value $1 + 2\alpha$ for b_i and a_i , and thus

$$SC(a_i) = 1 + (n-1)(1 + 2\alpha + 2\delta)$$

 $SC(b_i) = \alpha + \delta + (n-1)(1+2\alpha),$

and

 $SC(b_{i-1}) = n(\alpha + \delta).$

As *n* becomes arbitrarily large and δ becomes arbitrarily small, since $\alpha < 1 + \sqrt{2}$, we obtain a lower bound of $2 + 1/\alpha \ge 2 + 1/(1 + \sqrt{2}) = 1 + \sqrt{2}$.

It is an interesting open problem whether there exists a mechanism that inputs only the α -TAS and the ordinal preferences and achieves a distortion of $1 + \sqrt{2}$ in general metric spaces. In the next subsection, we will show that, for the line metric, there is a mechanism with these informational inputs that achieves the best possible distortion, even among mechanisms that use all three types of information.

Before we do that however, we show that any mechanism that uses only the α -TAS (i.e., without having access to the distances between alternatives or the ordinal preferences) is bound to have a very high distortion.

Theorem 3.7. For the social cost objective, the distortion of mechanisms with access to TAS is $\Theta(n)$.

Proof. Clearly, if we choose $\alpha = 1$, the α -TAS of each agent consists of her most-preferred alternative (or most-preferred alternatives in case there are more than one at equal minimum distance from the agent). To obtain an asymptotically linear upper bound, it suffices to consider the very simple mechanism that uses $\alpha = 1$, and outputs the most-preferred alternative o_i of an arbitrary agent *i* as the winner *w*. Denoting by *o* an optimal alternative, and using the triangle inequality, we have

$$\begin{split} \mathrm{SC}(w) &= \sum_{j \in N} d(j, o_i) \\ &\leq \sum_{j \in N} \left(d(j, o) + d(i, o) + d(i, o_i) \right) \\ &\leq \sum_{j \in N} \left(d(j, o) + 2d(i, o) \right) \\ &\leq (1 + 2n) \cdot \mathrm{SC}(o), \end{split}$$

and thus the distortion is $\Theta(n)$ when we only have access to the α -TAS.

For the lower bound, consider an arbitrary deterministic mechanism that has access only to the α -TAS, and the following instance with *n* agents and *n* alternatives $\{a_1, \ldots, a_n\}$. Each agent *i* approves only alternative a_i , that is, $A_i = \{a_i\}$. Since the given information is symmetric, the mechanism can choose any of the alternatives as the winner, say a_1 . Let $\varepsilon > 0$ be an infinitesimal. We define the following metric space:

- Agent 1 and alternative a_1 are both located at 1 of the line of real numbers.
- Agent $i \in \{2, ..., n-1\}$ and alternative a_i are located at a unique point within the ε -neighborhood of 0, but not at 0.
- Agent *n* and alternative a_n are located at 0.

Since every agent *i* is located exactly where their most-preferred alternative a_i is located, this metric space is consistent with the α -TAS given as input, for any α . As ε tends to 0, we have that $SC(a_1) = n - 1$, whereas $SC(a_n) = 1$, leading to a lower bound of n - 1.

3.3. Refined results for the line metric

When the metric space is a line, we will show refined tight bounds for mechanisms that use different types of information. The best possible distortion achievable here is $2\sqrt{2} - 1 \approx 1.83$. The lower bound follows from the work of Abramowitz et al. [2], via a connection between their setting and ours, when there are only two alternatives; we provide a proof below using our terminology for completeness. More interestingly, by combining properties shown in previous work, we propose a mechanism, coined α -ELIMINATION-WEIGHTED-MAJORITY, that achieves the distortion of $2\sqrt{2} - 1 \approx 1.83$ as an upper bound for $\alpha = \sqrt{2} - 1$, and is thus best possible. In fact, our mechanism does not need to have access to the distances between alternatives, only to the α -TAS and the ordinal preferences of the agents.

Since the metric space is a line, using the ordinal preferences of the agents, we can determine an ordering of the agents on the line (which is unique up to permutations of identical agents and reversal) and the unique ordering of the alternatives that are between the most-preferred alternatives of the leftmost and rightmost agents [28]. Using this, Anshelevich and Postl [6] showed that the optimal alternative is either the most-preferred alternative of the median agent or one of the alternatives directly to the left or right of that. Our mechanism exploits this property to eliminate all but two alternatives, and then runs a weighted majority over these alternatives to determine the winner. In particular, our mechanism works in the following two steps:

- *Elimination step*: It identifies the most-preferred alternative x of the median agent, the alternative ℓ that is directly the left of x, and the alternative r that is directly to the right of x. It then eliminates one of ℓ and r by comparing the number $n(\ell, x)$ of agents that prefer ℓ to x and the number n(r, x) of agents that prefer r to x as in the work of Anshelevich and Postl [6]. In particular, if $n(\ell, x) \leq n(r, x)$, then it eliminates ℓ , otherwise it eliminates r, and stores the non-eliminated alternative as y.
- Weighted-majority step: Afterwards, the mechanism runs a weighted majority between x and y by assigning a weight of 1 to each agent *i* such that $x, y \in A_i$, and a weight of $\frac{\alpha+1}{\alpha-1}$ to every other agent. This step will be shown to be equivalent to the algorithm used by Abramowitz et al. [2] to show a bound of $2\sqrt{2} 1 \approx 1.83$ in their setting when there are two alternatives.

See Mechanism 2 for a description using pseudocode.

Theorem 3.8. When the metric space is a line, for the social cost objective, the distortion of α -ELIMINATION-WEIGHTED-MAJORITY is at most $\max\{\frac{3\alpha-1}{\alpha+1}, 1+\frac{2}{\alpha}\}$.

Mechanism 2: α-ELIMINATION-WEIGHTED-MAJORITY.

Input: Ordinal preferences, α -TAS; Output: Winner w: $x \leftarrow$ most-preferred alternative of median agent; $\ell \leftarrow$ alternative directly to the left of *x*; $r \leftarrow$ alternative directly to the right of x; // Elimination step $n(\ell, x) \leftarrow |\{i \in N : \ell \succ_i x\}|;$ $n(r, x) \leftarrow |\{i \in N : r \succ_i x\}|;$ if $n(r, x) \ge n(\ell, x)$ then $y \leftarrow r;$ else $y \leftarrow \ell;$ for $i \in N$ do if $x, y \in A_i$ then $w_i \leftarrow 1;$ else $w_i \leftarrow \frac{\alpha+1}{\alpha-1};$ // Weighted-majority step $v_x \leftarrow \sum_{i:x \succ_i y} w_i;$ $v_y \leftarrow \sum_{i:x \succ_i y} w_i;$ if $v_x \ge v_y$ then $w \leftarrow x;$ else $w \leftarrow y;$

Proof. Anshelevich and Postl [6] first showed that one of the alternatives in $\{x, \ell, r\}$ must be the optimal alternative. In addition, they showed that ℓ cannot be better than x in terms of social cost in case $n(\ell, x) \le n(r, x)$, and r cannot be better than x in case if $n(\ell, x) > n(r, x)$. Hence, by appropriately eliminating one of ℓ or r depending on whether $n(\ell, x) \le n(r, x)$ or not, we have that the optimal alternative is one of x or y (the non-eliminated alternative among ℓ and r). Without loss of generality, we assume that y = r, and thus y is at the right of x on the line. We also assume that the mechanism chooses x as the winner, in which case, to have a distortion larger than 1, y must be the optimal alternative; clearly, the case in which the mechanism chooses y as the winner and x is the optimal alternative can be handled with similar arguments.

We partition the agents in the following sets:

- $N_{xy} = \{i : x, y \in A_i\}$ is the set of agents that approve both x and y;
- $N_{x\overline{y}} = \{i : x \in A_i, y \notin A_i\}$ is the set of agents that approve only x;
- $N_{\overline{xy}} = \{i : x \notin A_i, y \in A_i\}$ is the set of agents that approve only *y*;
- $N_{\overline{x} > \overline{y}} = \{i : x, y \notin A_i, x >_i y\}$ is the set of agents that approve none of x and y, but prefer x over y.
- $N_{\overline{y \succ x}} = \{i : x, y \notin A_i, y \succ_i x\}$ is the set of agents that approve none of x and y, but prefer y over x.

Observe that the most-preferred alternative of any agent $i \in N_{\overline{x} \succ \overline{y}}$ must be at the left of x, and thus agent i must also be at the left of x. Similarly, any agent $i \in N_{\overline{y} \succ \overline{x}}$ must be at the right of y. We argue that, in a worst-case instance, all agents lie in the interval [x, y], and thus $N_{\overline{x} \succ \overline{y}} = N_{\overline{y} \succ \overline{x}} = \emptyset$. Let i be an agent that is positioned at the left of x; clearly, agent $i \in N_{xy} \cup N_{\overline{x} \overline{y}} \cup N_{\overline{x} \succ \overline{y}}$. The distortion can be written as

$$\frac{\operatorname{SC}(x)}{\operatorname{SC}(y)} = \frac{\sum_{j \in N \setminus \{i\}} d(j, x) + d(i, x)}{\sum_{j \in N \setminus \{i\}} d(j, y) + d(i, y)}$$

Since each agent *i* is at the left of *x*, d(i, y) = d(i, x) + d(x, y), and thus

$$\frac{\mathrm{SC}(x)}{\mathrm{SC}(y)} = \frac{\sum_{j \in N \setminus \{i\}} d(j, x) + d(i, x)}{\sum_{j \in N \setminus \{i\}} d(j, y) + d(i, x) + d(x, y)} \le \frac{\sum_{j \in N \setminus \{i\}} d(j, x)}{\sum_{j \in N \setminus \{i\}} d(j, y) + d(x, y)}$$

where the inequality follows since d(i, x) appears in the numerator and the denominator of the ratio, and the distortion is, by definition, at least 1. Hence, we can obtain a worse instance in terms of the distortion by moving any potential agent that are positioned at the left of x to x; in this new instance, such agents can only approve x, and are thus all moved to the set $N_{x\overline{y}}$, leaving $N_{\overline{x} > \overline{y}}$ empty. A similar argument can be used to show that, in a worst-case instance, there are no agents at the right of y, and thus $N_{\overline{y} > \overline{x}} = \emptyset$.

Now, the distortion bound of $2\sqrt{2} - 1$ follows by the corresponding bound shown by Abramowitz et al. [2]. In their model, when there are two alternatives *a* and *b*, an agent *i* that prefers alternative *a* over *b* is said to have a preference strength of $\frac{d(i,b)}{d(i,a)}$. Abramowitz et al. considered a weighted majority mechanism that assigns weight 1 to all agents with preference strength at most τ and weight $\frac{\tau+1}{\tau-1}$ to all other agents (with preference strength strictly larger than τ), for some $\tau \ge 1$, and showed that this mechanism achieves a distortion of at most max $\{\frac{3\tau-1}{\tau+1}, 1+\frac{2}{\tau}\}$.

In the worst-case instance of our mechanism, since all agents are positioned in the interval [x, y], we have:

• For any agent $i \in N_{xy}$, the preference strength of *i* is at most α . In particular, if *i* prefers *x* over *y*,

$$d(i, y) \le \alpha d(i, x) \Leftrightarrow \frac{d(i, y)}{d(i, x)} \le \alpha.$$

Otherwise,

$$d(i, x) \le \alpha d(i, y) \Leftrightarrow \frac{d(i, x)}{d(i, y)} \le \alpha.$$

• For any agent $i \in N_{xy}$, the preference strength of *i* is strictly larger than α since

$$d(i, y) > \alpha d(i, x) \Leftrightarrow \frac{d(i, y)}{d(i, x)} > \alpha.$$

• For any agent $i \in N_{\overline{x}v}$, the preference strength of *i* is strictly larger than α since

$$d(i, x) > \alpha d(i, y) \Leftrightarrow \frac{d(i, x)}{d(i, y)} > \alpha.$$

Consequently, in the worst-case instance, the weighted majority that our mechanism runs in Step 3 (which assigns a weight of 1 to each agent in N_{xy} and a weight of $\frac{\alpha+1}{\alpha-1}$ to any other agent) is exactly the same as the weighted majority mechanism of Abramowitz et al. [2] with $\tau = \alpha$, and thus the distortion is at most $\max\{\frac{3\alpha-1}{\alpha+1}, 1+\frac{2}{\alpha}\}$.

By optimizing over α , we obtain the following corollary.

Corollary 3.9. When the metric space is a line, for the social cost objective, the distortion of $(1 + \sqrt{2})$ -ELIMINATION-WEIGHTED-MAJORITY is at most $2\sqrt{2} - 1 \approx 1.83$.

Next, we present a matching lower bound of $2\sqrt{2} - 1$ on the distortion of any mechanism that uses all three types of information, which establishes the tightness of the distortion proven above. As mentioned before, this bound already follows by the work of Abramowitz et al. [2], whose model coincides with ours when there are only two alternatives.

Theorem 3.10. When the metric space is a line, for the social cost objective, the distortion of any mechanism M with access to ORD \land DIST \land TAS is at least $2\sqrt{2} - 1$.

Proof. Let $\varepsilon > 0$ be an infinitesimal. We consider instances with two alternatives *a* and *b*, and two agents 1 and 2. If $\alpha \le 1 + \sqrt{2}$, suppose that $A_1 = \{a\}, a >_1 b, A_2 = \{b\}$, and $b >_2 a$. Given this information, the mechanism can choose any of the two alternatives as the winner, say *a*. However, the positions of the agents might be the following ones: The two alternatives are at 0 and 1, and the two agents are at $\frac{1}{\alpha+1} - \varepsilon$ and 1; it is not hard to verify that the distances are consistent with the information provided to the mechanism. Then, the social cost of *a* is approximately $\frac{2+\alpha}{1+\alpha}$, whereas the social cost of *b* is $\frac{\alpha}{1+\alpha}$. Since $\alpha \le 1 + \sqrt{2}$, the distortion of at least $1 + \frac{2}{\alpha} \ge 2\sqrt{2} - 1$.

For the case $\alpha > 1 + \sqrt{2}$, suppose that $A_1 = \{a, b\}$, $a >_1 b$, $A_2 = \{a, b\}$, and $b >_2 a$. Given this information, the mechanism can again choose any of the two alternative as the winner, say *a*. However, the positions of the agents might be the following ones: The two alternatives are at 0 and 1, and the two agents are at $1/2 - \epsilon$ and $\frac{\alpha}{1+\alpha}$; it is not hard to verify that the distances are consistent with the information provided to the mechanism. Then, the social cost of *a* is approximately $\frac{3\alpha-1}{2(\alpha-1)}$, whereas the social cost of *b* is $\frac{1+\alpha}{2(\alpha-1)}$. Since $\alpha > 1 + \sqrt{2}$, the distortion is at least $\frac{3\alpha-1}{1+\alpha} > 2\sqrt{2} - 1$.

We next consider the case of mechanisms that have access to the α -TAS and the distances between alternatives, but not the ordinal preferences. We prove that the 2-MINISUM-TAS-DISTANCE mechanism (recall its definition in Mechanism 1) achieves a distortion of 2 when the metric is a line, and this is the best possible among all mechanisms that use this type of information.

Theorem 3.11. When the metric space is a line, the distortion of α -MINISUM-TAS-DISTANCE is at most $\max\{\alpha, 1+\frac{2}{2}\}$.

Proof. Since the metric space is a line, the sets A_i define contiguous intervals. Let ℓ_i and r_i denote the leftmost and rightmost alternatives in the α -TAS A_i of agent *i*, respectively. Let *o* be an optimal alternative, and assume, without loss of generality, that *o* is to the right of the alternative *w* chosen by the mechanism.

Consider the set $L = \{i : r_i \le w\}$ of all agents for whom A_i ends before w or at w itself, and the set $R = \{i : \ell_i > w\}$ of all agents whose intervals A_i begin after w. Then, we claim that $|R| \le |L|$. To see this, consider how the quantity $\sum_{i \in N} \min_{i \in A_i} d(j, w)$

compares to $\sum_{i \in N} \min_{j \in A_i} d(j, w')$, where w' is the alternative directly to the right of w. When moving from w to w', this quantity decreases by d(w, w') for each agent in R, increases by d(w, w') for each agent in L, and stays the same for any other agent i (since w and w' belong to A_i). Thus, by the definition of the mechanism, it must be that $|R| \le |L|$, as otherwise w' would be selected instead of w.

We now proceed to prove the distortion bound. Since $|R| \le |L|$, let $f : R \to L$ be a one-to-one mapping. We make the following observations:

• For every agent $i \in R$, let $j = f(i) \in L$. If j is to the left of w, then

 $d(j, w) + d(i, w) \le d(j, w) + d(i, o) + d(w, o)$

= d(j, o) + d(i, o).

If, instead, *j* is to the right of *w*, then it must be that *w* is the closest alternative to *j*, by the definition of *L*. Therefore, since $o \notin A_j$, we have that $d(j, w) \leq \frac{1}{\alpha} d(j, o)$. Combining this with the triangle inequality, we obtain

$$\begin{aligned} d(j,w) + d(i,w) &\leq 2d(j,w) + d(j,o) + d(i,o) \\ &\leq \left(1 + \frac{2}{\alpha}\right) d(j,o) + d(i,o). \end{aligned}$$

• For every agent *i* with $w \in A_i$, we have that

$$d(i, w) \le \alpha d(i, o_i) \le \alpha d(i, o).$$

Notice that this must be true for every agent $i \notin R \cup L$.

Combing the above, we obtain

$$\begin{split} \mathrm{SC}(w) &= \sum_{i \in L} d(i, w) + \sum_{i \in R} d(i, w) + \sum_{i \notin R \cup L} d(i, w) \\ &\leq \left(1 + \frac{2}{\alpha}\right) \sum_{i \in L} d(i, o) + \sum_{i \in R} d(i, o) + \alpha \sum_{i \notin R \cup L} d(i, o) \\ &\leq \max\left\{\alpha, 1 + \frac{2}{\alpha}\right\} \cdot \mathrm{SC}(o). \end{split}$$

Consequently, the distortion is at most max $\left\{ \alpha, 1 + \frac{2}{\alpha} \right\}$.

By optimizing over α , we obtain an upper bound of 2.

Corollary 3.12. When the metric space is a line, the distortion of 2-MINISUM-TAS-DISTANCE is at most 2.

Below we prove the lower bound that establishes that 2-MINISUM-TAS-DISTANCE is best possible.

Theorem 3.13. When the metric space is a line, for the social cost objective, the distortion of any mechanism M with access to $DIST \wedge TAS$ is at least 2.

Proof. We consider instances with alternative locations 0 and $1 + \alpha$. If $\alpha \ge 2$, consider an instance where all agents include both alternatives in their α -approval sets. As there is no way of distinguishing the two alternatives, the mechanism can choose any of them as the winner. If it chooses the one at 0, then the agents may be located at α , leading to a distortion of α . If the mechanism chooses the alternative at $1 + \alpha$, then the agents may be located at 1, leading again to a distortion of α . Overall, the distortion of the mechanism is at least $\alpha \ge 2$.

For the case $\alpha < 2$, consider an instance with two agents that have $A_1 = \{0\}$ and $A_2 = \{1 + \alpha\}$. Consider the following two cases depending on the decision made by the mechanism:

- The winner is the alternative at 0: The first agent is located at 1ϵ (so that the distance from 0 is 1ϵ while the distance from $1 + \alpha$ is $\alpha + \epsilon$, leading to the alternative at 0 being the only one in A_1) and the second agent is located at $1 + \alpha$. Hence, the social cost of 0 is approximately $1 + 1 + \alpha = 2 + \alpha$, whereas the social cost of $1 + \alpha$ is approximately α , leading to a distortion of at least $1 + 2/\alpha$.
- The winner is the alternative at $1 + \alpha$: The first agent is located at 0, and the second agent is located at $\alpha + \epsilon$ (so that the distance from 0 is $\alpha + \epsilon$ while the distance from $1 + \alpha$ is 1ϵ , leading to the alternative at $1 + \alpha$ being the one in A_2). Hence, the social cost of 0 is approximately α , whereas the social cost of $1 + \alpha$ is approximately $1 + \alpha + 1 = 2 + \alpha$, leading to a distortion of at least $1 + 2/\alpha$.

Overall, the distortion of the mechanism is least $1 + 2/\alpha > 2$.

Finally, we remark that the proof of Theorem 3.7 already uses an instance that is on the line metric, so we have the following straightforward corollary.

Corollary 3.14. For the line metric, for the social cost objective, the distortion of any mechanism M with access to TAS is $\Theta(n)$.

4. Maximum cost

In this section we consider the other most natural objective, i.e., to minimize the maximum cost of any agent, abbreviated as MC. For mechanisms that have access to ORD \land DIST \land TAS we show a tight bound of $1 + \sqrt{2}$. In the real line metric, we show that there in fact exist mechanisms with access to either ORD \land TAS or DIST \land TAS that achieve the $1 + \sqrt{2}$ distortion guarantee, without having access to the third type of information. We also show that the best distortion achievable by any mechanism with access to TAS is 3.

4.1. Results for general metric spaces

We start by proving an upper bound of $1 + \sqrt{2}$ for general metric spaces. This is achieved by a novel mechanism with access to ORD \wedge DIST \wedge TAS, which we refer to as α -MOST-COMPACT-SET, for $\alpha = 1 + \sqrt{2}$. In fact, the mechanism requires a bit less information, as it only uses the most-preferred alternatives of the agents, rather than their complete ordinal preferences.

Before we define the mechanism, we will prove the following useful lemma relating the choices of any mechanism with the distortion for the maximum cost objective. To differentiate between alternatives that lie at minimum distance from an agent *i* and alternatives that appear first in the preference ranking of the agent (which might be different due to ties), we will refer to the former as *agent i's min-distance alternative* (rather than agent *i*'s most-preferred alternative). In other words, a min-distance alternative for agent *i* is an alternative x for which $d(i, x) \le d(i, x')$ for any $x' \in A$.

Lemma 4.1. Consider a mechanism M that chooses the winner to be an alternative w that satisfies at least one of the following two conditions:

- 1. $w \in A_i$ for all $j \in N$.
- 2. w is a min-distance alternative for some agent $i \in N$, and $o \notin A_i$, where o is an alternative that minimizes the maximum cost.

Then, for the maximum cost objective, the distortion of M is at most $\max\{\alpha, 2 + \frac{1}{\alpha}\}$.

Proof. For the first condition, if $w \in A_j$ for all $j \in N$, then $d(j, w) \le \alpha \cdot d(j, o_j) \le \alpha \cdot d(j, o)$, where o_j is the most-preferred alternative of agent j. This implies that $MC(w) \le \alpha \cdot MC(o)$. For the second condition, let t be an agent that is at maximum distance from w. We have that

$$\begin{split} \mathsf{MC}(w) &= d(t, w) \leq d(t, o) + d(o, i) + d(i, w) \\ &< d(t, o) + d(i, o) + \frac{1}{\alpha} \cdot d(i, o) \\ &\leq \left(2 + \frac{1}{\alpha}\right) \mathsf{MC}(o), \end{split}$$

where the first inequality is due to the triangle inequality, and the second inequality follows from the fact that $o \notin A_i$.

Now, our mechanism, coined α -MOST-COMPACT-SET, works as follows: It first checks if there exists some alternative x that is approved by all agents, and, if this is indeed the case, it then chooses x as the winner w. Otherwise, for each agent i, it computes the maximum distance ρ_i of the most-preferred alternative o_i of i to any other alternative in A_i ; then it chooses the most-preferred alternative o_i of the agent i with minimum ρ_i . Essentially, in the otherwise case, the mechanism aims to choose the most-preferred alternative of an agent in order to minimize the radius to all other alternatives approved by the agent. See Mechanism 3 for a description using pseudocode.

Theorem 4.2. In general metric spaces, for the max cost objective, the distortion of α -MOST-COMPACT-SET is at most $\max\{\alpha, 2+\frac{1}{\alpha}\}$.

Proof. If there is an alternative *w* such that $w \in A_j$ for all $j \in N$, the mechanism chooses it as the winner *w*. In that case, by Lemma 4.1, the distortion of α -MOST-COMPACT-SET is at most max $\{\alpha, 2 + \frac{1}{\alpha}\}$. If such an alternative does not exist, then the mechanism chooses the most-preferred alternative of some agent *i*, in particular the agent with the smallest ρ_i , i.e., the smallest maximum distance from any of her approved alternatives to her most-preferred alternative o_i . If $o \notin A_i$, then by Lemma 4.1, the distortion of the mechanism is again at most max $\{\alpha, 2 + \frac{1}{\alpha}\}$. Therefore we will consider the case when $o \in A_i$. By definition of ρ_i , this implies that $d(o, o_i) \leq \rho_i$.

Mechanism 3: α-MOST-COMPACT-SET.

Input: Most-preferred alternatives of agents, distances between alternatives, α -TAS; Output: Winner w; if $\exists x \in A_i, \forall i \in N$ then $\psi \leftarrow x$; else for $i \in N$ do $\rho_i \leftarrow \max_{x \in A_i} d(x, o_i)$; $w \in \arg\min_{i \in N} \rho_i$;

Since there is no alternative x such that $x \in A_i$ for all $i \in N$, there must exist at least one agent j such that $o \notin A_j$. Also, let t be the agent that has the maximum distance from w. We have that

 $\mathrm{MC}(w) = d(t, o_i) \le d(t, o) + d(o, o_i) \le \mathrm{MC}(o) + \rho_i \le \mathrm{MC}(o) + \rho_i.$

where the last inequality follows from that fact that *i* was chosen to be an agent that minimizes $\rho_k = \max_{x \in A_k} d(x, o_k)$ by the mechanism. Let *y* be an alternative in A_i with maximum distance from o_i . Then,

$$\rho_j = d(y, o_j) \le d(j, o_j) + d(j, y) \le \left(1 + \frac{1}{\alpha}\right) \operatorname{MC}(o),$$

where the first inequality is due to the triangle inequality, and the second one follows since $d(j, o_j) \le d(j, o) \le MC(o)$, and since $o \notin A_j$ means that $d(j, o_j) < \frac{1}{\alpha} MC(o)$. Putting everything together, we obtain

$$\operatorname{MC}(w) \le \left(2 + \frac{1}{\alpha}\right) \operatorname{MC}(o).$$

Overall, the distortion of the mechanism is at most $\max\{\alpha, 2 + \frac{1}{\alpha}\}$, as desired.

By optimizing over α , we obtain the following corollary.

Corollary 4.3. In general metric spaces, for the max cost objective, the distortion of $(1 + \sqrt{2})$ -MOST-COMPACT-SET is at most $1 + \sqrt{2}$.

Next, we prove that the bound of $1 + \sqrt{2}$ is the best possible for the maximum cost objective, even when all three types of information are available, and even when the metric space is as simple as a line.

Theorem 4.4. For the max cost, the distortion of any mechanism M with access to $ORD \wedge DIST \wedge TAS$ is at least $1 + \sqrt{2}$, even on a line metric.

Proof. Consider any mechanism M that has access to all three types of information. To show the lower bound, we will consider the following two cases depending on the value of α that M uses.

Case 1: $\alpha \le 1 + \sqrt{2}$. Consider an instance I_1 with two agents $\{1,2\}$ and two alternatives $\{a_1, a_2\}$. We have $A_1 = \{a_1\}$, $A_2 = \{a_2\}$, $a_1 >_1 a_2$, and $a_2 >_2 a_1$. The distance between the two alternatives is equal to $1 + \alpha + \varepsilon$, for some infinitesimal $\varepsilon > 0$. Given this information, M may choose any of the two alternatives as the winner, say a_1 . We define the following line metric (see Fig. 5a):

- Alternative a_1 is at 0;
- Agent 1 is at 1;
- Alternative a_2 is at $1 + \alpha + \varepsilon$;
- Agent 2 is at $1 + 2\alpha + \epsilon$.

It is not hard to verify that these positions of the agents and the alternatives are consistent with the information given to the mechanism for any $\alpha \le 1 + \sqrt{2}$: Agent 2 approves only alternative a_2 as $d(2, a_1) = 1 + 2\alpha + \epsilon > \alpha^2 = \alpha \cdot d(2, a_2)$; all other information is clearly consistent. We have that $MC(a_1) = 1 + 2\alpha + \epsilon$ and $MC(a_2) = \alpha + \epsilon$, thus obtaining a lower bound of $2 + 1/\alpha \ge 1 + \sqrt{2}$ when ϵ tends to 0.

Case 2: $\alpha > 1 + \sqrt{2}$. Consider an instance I_2 with two agents $\{1,2\}$ and two alternatives $\{a_1, a_2\}$. We have $A_1 = A_2 = \{a_1, a_2\}$, $a_1 >_1 a_2$ and $a_2 >_2 a_1$. The distance between the two alternatives is equal to 1. Given this information, M may choose any of the two alternatives as the winner, say a_1 . We define the following line metric (see Fig. 5b):

- Alternative a_1 is at 0;
- Agent 1 is at 1/2;





Fig. 5. The line metrics used in the proof of Theorem 4.4. Agents are represented by rectangles and alternatives are represented by circles. A value above an edge represents the distance between the edge's endpoints.

- Alternative a_2 is at 1;
- Agent 2 is at min{ $3/2, \frac{\alpha}{\alpha-1}$ }.

It is not hard to verify that these positions of the agents and the alternatives are consistent with the information given to the mechanism for any $\alpha > 1 + \sqrt{2}$: Agent 1 is at distance 1/2 from both a_1 and a_2 . When $\frac{\alpha}{\alpha-1} \le 3/2 \Leftrightarrow \alpha \le 3$, agent 2 is at distance $\frac{\alpha}{\alpha-1} - 1 = \frac{1}{\alpha-1}$ from a_2 and distance $\frac{\alpha}{\alpha-1}$ from a_1 ; clearly, the ratio of these distances is exactly α . When $\frac{\alpha}{\alpha-1} > 3/2 \Leftrightarrow \alpha > 3$, agent 2 is at distance 1/2 from a_2 and distance 3/2 from a_1 ; the ratio of these distances is exactly α . When $\frac{\alpha}{\alpha-1} > 3/2 \Leftrightarrow \alpha > 3$, agent 2 is at distance 1/2 from a_2 and distance 3/2 from a_1 ; the ratio of these distances is exactly 3, and thus at most α . Since $\frac{1}{\alpha-1} < 1/2 < \frac{\alpha}{\alpha-1}$ for any $\alpha > 1$, we have that $MC(a_1) = \min\{3/2, \frac{\alpha}{\alpha-1}\}$ and $MC(a_2) = 1/2$, thus obtaining a lower bound of 3 when $\alpha > 3$ and $\frac{2\alpha}{\alpha-1} \ge \alpha \ge 1 + \sqrt{2}$ when $\alpha \in (1 + \sqrt{2}, 3]$.

We conclude the exposition of our results for general metric spaces by showing that the very simple mechanism ANY-APPROVED which chooses any alternative in the α -TAS of some agent achieves a distortion bound of $2 + \alpha$ in terms of the maximum cost objective. By setting $\alpha = 1$, we obtain a distortion bound of 3, which is the best possible for all mechanisms with access to TAS, and also matches the bound of 3 that is best possible for any mechanism with access to ORD \wedge DIST.

Theorem 4.5. In general metric spaces, for the max cost objective, the distortion of ANY-APPROVED is at most $2 + \alpha$.

Proof. Consider an arbitrary instance and let w be the alternative chosen by M on this instance. By definition, w is in the α -TAS of some i. Let j be an agent that has the maximum distance from w. Let o be an optimal alternative, i.e., the one that minimizes the maximum cost. Let o_i be the alternative that is closest to the position of agent i. We have that

$$\begin{split} \mathrm{MC}(w) &= d(j,w) \leq d(j,o) + d(i,o) + d(i,w) \\ &\leq 2\mathrm{MC}(o) + \alpha \cdot d(i,o_i) \\ &\leq (2+\alpha)\mathrm{MC}(o), \end{split}$$

where the first inequality holds by the triangle inequality, the second inequality follows from the definition of *o* and the fact that $w \in A_i$, and the last inequality follows since $d(i, o_i) \le d(i, o)$ by the definition of o_i .

The tight lower bound for all mechanisms with access to TAS is given below.

Theorem 4.6. For the max cost objective, the distortion of any mechanism M with access to TAS is at least 3, even when the metric space is a line.

Proof. Consider an arbitrary mechanism M, and an instance with two agents $\{1, 2\}$, and five alternatives $\{a_1, \dots, a_5\}$. That α -TAS of the agents are $A_1 = \{a_1, a_2\}$ and $A_2 = \{a_3, a_4\}$. Let ε be an infinitesimal. If M chooses alternative a_5 as the winner, then the distortion can easily be seen to be unbounded as the metric space might be the following:

- Agent 1 and alternatives *a*₁, *a*₂ are at 0;
- Agent 2 and alternatives a_3, a_4 are at ϵ ;
- Alternative a_5 is at 1.

Hence, the optimal max cost is ε , whereas the max cost of the mechanism is 1.

So, suppose that *M* selects one of the alternatives that is approved by one of the agents, say $w = a_1$. We define the following line metric (see also Fig. 6):

```
• Agent 1 is at \alpha + \varepsilon;
```



Fig. 6. The line metric used in the proof of Theorem 4.6.

- Alternative a_2 is at $\alpha + 1 + \varepsilon$;
- Alternative a_1 is at $2\alpha + \epsilon$
- Agent 2 is at $-\alpha \varepsilon$;
- Alternative a_3 is at $-\alpha 1 \varepsilon$;
- Alternative a_4 is at $-2\alpha \varepsilon$;
- Alternative a_5 is at 0.

It is not hard to verify that the positions of the agents and the alternatives on the line are consistent with the α -TAS: Agent 1 is at distance 1 from a_2 , α from a_1 , and strictly larger than α from any other alternative. Similarly, agent 2 is at distance 1 from a_3 , α from a_4 , and strictly larger than α from any other alternative. The cost of the mechanism is $MC(w) = d(2, a_1) \approx 3\alpha$, whereas the optimal cost if $MC(a_5) = d(1, a_5) = d(2, a_5) \approx \alpha$, leading to a distortion of at least 3.

4.2. Refined upper bounds for the line metric

The lower bounds of $1 + \sqrt{2}$ for mechanisms with access to all three types of information (Theorem 4.4) and 3 for mechanisms in with access to only TAS (Theorem 4.6) already apply when the metric space is a line. For both cases, tight upper bounds can be achieved for general spaces (Theorem 4.2 and Theorem 4.5). In this section, we show that, when the metric space is a line, the bound of $1 + \sqrt{2}$ can be achieved by several mechanisms with access to ORD \wedge TAS and DIST \wedge TAS (i.e., without having access to the third type of information). Whether such mechanisms exist for general metric spaces is an intriguing open problem.

From a technical perspective, these mechanisms, although different in nature, establish the same kind of properties that are needed for the distortion guarantees to hold. For this reason, we first present all these mechanisms, and then prove their distortion bounds with one, unified proof. Before that, we start with a simple known observation for the maximum cost objective on the line.

Observation 4.7. Any mechanism that always chooses the winner to be an alternative that lies in the interval defined by the leftmost and rightmost agents has distortion at most 2.

Proof. Let ℓ' and r be the positions of the leftmost and rightmost agent, respectively. Clearly, the max cost of the mechanism is $MC(w) \le d(\ell, r)$, whereas the optimal max cost is $MC(o) \ge \frac{1}{2} \cdot d(\ell, r)$, and thus the distortion is at most 2.

We now present the mechanisms, starting with one that uses ORD \wedge TAS, to which we refer as TAS-OPTIMAL. The mechanism chooses an (arbitrary) alternative that is *TAS-optimal*, a notion of dominance reminiscent of Pareto dominance, adapted to take the α -TAS into account.

Definition 4.8 (*TAS-optimality*). An alternative *x TAS-dominates* an alternative *y* if for all agents $i \in N$ such that $\{x, y\} \nsubseteq A_i$, we have $x >_i y$. An alternative *x* that is not TAS-dominated by any other alternative is called *TAS-optimal*.

In other words, *x* TAS-dominates *y* if *x* is preferred over *y* by all agents that do not consider both of them acceptable. Clearly, a TAS-optimal alternative can be found using the α -TAS and the ordinal preferences of the agents.

The next two mechanisms we consider require access to DIST \wedge TAS. The first of them is the α -MINISUM-TAS-DISTANCE that we defined and used in Section 3.1, which chooses the winner w to minimize the sum of distances from all α -TAS; recall its definition in Mechanism 1. The second one is a mechanism to which we refer as α -MINIMAX-TAS-DISTANCE. This mechanism is the natural translation of α -MINISUM-TAS-DISTANCE to the maximum cost objective, as it chooses w to minimize the maximum distance from any α -TAS. See Mechanism 5 for its description.

Distortion upper bounds

We now present a proof that applies to all the mechanisms defined above at the same time, by establishing a series of claims. For brevity, we will refer to the mechanisms as M_1 for TAS-OPTIMAL, M_2 for α -MINISUM-TAS-DISTANCE, and M_3 for α -MINIMAX-TAS-DISTANCE.

Theorem 4.9. When the metric space is a line, for the maximum cost objective, the distortion of any mechanism $M \in \{M_1, M_2, M_3\}$ is at most $\max\{\alpha, 2 + \frac{1}{\alpha}\}$.

Proof. Assume by contradiction that the distortion of *M* is larger than $\max\{\alpha, 2 + \frac{1}{\alpha}\}$. Consider an arbitrary instance, and let i_1 and i_n be the positions of the leftmost and rightmost agent on the line respectively (breaking ties arbitrarily). Let *w* be the alternative chosen by *M* for the given instance, and let *o* be an optimal alternative that minimizes the maximum cost. Clearly, $w \neq o$, as, otherwise,

Artificial Intelligence 341 (2025) 104295

Mechanism 4: TAS-OPTIMAL.	
Input: Ordinal preferences, α -TAS;	
Output: Winner <i>w</i> ;	
$w \leftarrow$ any TAS-optimal alternative.	

Mechanism 5: α -MINIMAX-TAS-DISTANCE.				
Input: Distances between alternatives, α -TAS;				
Output : Winner <i>w</i> ;				
$w \in \arg\min_{x \in A} \left\{ \max_{i \in N} \min_{j \in A_i} d(j, x) \right\};$				

the distortion of M would be 1. Additionally, $w \notin [i_1, i_n]$, as, otherwise, by Observation 4.7 the distortion of M would be at most 2. Without loss of generality, we assume that $w > i_n$, and thus o < w. Finally, there exists some agent $j \in N$ such that $w \notin A_j$, as, otherwise, the distortion of M would be most $\max\{\alpha, 2 + \frac{1}{\alpha}\}$ by Lemma 4.1.

For any alternative x, let $\ell(x)$ denote the alternative directly to the left of x on the line, which does not coincide with x (breaking ties arbitrarily). We state and prove the following useful claim.

Claim 4.10. *w* is a min-distance alternative for agent i_n .

Proof. Assume by contradiction that w is not a min-distance alternative for i_n . Then $\ell(w)$ has to be a min-distance alternative for agent i_n ; note that such an alternative exists since o < w. This implies that $\ell(w) >_{i_n} w$, which also implies that $\ell(w) >_i w$ for all $i \in N$, since i_n is the rightmost agent and $w > i_n$. Since there exists some agent $j \in N$ such that $w \notin A_j$, it follows that $\ell(w)$ TAS-dominates w. This establishes the proof of the claim for $M = M_1$, since it contradicts the choice of w as the winner by the mechanism.

Next we prove the claim for $M \in \{M_2, M_3\}$. Observe that for every agent $i \le i_n$ we have that $w \in A_i \Rightarrow \ell(w) \in A_i$, since $w > \ell(w)$. In other words, all agents that find w acceptable also find $\ell(w)$ acceptable; for those agents we have $d(\ell(w), A_j) = d(w, A_j) = 0$. Now consider any agent $j \in N$ such that $w \notin A_j$ and let $S^{\neg w}$ be the set of those agents. Since w is not acceptable for all agents, we have that $|S^{\neg w}| \ge 1$. For any agent $j \in S^{\neg w}$, let x_r^j be the rightmost alternative smaller than w in A_j (breaking ties arbitrarily), and observe that $x_r^j \le \ell(w)$. We have that $d(x_r^j, \ell(w)) < d(x_r^j, w)$. Since this holds for every agent $j \in S^{\neg w}$, this means that $\ell(w)$ has both a smaller total and maximum distance from the α -TAS, contradicting the choice of w by the mechanism. \Box

Given Claim 4.10, by Lemma 4.1, it suffices to prove that $o \notin A_{i_n}$, as, otherwise, the distortion would be upper bounded by $\max\{2 + \frac{1}{n}\}$, leading to a contradiction. This is established by the following claim, the proof of which is similar to that of Claim 4.10.

Claim 4.11. $o \notin A_{i_n}$.

Proof. Assume by contradiction that $o \in A_{i_n}$. Let S^w be the set of agents for which w is a min-cost alternative, and notice that $i_n \in S^w$ be Claim 4.10. From the above, it follows that for every agent $i \in S^w$, and for any alternative $x \in [o, w]$ we have that $x \in A_i$; in particular, $\ell(w) \in A_i$, which implies that $\{\ell(w), w\} \subseteq A_i$ for those agents. For agents $i \notin S^w$, it holds that $d(i, \ell(w)) < d(i, w)$, as (a) either $i \ge \ell(w)$ and $\ell(w)$ is a min-cost alternative for i, or (b) $i < \ell(w)$. This clearly implies that $\ell(w) >_i w$ for each agent $i \notin S^w$. Overall, each agent $i \in N$ either finds both w and $\ell(w)$ acceptable, or prefers $\ell(w)$ to w in her preference ranking. Since there exists some agent that does not find w acceptable, $\ell(w)$ TAS-dominates w. For the case of $M = M_1$, this contradicts the choice of w as the winner by the mechanism. The proof for $M \in \{M_2, M_3\}$ is identical to that of Claim 4.10, as, again, we have that for any agent $j \in N$, $w \in A_j \Rightarrow \ell(w) \in A_j$, similarly to Claim 4.10. \Box

This completes the proof.

By optimizing over α , we obtain the following corollary, establishing that all these mechanisms are best possible when the metric space is a line.

Corollary 4.12. When the metric space is a line, for the maximum cost objective, the distortion of any mechanism $M \in \{M_1, M_2, M_3\}$ using $\alpha = 1 + \sqrt{2}$ is at most $1 + \sqrt{2}$.

Remark 4.13. One can obtain the distortion bound of Theorem 4.9 using a mechanism with access to DIST \land TAS by appropriately modifying a mechanism of Anshelevich et al. [13]. Anshelevich et al. considered a different setting, that of *distributed metric social choice*, where the agents are partitioned into districts, and local decisions within districts are then aggregated at a global level. In their setting, they define a mechanism called α -ACCEPTABLE-RIGHTMOST-LEFTMOST (α -ARL), which first selects, for each district, the rightmost alternative that the agents therein find " α -acceptable" and then, the leftmost among these alternatives is chosen as the winner. If we consider districts of size 1 in their setting, α -acceptability reduces to acceptability according to the α -TAS in our

setting. Note that their mechanism requires the ordering of the alternatives on the line to be able to identify the rightmost and leftmost of them. We can infer this information from the distances between alternatives on the line, and then appropriately modify their mechanism. Theorem 4.8 from their work would then yield the upper bound of our Theorem 4.9. Still, we believe that the mechanisms that we consider in this section are much more natural for our setting and much easier to describe. \Box

5. Future directions

The main problem that our work leaves open is that of identifying the best possible distortion bound in terms of the social cost for general metric spaces and mechanisms that use all three types of information, for which we showed an upper bound of $1 + \sqrt{2}$ and a lower bound of 2. Another interesting open question is to show whether $1 + \sqrt{2}$ can be achieved with a mechanism with access to ORD \wedge TAS for general metric spaces. For the maximum cost objective, we were able to obtain tight bounds for the general class of mechanisms, so the open question is whether these bounds can be achieved by mechanisms with access to ORD \wedge TAS and DIST \wedge TAS, as we proved to be the case on the real line. For DIST \wedge TAS, the $(1 + \sqrt{2})$ -MINIMAX-TAS-DISTANCE mechanism, which we showed to achieve the tight bound of $1 + \sqrt{2}$ on the line, seems like a very natural candidate. Unfortunately, however, we can construct counterexamples where the distortion of this mechanism is lower-bounded by 3 in general metric spaces. For ORD \wedge TAS, TAS-OPTIMAL falls short once considered in general metric spaces, as similar examples to the aforementioned one show a lower bound of 3 on its distortion.

More generally, our paper advocates the study of the interplay between information and efficiency in the metric social choice setting, building on a growing literature that has recently gained significant momentum. As a next step, one could consider eliciting additional cardinal information, for example, by using multiple approval thresholds $\alpha_1, \ldots, \alpha_k$, in conjunction with the other types of information that we use here, and quantify the effect on the distortion. More generally, it would make sense to consider different types of information structures and combinations between them; note, for example, that our α -MOST-COMPACT-SET mechanism for the maximum cost uses the distances between alternatives, the α -TAS, and information about the top-ranked alternative of each agent. Finally, another very meaningful avenue could be to consider how randomization might aid in achieving even more improved distortion bounds in the presence of the α -TAS, possibly by building on top of the recent breakthrough result of Charikar et al. [23] who showed a randomized constant upper bound strictly smaller than 3 for when only ordinal information about the preferences of the agents is available.

CRediT authorship contribution statement

Elliot Anshelevich: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. Aris Filos-Ratsikas: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. Christopher Jerrett: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. Alexandros A. Voudouris: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization.

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.

Acknowledgements

Aris Filos-Ratsikas was supported by EPSRC Grant EP/Y003624/1 "Algorithms and Complexity for Economic Environments (ACEE)". We would like to thank the anonymous reviewers for their valuable feedback.

Data availability

No data was used for the research described in the article.

References

- 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] Ben Abramowitz, Elliot Anshelevich, Wennan Zhu, Awareness of voter passion greatly improves the distortion of metric social choice, in: Proceedings of the 15th International Conference on Web and Internet Economics (WINE), 2019, pp. 3–16.
- [3] 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.
- [4] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros A. Voudouris, A few queries go a long way: information-distortion tradeoffs in matching, J. Artif. Intell. Res. 74 (2022) 227–261.
- [5] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros A. Voudouris, Don't roll the dice, ask twice: the two-query distortion of matching problems and beyond, SIAM J. Discrete Math. 38 (1) (2024) 1007–1029.
- [6] Elliot Anshelevich, John Postl, Randomized social choice functions under metric preferences, J. Artif. Intell. Res. 58 (2017) 797–827.

- [7] Elliot Anshelevich, Shreyas Sekar, Blind, greedy, and random: algorithms for matching and clustering using only ordinal information, in: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI), 2016, pp. 390–396.
- [8] Elliot Anshelevich, Shreyas Sekar, Truthful mechanisms for matching and clustering in an ordinal world, in: Proceedings of the 12th Conference on Web and Internet Economics (WINE), 2016, pp. 265–278.
- [9] Elliot Anshelevich, Wennan Zhu, Tradeoffs between information and ordinal approximation for bipartite matching, in: Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT), 2017, pp. 267–279.
- [10] Elliot Anshelevich, Wennan Zhu, Ordinal approximation for social choice, matching, and facility location problems given candidate positions, ACM Trans. Econ. Comput. 9 (2) (2021) 1–24.
- [11] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, Piotr Skowron, Approximating optimal social choice under metric preferences, Artif. Intell. 264 (2018) 27–51.
- [12] 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, 2021, pp. 4294–4301.
- [13] Elliot Anshelevich, Aris Filos-Ratsikas, Alexandros A. Voudouris, The distortion of distributed metric social choice, Artif. Intell. 308 (2022) 103713.
- [14] Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, Nisarg Shah, Preference elicitation for participatory budgeting, Manag. Sci. 67 (5) (2021) 2813–2827.
- [15] Umang Bhaskar, Varsha Dani, Abheek Ghosh, Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32, 2018.
- [16] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, Or Sheffet, Optimal social choice functions, Artif. Intell. 227 (C) (2015) 190–213.
- [17] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia, Handbook of Computational Social Choice, Cambridge University Press, 2016.
- [18] Ioannis Caragiannis, Ariel D. Procaccia, Nisarg Shah, Truthful univariate estimators, in: International Conference on Machine Learning, 2016, pp. 127–135.
- [19] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, Nisarg Shah, Subset selection via implicit utilitarian voting, J. Artif. Intell. Res. 58 (2017) 123–152.
 [20] Ioannis Caragiannis, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Kristoffer Arnsfelt Hansen, Zihan Tan, Truthful facility assignment with resource augmentation: an exact analysis of serial dictatorship, Math. Program. (2022) 1–30.
- [21] Joannis Caragiannis, Aris Filos-Ratsikas, Swaprava Nath, Alexandros A. Voudouris, Truthful ownership transfer with expert advice, Math. Program. (2022) 1–30.
 [22] Joannis Caragiannis, Nisarg Shah, Alexandros A. Voudouris, The metric distortion of multiwinner voting, in: Proceedings of the 36th AAAI Conference on Artificial
- Intelligence (AAAI), 2022, pp. 4900–4907.
- [23] Moses Charikar, Prasanna Ramakrishnan, Kangning Wang, Hongxun Wu, Breaking the metric voting distortion barrier, in: Proceedings of the 35th ACM-SIAM Symposium on Discrete Algorithms (SODA), 2024, pp. 1621–1640.
- [24] Yiling Chen, Chara Podimata, Ariel D. Procaccia, Nisarg Shah, Strategyproof linear regression in high dimensions, in: Proceedings of the 2018 ACM Conference on Economics and Computation, 2018, pp. 9–26.
- [25] Yiling Chen, Yiheng Shen, Shuran Zheng, Truthful data acquisition via peer prediction, Adv. Neural Inf. Process. Syst. 33 (2020) 18194–18204.
- [26] Cynthia Dwork, Ravi Kumar, Moni Naor, Dandapani Sivakumar, Rank aggregation methods for the web, in: Proceedings of the 10th International Conference on World Wide Web, 2001, pp. 613–622.
- [27] Soroush Ebadian, Anson Kahng, Dominik Peters, Nisarg Shah, Optimized distortion and proportional fairness in voting, in: Proceedings of the 23rd ACM Conference on Economics and Computation (EC), 2022, pp. 563–600.
- [28] Edith Elkind, Piotr Faliszewski, Recognizing 1-euclidean preferences: an alternative approach, in: Proceedings of the 7th International Symposium on Algorithmic Game Theory, 2014, pp. 146–157.
- [29] Eithan Ephrati, Jeffrey S. Rosenschein, The Clarke tax as a consensus mechanism among automated agents, in: AAAI, vol. 91, 1991, pp. 173–178.
- [30] Brandon Fain, Ashish Goel, Kamesh Munagala, Nina Prabhu, 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.
- [31] 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.
- [32] Aris Filos-Ratsikas, Peter Bro Miltersen, Truthful approximations to range voting, in: Proceedings of the 10th International Conference on Web and Internet Economics (WINE), 2014, pp. 175–188.
- [33] 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.
- [34] Aris Filos-Ratsikas, Evi Micha, Alexandros A. Voudouris, The distortion of distributed voting, Artif. Intell. 286 (2020) 103343.
- [35] Aris Filos-Ratsikas, Panagiotis Kanellopoulos, Alexandros A. Voudouris, Rongsen Zhang, The distortion of distributed facility location, Artif. Intell. 328 (2024) 104066.
- [36] Mohammad Ghodsi, Mohamad Latifian, Masoud Seddighin, On the distortion value of the elections with abstention, J. Artif. Intell. Res. 70 (2021) 567–595.
- [37] Sumit Ghosh, Manisha Mundhe, Karina Hernandez, Sandip Sen, Voting for movies: the anatomy of a recommender system, in: Proceedings of the Third Annual Conference on Autonomous Agents, 1999, pp. 434–435.
- [38] 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.
- [39] Stephen Gross, Elliot Anshelevich, Lirong Xia, Vote until two of you agree: mechanisms with small distortion and sample complexity, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 31, 2017.
- [40] Anson Kahng, Min Kyung Lee, Ritesh Noothigattu, Ariel Procaccia, Christos-Alexandros Psomas, Statistical foundations of virtual democracy, in: International Conference on Machine Learning (ICML), PMLR, 2019, pp. 3173–3182.
- [41] Anson Kahng, Gregory Kehne, Ariel Procaccia, Strategyproof mean estimation from multiple-choice questions, in: International Conference on Machine Learning (ICML), 2020, pp. 5042–5052.
- [42] David Kempe, Communication, distortion, and randomness in metric voting, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, 2020, pp. 2087–2094.
- [43] Fatih Erdem Kizilkaya, David Kempe, Plurality veto: a simple voting rule achieving optimal metric distortion, in: Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI), 2022, pp. 349–355.
- [44] Fatih Erdem Kizilkaya, David Kempe, Generalized veto core and a practical voting rule with optimal metric distortion, in: Proceedings of the 24th ACM Conference on Economics and Computation (EC), 2023, pp. 913–936.
- [45] Mohamad Latifian, Alexandros A. Voudouris, The distortion of threshold approval matching, in: Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI), 2024, pp. 2851–2859.
- [46] 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.
- [47] Debmalya Mandal, Nisarg Shah, David P. Woodruff, Optimal communication-distortion tradeoff in voting, in: Proceedings of the 21st ACM Conference on Economics and Computation, 2020, pp. 795–813.
- [48] Andrew Mao, Ariel D. Procaccia, Yiling Chen, Social choice for human computation, in: Workshops at the 26th AAAI Conference on Artificial Intelligence, 2012.
- [49] 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.

- [50] Ritesh Noothigattu, Snehalkumar Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, Ariel Procaccia, A voting-based system for ethical decision making, in: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), vol. 32, 2018.
- [51] Dominik Peters, Ariel D. Procaccia, Alexandros Psomas, Zixin Zhou, Explainable voting, Adv. Neural Inf. Process. Syst. 33 (2020) 1525–1534.
- [52] Jannik Peters, A note on rules achieving optimal metric distortion, CoRR, arXiv:2305.08667, 2023.
- [53] Ariel D. Procaccia, Jeffrey S. Rosenschein, The distortion of cardinal preferences in voting, in: Cooperative Information Agents X, 10th International Workshop, CIA 2006 Edinburgh, UK, September 11-13, 2006 Proceedings 10, Springer, 2006, pp. 317–331.
- [54] John Von Neumann, Oskar Morgenstern, Theory of games and economic behavior, in: Theory of Games and Economic Behavior, Princeton University Press, 2007.
- [55] Alexandros A. Voudouris, Tight distortion bounds for distributed metric voting on a line, Oper. Res. Lett. 51 (2023) 266-269.
- [56] Lirong Xia, Group decision making under uncertain preferences: powered by AI, empowered by AI, Ann. N.Y. Acad. Sci. 1511 (1) (2022) 22–39.
- [57] Zhibing Zhao, Lirong Xia, Learning mixtures of Plackett-Luce models from structured partial orders, Adv. Neural Inf. Process. Syst. 32 (2019).