

Settling the Distortion of Distributed Facility Location

Aris Filos-Ratsikas
University of Edinburgh
Edinburgh, United Kingdom
aris.filos-ratsikas@ed.ac.uk

Panagiotis Kanellopoulos
University of Essex
Colchester, United Kingdom
panagiotis.kanellopoulos@essex.ac.uk

Alexandros A. Voudouris
University of Essex
Colchester, United Kingdom
alexandros.voudouris@essex.ac.uk

Rongsen Zhang
University of Essex
Colchester, United Kingdom
rz19109@essex.ac.uk

ABSTRACT

We study the *distributed facility location problem*, where a set of agents with positions on the line of real numbers are partitioned into disjoint districts, and the goal is to choose a point to satisfy certain criteria, such as optimize an objective function or avoid strategic behavior. A mechanism in our distributed setting works in two steps: For each district it chooses a point that is representative of the positions reported by the agents in the district, and then decides one of these representative points as the final output. We consider two classes of mechanisms: *Unrestricted* mechanisms which assume that the agents directly provide their true positions as input, and *strategyproof* mechanisms which deal with strategic agents and aim to incentivize them to truthfully report their positions. For both classes, we show tight bounds on the best possible approximation in terms of several minimization social objectives, including the well-known social cost (total distance of agents from chosen point) and max cost (maximum distance among all agents from chosen point), as well as other fairness-inspired objectives that are tailor-made for the distributed setting.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic mechanism design; Design and analysis of algorithms.**

KEYWORDS

Facility location; Distortion; Strategyproofness

ACM Reference Format:

Aris Filos-Ratsikas, Panagiotis Kanellopoulos, Alexandros A. Voudouris, and Rongsen Zhang. 2023. Settling the Distortion of Distributed Facility Location. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

The theory of social choice deals with the fundamental question of how to aggregate the opinions or *preferences* of diverse individuals into a collective decision. The quality of such a social decision can be measured in several ways, such as based on axiomatic properties, as is usually the case in economics, or qualitative metrics, an approach mainly stemming from the literature in computer science. The most

prominent such metric is that of *distortion* [36], which captures precisely the (in)efficiency of a social choice rule, or a class of such rules that often operate under some restrictions, such as the lack of expressive elicitation of the preferences of the agents.

The distortion of social choice rules (or *mechanisms*) has been a focal point of research over the past decade, for many different settings; see the recent survey of Anshelevich et al. [7] for an overview. The vast majority of previous works assume a basic setting in which a set of agents have cardinal (i.e., numerical) preferences over a set of possible outcomes (alternatives), and the goal is to quantify the best possible distortion of mechanisms that are given as input limited information about the preferences of the agents (usually rankings that are consistent with the cardinal values) in terms of the *social welfare* objective, the total value of the agents for the chosen outcome [6, 11, 18, 27].

In many cases, however, the situation is often not that simple. For example, in elections, the agents (now voters) are naturally or artificially partitioned into districts, which elect their representatives, and based on these representatives only a final winner is chosen. More generally, the decision-making process is often *distributed*, in the sense that decisions are first made at a local level, among disjoint sets of agents, and then these decisions are aggregated into a collective outcome. These types of situations are not captured by the simple setting laid out above, and bring forward important challenges and complications when measuring the efficiency of social choice mechanisms.

To capture problems of a more complex nature like the ones mentioned above, Filos-Ratsikas et al. [24] initiated the study of the distortion in *distributed social choice*, where decisions are made by mechanisms that operate as follows: The mechanism first chooses a representative alternative for each district according to a local election with the agents of the district as voters, and then chooses one of the representatives as the winner. In their work, Filos-Ratsikas et al. [24] considered a setting with agents that have normalized cardinal valuations over the possible outcomes. In follow-up work, [8] studied the same question in the very popular *metric social choice* setting, where agents and alternatives are points on a metric space, and distances capture either physical or ideological distances along different axes. Their identify mechanisms with low distortion bounds not only for the social welfare but also for several other objectives which are appropriate for the distributed setting, and explore the limitations in the design of distributed mechanisms via (almost) matching lower bounds; see also [39].

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Importantly, the work of [8] only considers a *discrete* social choice setting, in which there is a finite set of alternatives over which the agents are required to choose. Many real-world problems are better modeled as settings where there is a *continuum* of alternatives (e.g., captured by the line of real numbers). Traditionally, this setting has become known as *facility location* [37] and its centralized variant is one of the most well-studied topics in social choice theory; see the recent survey of Chan et al. [15] for a detailed overview. The distributed variant of the continuous setting was first studied by Filos-Ratsikas and Voudouris [25], who provided upper and lower bounds on the distortion of mechanisms for the social cost objective, the sum of costs of all the agents. Filos-Ratsikas and Voudouris [25] considered two types of mechanisms: (a) mechanisms that are only constrained by the fact that they operate in a distributed environment, and (b) mechanisms that are also constrained to be *strategyproof*, i.e., they do not provide incentives to the agents to lie about their preferences. While for the latter case the authors identified the mechanisms with the best possible distortion, for the former case they only managed to show that the distortion lies in the interval $[2, 3]$ leaving open the question of whether a mechanism with distortion 2 is actually possible. Our first contribution is to settle this open question in the affirmative: we design a novel mechanism for distributed facility location with a distortion of 2 for the social cost objective.

Besides the social cost objective, [8] identified three more objectives in the discrete setting which are particularly meaningful for the distributed setting, namely the maximum cost of any agent within any district, the maximum of the sum of costs of the agents in each district, and the sum of the maximum costs of the agents in each district. Following a similar approach, we study these four objectives in the continuous setting and provide upper and lower bounds on the distortion of distributed mechanisms, both with and without the strategyproofness requirement. All of our bounds are *tight*, meaning that we completely settle the distortion of distributed facility location on the real line.

1.1 Setting and Results

We consider a facility location setting with a set of *agents* that are positioned in the line of real numbers and are partitioned into disjoint *districts*. A *distributed mechanism* takes as input the positions of the agents and outputs a single point of the line where a public facility is to be located. This decision is made as follows: For each district, the mechanism chooses a location that is representative of the positions of the agents therein. Afterwards, it chooses the output to be one of the locations that represent the districts. The mechanism is distributed in the sense that the choice of the representative location of each district depends only on the positions reported by the agents that belong to the district.

We design deterministic distributed mechanisms that satisfy various criteria of interest and achieve the best possible distortion bounds. First, we aim to design distributed mechanisms to approximately optimize social objectives that are functions of the distances between the chosen locations and the positions of the agents. Following the work of [8], we focus on the following objectives:

- The total distance of the agents (*social cost*).
- The maximum distance among all agents (*max cost*).

Table 1: Overview of our tight distortion bounds for deterministic distributed mechanisms. The bound of 3 for the social cost and the class of strategyproof mechanisms marked with a \star is due to Filos-Ratsikas and Voudouris [25].

	Unrestricted	Strategyproof
Social cost	2 (Section 3)	3^\star
Max cost	2 (Section 4)	2 (Section 4)
Sum-of-Max	1 (Section 5.1)	$1 + \sqrt{2}$ (Section 5.2)
Max-of-Sum	2 (Section 6.1)	$1 + \sqrt{2}$ (Section 6.2)

- The total maximum agent distance in each district (*Sum-of-Max cost*)
- The maximum total agent distance in each district (*Max-of-Sum cost*).

To account for the possibly different sizes of the districts, Anshelevich et al. [8] considered the average total distance of the agents in their objectives (in and over the districts) instead of the sum of distances (i.e., they studied the Average-of-Average, Average-of-Max, and Max-of-Average objectives). To clearly demonstrate the arguments in the proofs without overcomplicating the notation, in this paper we focus exclusively on symmetric districts that contain the same number of agents, in which case the average is equivalent to the sum. Similarly to the case of Anshelevich et al. [8], our proofs can be easily extended to the case of asymmetric districts by considering the average total distance in the objectives; for the Max cost no change is required at all.

Our first contribution is the design of a novel mechanism that achieves a distortion of 2 for the social cost; as mentioned above, this settles a question left open in the work of Filos-Ratsikas and Voudouris [25] in the affirmative, matching their lower bound of 2. For the remaining objectives we provide mechanisms as well as lower bounds establishing that these mechanisms achieve the best possible distortion. The precise bounds are shown in the first column of Table 1. Quite interestingly, and perhaps unexpectedly, our mechanism for the Sum-of-Max objective is *optimal*, that is, it achieves a distortion of 1. This demonstrates that for this particular objective, the distributed nature of the decision making does not influence the quality of the decision at all, and stands in contrast to the results of Anshelevich et al. [8] for the same objective in the discrete setting.

Next, we consider strategyproof mechanisms, i.e., mechanisms that do not incentivize the agents to misreport their locations. This type of mechanisms were considered by Filos-Ratsikas and Voudouris [25] who settled their distortion for the social cost. For the remaining three objectives, strategyproof mechanisms have not been previously studied, not even in the discrete setting of Anshelevich et al. [8]. We show tight bounds by carefully composing centralized statistics mechanisms for choosing the district representatives and the final location; in particular, depending on the objective at hand, we appropriately choose the values of two parameters p and q to define mechanisms that work by choosing the position of the q -th agent in a district as its representative, and then select the p -th representative as the output location. Our results for strategyproof mechanisms are shown in the second column of Table 1.

1.2 Related Work

The distortion was originally defined by Procaccia and Rosenschein [36] to quantify the loss in social welfare due to social choice mechanisms having access only to preference rankings over the possible outcomes, rather than to the complete cardinal structure of the preferences. The distortion of mechanisms has been studied for several social choice problems, including single-winner voting, multi-winner voting, participatory budgeting, and matching in both the normalized utilitarian setting [10, 11, 13, 18, 22], as well as the metric setting [6, 9, 14, 16, 27, 29, 30]. Recently, the notion of the distortion has been more broadly interpreted as capturing the deterioration of an aggregate objective due to limited information, giving rise to works on communication complexity [34, 35], query complexity [2–4, 33], and other tradeoffs between information and distortion [1], as well as the distortion of distributed mechanisms that we study in the present paper [8, 24, 25]. We refer the reader to the survey of [7] for a detailed exposition.

The literature on strategyproof facility location is also rather extensive. Procaccia and Tennenholtz [37] were the first to study strategyproof facility location problems as part of their approximate mechanism design agenda. Since then, several variants of the problem have been proposed and studied, including settings in which there are several facilities to locate [26, 31, 32], the space of possible locations is restricted [21, 28, 38], the agents have heterogeneous preferences over the facilities [5, 20, 40], only some of the available facilities can be located [17, 19], or the aim is to optimize different objectives [12, 23, 41]. We refer the reader to the survey of Chan et al. [15] for more details.

2 PRELIMINARIES

An instance of our problem is a tuple $I = (N, \mathbf{x}, D)$, where

- N is a set of n agents.
- $\mathbf{x} = (x_i)_{i \in N}$ is a vector containing the *position* $x_i \in \mathbb{R}$ of agent i on the line of real numbers.
- $D = \{d_1, \dots, d_k\}$ is a set of $k \geq 1$ *districts*. Each district $d \in D$ contains a set $N_d \subseteq N$ of agents such that $N_d \cap N_{d'} = \emptyset$ and $\bigcup_{d \in D} N_d = N$. We mainly focus on the case of *symmetric* districts so that $|N_d| := \lambda := n/k$ for every $d \in D$.

A *distributed mechanism* M is used to decide the location of a facility based on the positions reported by the agents and the composition of the districts. In particular, given an instance I , a distributed mechanism works by implementing the following two steps:

- Step 1: For each district $d \in D$, using only the positions of the agents in N_d , the mechanism chooses a *representative location* $y_d \in \mathbb{R}$ for the district.
- Step 2: Given the representative locations of the districts, the mechanism outputs a single location $M(I) \in \{y_d\}_{d \in D}$ as the *winner*.

If a location z is chosen, then the distance $\delta(x_i, z) = |x_i - z|$ between the position x_i of agent i and z is the *individual cost* of agent i for z .

2.1 Social Objectives and Strategyproofness

We want to design mechanisms that output locations which are efficient according to a social objective. Let $z \in \mathbb{R}$ be any location. We consider the following four social minimization objectives:

- The *Sum cost* (or *social cost*) of location z is the total individual cost of all agents for z :

$$\sum_{i \in N} \delta(x_i, z) = \sum_{d \in D} \sum_{i \in N_d} \delta(x_i, z).$$

- The *Max cost* of location z is the maximum individual cost over all agents for z :

$$\max_{i \in N} \delta(x_i, z) = \max_{d \in D} \max_{i \in N_d} \delta(x_i, z).$$

- The *Sum-of-Max cost* of location z is the sum over each district of the maximum individual cost therein:

$$\sum_{d \in D} \max_{i \in N_d} \delta(x_i, z).$$

- The *Max-of-Sum cost* of location z is the maximum over each district of the total individual cost therein:

$$\max_{d \in D} \sum_{i \in N_d} \delta(x_i, z).$$

To simplify our notation, whenever the social objective is clear from context, we will use $\text{cost}(z|I)$ to denote the cost of $z \in \mathbb{R}$ according to the objective function at hand in instance I .

Another goal is to design mechanisms that are resilient to strategic manipulation, that is, they do not allow the agents to unilaterally affect the outcome in their favor (i.e., lead to a location with smaller individual cost) by reporting false positions. Formally, a mechanism is *strategyproof* if for any pair of instances $I = (N, (\mathbf{x}_{-i}, x_i), D)$ and $J = (N, (\mathbf{x}_{-i}, x'_i), D)$ that differ in the position of a single agent i , it holds that $\delta(x_i, M(I)) \leq \delta(x_i, M(J))$.

2.2 Distortion of Mechanisms

The *distortion* of a distributed mechanism M with respect to some social objective (which defines the cost of each possible location) is the worst case (over all instances) of the ratio between the cost of the location chosen by the mechanism and the minimum cost of any location:

$$\sup_{I=(N,\mathbf{x},D)} \frac{\text{cost}(M(I)|I)}{\min_{z \in \mathbb{R}} \text{cost}(z|I)}$$

By definition, the distortion is always at least 1. When the numerator is positive and the denominator is (extremely close to) 0, we will say that the distortion is unbounded. Our goal is to design distributed mechanisms that have an as low distortion as possible with respect to the social objectives defined above. We will consider both unrestricted mechanisms which assume that the agents act truthfully, as well as strategyproof mechanisms which aim to avoid strategic manipulations.

2.3 Useful Observations

Before we proceed with the presentation of our main technical results in the upcoming sections, we first state some useful properties. The bounds on the distortion of some of our mechanisms will follow by characterizing worst-case instances, and for that we will need the inequality

$$\frac{\alpha + \gamma}{\beta + \gamma} < \frac{\alpha}{\beta}, \quad (1)$$

which holds for any $\alpha > \beta \geq 0$ and $\gamma > 0$.

Filos-Ratsikas and Voudouris [25] observed that any distributed mechanism with finite distortion with respect to the social cost (sum objective) must be cardinally unanimous. We extend this result by showing this is true for any of the social objectives we consider in this paper. Formally, a mechanism is *cardinally-unanimous* if it chooses the representative location of a district to be z whenever all agents in the district are positioned at z . Due to space constraints, the proof of the following lemma and of a few other statements can be found in the supplementary material.

LEMMA 2.1. *Any distributed mechanism with finite distortion in terms of a social objective $F \in \{\text{Sum}, \text{Max}, \text{Sum-of-Max}, \text{Max-of-Sum}\}$ must be cardinally-unanimous.*

We next show that each member of a class of intuitive distributed mechanisms is strategyproof. Let $p \in [k]$ and $q \in [\lambda]$. The p -Statistic-of- q -Statistic mechanism first chooses the representative location of each district to be the position of the q -th ordered agent therein, and then outputs the p -th ordered representative location as the winner. For example, if $p = \lfloor (k+1)/2 \rfloor$ and $q = \lfloor (\lambda+1)/2 \rfloor$, the mechanism selects the position of the (leftmost) median agent in each district to be its representative location and then selects the (leftmost) median representative location as the winner. All strategyproof mechanisms that achieve the best possible distortion for the various social objectives we consider are members of this class, so in the next sections we will only focus on bounding their distortion.

LEMMA 2.2. *For any $p \in [k]$ and $q \in [\lambda]$, the p -Statistic-of- q -Statistic mechanism is strategyproof.*

3 SOCIAL COST

We begin with the social cost (Sum) objective. In previous work, Filos-Ratsikas and Voudouris [25] showed that the MEDIAN-OF-MEDIANS mechanism ($\lfloor \lambda/2 \rfloor$ -STATISTIC-OF- $\lfloor k/2 \rfloor$ -STATISTIC) has distortion at most 3, and this is best possible strategyproof mechanism. For the class of unrestricted mechanisms, they showed a lower bound of 2, thus leaving a gap between 2 and 3. Here, we complete the picture by showing a tight bound of 2 for unrestricted mechanisms. We do this by considering the MEDIAN-OF-TRUNCATEDAVG mechanism which works as follows: For each district, the mechanism considers a set of $\lambda/2$ agents ranging from the $(\lambda/4 + 1)$ -th leftmost to the $(3\lambda/4)$ -th leftmost¹, and chooses their average as the representative location of the district. Then, it chooses the median representative location as the final location. See Mechanism 1 for a detailed description.

Mechanism 1: MEDIAN-OF-TRUNCATEDAVG

```

for each district  $d \in D$  do
   $S_d := \{i \in N_d : i \text{ is at least the } (\lambda/4 + 1)\text{-th and at most the } (3\lambda/4)\text{-th leftmost agent}\};$ 
   $y_d := \frac{\sum_{i \in S_d} x_i}{|S_d|};$ 
return  $w := \text{Median}_{d \in D}\{y_d\};$ 

```

¹For simplicity, we present the mechanism assuming that the number of agents in each district is a multiple of 4; extending the description of the mechanism and the proof is straightforward.

To bound the distortion of MEDIAN-OF-TRUNCATEDAVG, we will characterize the structure of worst-case instances, where the distortion of the mechanism is maximized and is strictly larger than 1. Let w be the location chosen by the mechanism when given as input a worst-case instance, and denote by o the optimal location; since the objective is the social cost, o is the position of the median agent (or any point between the positions of the median agents in case of an even total number of agents). Without loss of generality, we assume that $w < o$; the case $w > o$ is symmetric.

We first show that there are cases where, starting from an instance with distortion strictly larger than 1, moving particular agents to appropriate intervals, leads to new instances that have strictly worse distortion. This transformation will be useful when characterizing the worst-case instances for the mechanism. Due to space constraints, the proof of the next lemma and of other statements are omitted.

LEMMA 3.1. *Let I and J be two instances that differ on the position of a single agent i , such that w is the location chosen by the mechanism and o is the optimal location for both instances. The distortion of the mechanism when given J as input is strictly larger than its distortion when given I as input in the following cases:*

- (a) i is positioned at $x_i < o$ in I , and at $x'_i \in (x_i, o]$ in J ;
- (b) i is positioned at $x_i > o$ in I , and at $x'_i \in [o, x_i)$ in J .

Using Lemma 3.1, we can show the following structural properties of worst-case instances. The main proof idea is that we start with a worst-case instance not satisfying the desired property, and then reach a contradiction by appropriately moving agents to define instances that end up with larger distortion (which should not be possible).

LEMMA 3.2. *There are no district representatives to the left of w .*

LEMMA 3.3. *Besides w , the only other district representative can be o , and all agents in such districts are positioned on o .*

LEMMA 3.4. *There exists a worst-case instance with one district d_w represented by w and one district d_o represented by o .*

Let ℓ and r be the locations of the leftmost and rightmost location among agents in S_{d_w} .

LEMMA 3.5. *In district d_w , $\ell \leq w \leq r$ and $r \leq o$.*

LEMMA 3.6. *In district d_w , the $\lambda/4$ leftmost agents are positioned at ℓ and the $\lambda/4$ rightmost agents are positioned at o .*

Using the above structural lemmas, we are finally ready to prove the main result of this section.

THEOREM 3.7. *The distortion of MEDIAN-OF-TRUNCATEDAVG is at most 2.*

PROOF. By Lemmas 3.3, 3.4 and 3.6, we have that the 2λ agents in the worst-case instance I are distributed on the line as follows: $\lambda/4$ agents are positioned at ℓ , $5\lambda/4$ agents are positioned at o (λ agents from d_o and $\lambda/4$ agents from d_w), and $\lambda/2$ agents are positioned in $[\ell, r]$. We partition the $\lambda/2$ agents in S_{d_w} into two sets: $L = \{i \in S_{d_w} : x_i \leq w\}$ and $R = \{i \in S_{d_w} : x_i > w\}$. Since

$r \leq o$ (due to Lemma 3.5) and $w = \sum_{i \in L \cup R} x_i$ (by definition), the optimal cost is

$$\begin{aligned} \text{cost}(o|I) &= \frac{\lambda}{4}(o - \ell) + \sum_{i \in L} (o - x_i) + \sum_{i \in R} (o - x_i) \\ &= \frac{\lambda}{4}(o - \ell) + \frac{\lambda}{2}(o - w) = \frac{\lambda}{4}(w - \ell) + \frac{3\lambda}{4}(o - w). \end{aligned} \quad (2)$$

Similarly, the cost of the mechanism is

$$\begin{aligned} \text{cost}(w|I) &= \frac{\lambda}{4}(w - \ell) + \sum_{i \in L} (w - x_i) \\ &\quad + \sum_{i \in R} (x_i - w) + \frac{5\lambda}{4}(o - w). \end{aligned} \quad (3)$$

By the definition of w , $\sum_{i \in L} (w - x_i) = \sum_{i \in R} (x_i - w)$. Also, again by definition, $|L| \geq 1$. If $R = \emptyset$, it must be the case that $\ell = w = r$, and the distortion is at most $5/3$ as Equations (2) and (3) are simplified to $\text{cost}(o|I) = 3(o - w)/4$ and $\text{cost}(w|I) = 5(o - w)/4$, respectively. Hence, in the rest of the proof we will assume that $|R| \geq 1$.

Since $x_i \leq o$ for each agent $i \in R$ and $|L| + |R| = \lambda/2$, we have

$$\sum_{i \in L} (w - x_i) = \sum_{i \in R} (x_i - w) \leq |R|(o - w) \Leftrightarrow o - w \geq \frac{\sum_{i \in L} (w - x_i)}{\lambda/2 - |L|}.$$

Similarly, as $x_i \geq \ell$ for each agent $i \in L$, we obtain

$$\sum_{i \in L} (w - x_i) \leq |L|(w - \ell) \Leftrightarrow w - \ell \geq \frac{\sum_{i \in L} (w - x_i)}{|L|}.$$

Let $o - w = \frac{\sum_{i \in L} (w - x_i)}{\frac{\lambda}{2} - |L|} + \xi_1$ and $w - \ell = \frac{\sum_{i \in L} (w - x_i)}{|L|} + \xi_2$, where $\xi_1, \xi_2 \geq 0$. Therefore, Equations (2) and (3) can be rewritten as

$$\begin{aligned} \text{cost}(o|I) &= \frac{\lambda}{4} \left(\frac{\sum_{i \in L} (w - x_i)}{|L|} + \xi_2 \right) + \frac{3\lambda}{4} \left(\frac{\sum_{i \in L} (w - x_i)}{\frac{\lambda}{2} - |L|} + \xi_1 \right) \\ \text{cost}(w|I) &= \frac{\lambda}{4} \left(\frac{\sum_{i \in L} (w - x_i)}{|L|} + \xi_2 \right) + \frac{5\lambda}{4} \left(\frac{\sum_{i \in L} (w - x_i)}{\frac{\lambda}{2} - |L|} + \xi_1 \right) \\ &\quad + 2 \sum_{i \in L} (w - x_i). \end{aligned}$$

It is not hard to see that, unless the distortion is at most $5/3$ and the claim holds trivially, the ratio is maximized when $\xi_1 = \xi_2 = 0$. We can then obtain the following upper bound on the distortion.

$$\begin{aligned} \frac{\text{cost}(w|I)}{\text{cost}(o|I)} &\leq \frac{\frac{\lambda}{4|L|} + \frac{5\lambda}{2\lambda - 4|L|} + 2}{\frac{\lambda}{4|L|} + \frac{3\lambda}{2\lambda - 4|L|}} \\ &\leq 2, \end{aligned}$$

where the last inequality follows since

$$\frac{\lambda}{4|L|} + \frac{5\lambda}{2\lambda - 4|L|} + 2 \leq 2 \left(\frac{\lambda}{4|L|} + \frac{3\lambda}{2\lambda - 4|L|} \right) \Leftrightarrow (\lambda - 4|L|)^2 \geq 0.$$

This concludes the proof. \square

4 MAX COST

We now consider the Max cost objective, for which we show a tight bound of 2 for both unrestricted and strategyproof mechanisms. We begin with the lower bound.

THEOREM 4.1. *For Max cost, the distortion of any mechanism (unrestricted or strategyproof) is at least 2.*

PROOF. Consider any mechanism and the following instance I with two districts. The agents in the first district are all positioned at -1 , while the agents in the second district are all positioned at 1 . Due to unanimity (Lemma 2.1), the representatives of the two districts must be -1 and 1 , respectively. Hence, the winner is either -1 or 1 . However, $\text{cost}(-1|I) = \text{cost}(1|I) = 2$, whereas $\text{cost}(0|I) = 1$, leading to a distortion of 2. \square

For the upper bound, we consider the ARBITRARY mechanism, which chooses the representative of each district to be the position of any agent therein, and then chooses any representative as the final winner. See Mechanism 2 for a specific implementation of this mechanism using the position of the leftmost agent from each district as the district representative, and then the leftmost representative as the final winner. Clearly, ARBITRARY is equivalent to some p -Statistic-of- q -Statistic mechanism depending on the choices within and over districts; the particular implementation of ARBITRARY as Mechanism 2 is equivalent to 1-Statistic-of-1-Statistic.

Mechanism 2: ARBITRARY (LEFTMOST-OF-LEFTMOST)

for each district d do

$y_d := \min_{i \in N_d} \{x_i\};$

return $w := \min_d \{y_d\};$

THEOREM 4.2. *For Max cost, the distortion of ARBITRARY is at most 2.*

PROOF. Given any instance I , let ℓ and r denote the positions of the leftmost and the rightmost agent, respectively. Clearly, the optimal location is $o = \frac{r - \ell}{2}$, and thus $\text{cost}(o|I) = \frac{r - \ell}{2}$. On the other hand, the ARBITRARY mechanism will necessarily return the location of some agent as the winner w , and hence $\text{cost}(w|I) \leq r - \ell$; the claim follows. \square

5 SUM-OF-MAX

Here, we focus on the Sum-of-Max objective; recall that for this objective we sum over each district the maximum agent cost therein. For unrestricted mechanisms, we show that, surprisingly, it is possible to achieve a distortion of 1, whereas, for strategyproof mechanisms, we show a tight bound of $1 + \sqrt{2}$.

5.1 Unrestricted Mechanisms

We will show that the MEDIAN-OF-MIDPOINTS mechanism optimizes the Sum-of-Max objective. This mechanism chooses the representative of each district to be the midpoint of the interval defined by the positions of the agents therein, and then chooses the median representative (breaking ties in favor of the leftmost median in case there are two) as the final winner. See Mechanism 3 for a detailed description.

To show the desired bound on the distortion of MEDIAN-OF-MIDPOINTS for Sum-of-Max, we again show some useful properties of worst-case instances. Without loss of generality, we will assume that the mechanism is applied on input a worst-case instance I

Mechanism 3: MEDIAN-OF-MIDPOINTS

for each district d do

$$\left[y_d := \frac{1}{2} \cdot \left(\max_{i \in N_d} x_i + \min_{i \in N_d} x_i \right); \right.$$

return $w := \text{Median}_{d \in D} \{y_d\}$;

where the chosen winner w is to the left of the optimal location o , that is, $w < o$.

LEMMA 5.1. *No midpoints are at the left of w or the right of o .*

LEMMA 5.2. *In each district represented by w , all agents are positioned at w .*

We are now ready to show that MEDIAN-OF-MIDPOINTS optimizes the Sum-of-Max objective.

THEOREM 5.3. *For Sum-of-Max, the distortion of MEDIAN-OF-MIDPOINTS is 1.*

PROOF. We will show that $\text{cost}(w|I) \leq \text{cost}(o|I)$; to simplify our notation, we drop I from cost for the rest of this proof. Let D_w be the set of districts represented by w , and denote by $\overline{D_w}$ the remaining districts. By Lemma 5.1 and the fact that w is the median midpoint, we have that $|D_w| \geq |\overline{D_w}|$. Also, by Lemma 5.2, we know that all agents in the districts of D_w are positioned at w . Hence,

$$\begin{aligned} \text{cost}(w) &= \sum_{d \in D_w} \max_{i \in N_d} \delta(x_i, w) + \sum_{d \notin D_w} \max_{i \in N_d} \delta(x_i, w) \\ &= \sum_{d \notin D_w} \max_{i \in N_d} \delta(x_i, w). \end{aligned}$$

We can write the cost of o as follows:

$$\begin{aligned} \text{cost}(o) &= \sum_{d \in D_w} \max_{i \in N_d} \delta(x_i, o) + \sum_{d \notin D_w} \max_{i \in N_d} \delta(x_i, o) \\ &= |D_w| \delta(w, o) + \sum_{d \notin D_w} \max_{i \in N_d} \delta(x_i, o). \end{aligned}$$

Now, let $d \notin D_w$ be a district with midpoint y , which is such that $y \in (w, o]$ due to Lemma 5.1. Let ℓ_d and r_d be the leftmost and rightmost agents of district d , respectively. We claim that the contribution of d to $\text{cost}(o)$ is the distance $\delta(\ell_d, o)$, and the contribution of d to $\text{cost}(w)$ is the distance $\delta(r_d, w)$. First observe that since $y \in (w, o]$, it cannot be the case that $\ell_d, r_d < w$ or $\ell_d, r_d > o$; in other words, we necessarily have that $w \leq r_d$ and $\ell_d \leq o$. If $\ell_d \leq r_d < o$, then our claim for $\text{cost}(o)$ follows immediately. So, suppose that $r_d > o$, and thus $\ell_d < o$. Then, the inequality $\frac{\ell_d + r_d}{2} \leq o$ due to the fact that $y \in (w, o]$ implies that $r_d - o \leq o - \ell_d$, and thus $\delta(\ell_d, o)$ is the contribution of d to $\text{cost}(o)$. Similarly, if $w \leq \ell_d \leq r_d$, our claim for $\text{cost}(w)$ follows immediately. Thus, we can suppose that $\ell_d < w$, and thus $r_d > w$. The inequality $\frac{\ell_d + r_d}{2} > w$ implies that $w - \ell_d < r_d - w$, and thus $\delta(r_d, w)$ is the contribution of d to $\text{cost}(w)$. By this, we have that

$$\text{cost}(w) = \sum_{d \notin D_w} \delta(r_d, w).$$

Furthermore, for each $d \notin D_w$,

$$\delta(\ell_d, o) = o - \ell_d$$

$$\begin{aligned} &= o - r_d + r_d - w + w - \ell_d \\ &= \delta(r_d, w) + (o + w) - (\ell_d + r_d) \\ &\geq \delta(r_d, w) + (o + w) - 2o \\ &= \delta(r_d, w) - (o - w) \\ &= \delta(r_d, w) - \delta(w, o), \end{aligned}$$

where the inequality follows since $\frac{\ell_d + r_d}{2} \leq o$. Consequently,

$$\begin{aligned} \text{cost}(o) &= |D_w| \delta(w, o) + \sum_{d \notin D_w} \delta(\ell_d, o) \\ &\geq |D_w| \delta(w, o) + \sum_{d \notin D_w} \left(\delta(r_d, w) - \delta(w, o) \right) \\ &= (|D_w| - |\overline{D_w}|) \delta(w, o) + \sum_{d \notin D_w} \delta(r_d, w) \\ &\geq \sum_{d \notin D_w} \delta(r_d, w) \\ &= \text{cost}(w), \end{aligned}$$

where the last inequality follows since $|D_w| \geq |\overline{D_w}|$. \square

5.2 Strategyproof Mechanisms

For strategyproof mechanisms, we will show a tight bound of $1 + \sqrt{2}$. We start by showing the lower bound on the distortion of all strategyproof mechanisms. The proof of the next theorem relies on arguing about how a strategyproof mechanism would need to make decisions on specific instances to achieve distortion less than $1 + \sqrt{2}$, and that such decisions lead to contradiction. We sketch the proof here.

THEOREM 5.4. *For Sum-of-Max, the distortion of any strategyproof mechanism is at least $1 + \sqrt{2} - \varepsilon$, for any $\varepsilon > 0$.*

PROOF SKETCH. We assume towards a contradiction that there is a strategyproof mechanism with distortion strictly smaller than $1 + \sqrt{2} - \varepsilon$, for any $\varepsilon > 0$. Without loss of generality, when there are two districts with different representatives we choose the leftmost as the final winner. The proof goes by first showing that there is a district with two agents positioned at 0 and 1 such that the mechanism must choose 1 as the representative of the district as otherwise the distortion would be at least 3. Next, we show that in instances with $\alpha + \beta$ districts, with $\beta/\alpha \approx 1 + \sqrt{2}$, such that $1/2$ is the representative of α districts and 1 is the representative of β districts, the overall winner must be 1 as otherwise the distortion would be $\beta/\alpha \approx 1 + \sqrt{2}$. The proof concludes by considering an instance containing $\alpha + \beta$ districts with $1/2$ as the representative of α districts and 1 as the representative of β districts, such that choosing 1 as the overall winner leads to a distortion of approximately $1 + \sqrt{2}$. \square

For the tight upper bound, we consider the $\left(1 - 1/\sqrt{2}\right)k$ -LEFTMOST-OF-RIGHTMOST mechanism, which chooses the representative of each district to be the position of the rightmost agent therein, and then chooses the $\left(1 - 1/\sqrt{2}\right)k$ -th leftmost representative as the final winner.² See Mechanism 4 for a detailed description. Clearly, the

²To be precise, the mechanism chooses the $\left[\left(1 - 1/\sqrt{2}\right)k\right]$ -leftmost representative as the winner. To simplify our notation and discussion, we drop the ceiling.

Mechanism 4: $\left(1 - 1/\sqrt{2}\right)k$ -LEFTMOST-OF-RIGHTMOST

for each district $d \in D$ **do**

$\lfloor y_d :=$ rightmost agent;

return $w := \left(1 - 1/\sqrt{2}\right)k$ -th leftmost representative;

mechanism is strategyproof as it is an implementation of p -Statistic-of- q -Statistic with $p = \left(1 - 1/\sqrt{2}\right)k$ and $q = \lambda$. So, it suffices to show that it achieves a distortion of at most $1 + \sqrt{2}$.

THEOREM 5.5. *For Sum-of-Max, the distortion of $\left(1 - 1/\sqrt{2}\right)k$ -LEFTMOST-OF-RIGHTMOST is at most $1 + \sqrt{2}$.*

PROOF. Consider any instance I . Let w be the location chosen by the mechanism, and o the optimal location. For each district d , let i_d be the most distant agent from w , and i_d^* the most distant agent from o . So, $\text{cost}(w|I) = \sum_{d \in D} \delta(i_d, w)$, and $\text{cost}(o|I) = \sum_{d \in D} \delta(i_d^*, o) \geq \sum_{d \in D} \delta(j, o)$ for any agent $j \in N_d$. We consider the following two cases depending on the relative positions of w and o .

Case 1: $o < w$. Let $S = \{d \in D : y_d \geq w\}$ be the set of district representatives to the right of w . By the definition of w , we have that $|S| \geq \frac{k}{\sqrt{2}}$. Since $o < w \leq y_d$ for every $d \in S$ and $y_d \in N_d$, we have that

$$\begin{aligned} \text{cost}(o|I) &\geq \sum_{d \in S} \delta(y_d, o) \geq |S| \cdot \delta(w, o) \geq \frac{k}{\sqrt{2}} \cdot \delta(w, o) \\ &\Leftrightarrow k \cdot \delta(w, o) \leq \sqrt{2} \cdot \text{cost}(o|I). \end{aligned}$$

By the triangle inequality and since $i_d \in N_d$, we have

$$\begin{aligned} \text{cost}(w|I) &= \sum_{d \in D} \delta(i_d, w) \leq \sum_{d \in D} \delta(i_d, o) + \sum_{d \in D} \delta(w, o) \\ &\leq \text{cost}(o|I) + k \cdot \delta(w, o) \\ &\leq (1 + \sqrt{2}) \cdot \text{cost}(o|I). \end{aligned}$$

Case 2: $w < o$. We partition the districts into a set L that includes $\left(1 - \frac{1}{\sqrt{2}}\right)k$ districts from the one with the leftmost representative until the one with the $\left(1 - \frac{1}{\sqrt{2}}\right)k$ -th leftmost representative (that is, w), and a set R that includes the remaining districts. By definition, we have that $|R|/|L| = 1 + \sqrt{2}$. For every district d , let ℓ_d and r_d be the leftmost and rightmost agents in d , respectively. We make the following observations:

- For every $d \in L$, since y_d is the rightmost agent of d and $y_d \leq w < o$, it must be the case that $i_d = i_d^* = \ell_d$. Due to the positions of ℓ_d , w and o , we have that $\delta(\ell_d, o) = \delta(\ell_d, w) + \delta(w, o)$.
- For every $d \in R$, by the triangle inequality, we have that $\delta(i_d, w) \leq \delta(i_d, o) + \delta(w, o)$. Since $\delta(i_d, o) \leq \delta(i_d^*, o)$ by the definition of i_d^* , we further have that $\delta(i_d, w) \leq \delta(i_d^*, o) + \delta(w, o)$.

Hence,

$$\text{cost}(w|I) = \sum_{d \in D} \delta(i_d, w) = \sum_{d \in L} \delta(\ell_d, w) + \sum_{d \in R} \delta(i_d, w)$$

$$\begin{aligned} &\leq \sum_{d \in L} \left(\delta(\ell_d, w) + \delta(w, o) \right) - |L| \delta(w, o) \\ &\quad + \sum_{d \in R} \left(\delta(i_d^*, o) + \delta(w, o) \right) \\ &= \text{cost}(o|I) + (|R| - |L|) \delta(w, o). \end{aligned}$$

Since $y_d \leq w < o$ for every $d \in L$ and $y_d \in N_d$, we have that

$$\text{cost}(o|I) \geq \sum_{d \in L} \delta(y_d, o) \geq |L| \cdot \delta(w, o) \Leftrightarrow \delta(w, o) \leq \frac{1}{|L|} \cdot \text{cost}(o|I).$$

Therefore, we obtain

$$\begin{aligned} \text{cost}(w|I) &\leq \text{cost}(o|I) + \frac{|R| - |L|}{|L|} \cdot \text{cost}(o|I) \\ &= \frac{|R|}{|L|} \cdot \text{cost}(o|I) = (1 + \sqrt{2}) \cdot \text{cost}(o|I), \end{aligned}$$

as desired. \square

6 MAX-OF-SUM

We now turn our attention to the last objective, Max-of-Sum, which is the maximum over each district of the total individual cost therein. We show a tight bound of 2 for unrestricted mechanisms and a tight bound of $1 + \sqrt{2}$ for strategyproof mechanisms.

6.1 Unrestricted Mechanisms

Since the lower bound of 2 for the Max cost objective holds even when there is a single agent in each district, it extends to the case of Max-of-Sum as well. For the upper bound, we consider the ARBITRARY-OF-AVG mechanism, which chooses the representative of each district to be the average of the positions of the agents in the district, and then chooses an arbitrary representative (e.g., the leftmost) as the final winner. See Mechanism 5 for a detailed description.

Mechanism 5: ARBITRARY-OF-AVG

for each district $d \in D$ **do**

$\lfloor y_d := (\sum_{i \in N_d} x_i) / \lambda$;

return $w := \min_{d \in D} y_d$;

THEOREM 6.1. *For Max-of-Sum, the distortion of ARBITRARY-OF-AVG is at most 2.*

PROOF. Consider any instance I . Let w be the location chosen by the mechanism, and o the optimal location; without loss of generality, we assume that $w < o$. Denote by d^* a district that defines the cost of w , that is, $d^* \in \arg \max_{d \in D} \sum_{i \in N_d} \delta(x_i, w)$. Also, denote by d_w a district represented by w , that is, $w = \frac{1}{\lambda} \sum_{i \in N_{d_w}} x_i \Leftrightarrow \sum_{i \in N_{d_w}} (w - x_i) = 0$. By the triangle inequality, we have that

$$\begin{aligned} \text{cost}(w|I) &= \sum_{i \in N_{d^*}} \delta(x_i, w) \\ &\leq \sum_{i \in N_{d^*}} \delta(x_i, o) + \sum_{i \in N_{d^*}} \delta(w, o) \leq \text{cost}(o|I) + \lambda \delta(w, o). \end{aligned}$$

By the definition of d_w and since it consists of λ agents, we have that

$$\begin{aligned} \lambda\delta(w, o) &= \lambda(o - w) = \lambda(o - w) + \sum_{i \in N_{d_w}} (w - x_i) \\ &= \sum_{i \in N_{d_w}} (o - x_i) \leq \sum_{i \in N_{d_w}} \delta(x_i, o) \leq \text{cost}(o|I), \end{aligned}$$

where the inequality follows since $\delta(x_i, o) = o - x_i$ when $x_i \leq o$ and $\delta(x_i, o) = x_i - o \geq o - x_i$ when $x_i \geq o$. Therefore, we obtain that $\text{cost}(w|I) \leq 2 \cdot \text{cost}(o|I)$, as desired. \square

6.2 Strategyproof Mechanisms

We now turn our attention to strategyproof mechanisms and first show a lower bound of $1 + \sqrt{2}$.

THEOREM 6.2. *For Max-of-Sum, the distortion of any strategyproof mechanism is at least $1 + \sqrt{2} - \varepsilon$, for any $\varepsilon > 0$.*

PROOF. Suppose towards a contradiction that there is a strategyproof mechanism with distortion strictly smaller than $1 + \sqrt{2} - \varepsilon$, for any $\varepsilon > 0$.

Property (P1): Consider a district with $(1 + \sqrt{2})x$ agents at 0 and x agents at 1, where x is an arbitrarily large integer.³ We claim that the mechanism must choose 0 as the representative of this district as otherwise the distortion would be at least $1 + \sqrt{2}$. Indeed, suppose that the representative is some $y \neq 0$. By moving one of the agents at 1 to y , we obtain a new district whose representative must still be y ; otherwise, in the instance that consists only of this new district, the agent at y would have incentive to misreport her position as 1, thus leading to the representative (and the final winner) to change to y . By induction, we obtain that y must be the representative of the district with $(1 + \sqrt{2})x$ agents at 0 and x agents at y . In the instance I that consists of only the latter district, the winner is y with $\text{cost}(y|I) = (1 + \sqrt{2})x \cdot |y|$, whereas $\text{cost}(0|I) = x \cdot |y|$, leading to a distortion of at least $1 + \sqrt{2}$.

Property (P2): Consider a district with x agents at 1 and $(1 + \sqrt{2})x$ agents at 2. We claim that the mechanism must choose 2 as the representative of this district as otherwise the distortion would be at least $1 + \sqrt{2}$. This follows by arguments similar to those for property (P1).

Reaching a contradiction: Consider the following instance J with two districts:

- In the first district, there are $(1 + \sqrt{2})x$ agents at 0 and x agents at 1.
- In the second district, there are x agents at 1 and $(1 + \sqrt{2})x$ agents at 2.

By properties (P1) and (P2), the representatives of the two districts must be 0 and 2, respectively, and thus one of these two locations is chosen as the final winner. However, $\text{cost}(0|J) = \text{cost}(2|J) = 2(1 + \sqrt{2})x + x$, while $\text{cost}(1|J) = (1 + \sqrt{2})x$, leading to a distortion of $2 + \frac{1}{1 + \sqrt{2}} = 1 + \sqrt{2}$. \square

³To be precise, since the number of agents must be an integer, we would need to have $\lceil (1 + \sqrt{2})x \rceil$ agents at 0. We simplify our notation by dropping the ceilings, but it should be clear that this does not affect our arguments.

For the tight upper bound, we consider the RIGHTMOST-OF- $(1 - 1/\sqrt{2})\lambda$ -LEFTMOST mechanism, which chooses the representative of each district to be the position of the $(1 - 1/\sqrt{2})\lambda$ -th leftmost agent therein, and then chooses the rightmost representative as the final winner. See Mechanism 6 for a detailed description. This mechanism is an implementation of p -Statistic-of- q -Statistic with $p = k$ and $q = (1 - 1/\sqrt{2})\lambda$, and is thus strategyproof. So, it suffices to show that it achieves a distortion of at most $1 + \sqrt{2}$.

Mechanism 6: RIGHTMOST-OF- $(1 - 1/\sqrt{2})\lambda$ -LEFTMOST

for each district $d \in D$ do

$y_d := (1 - 1/\sqrt{2})\lambda$ -th leftmost agent;

return $w :=$ rightmost representative;

THEOREM 6.3. *For Max-of-Sum, the distortion of RIGHTMOST-OF- $(1 - 1/\sqrt{2})\lambda$ -LEFTMOST is at most $1 + \sqrt{2}$.*

PROOF. Denote by d^* a district that gives the max sum for w , and by d_w a district represented by w . Also, for any district d , let $\text{cost}_d(x) = \sum_{i \in N_d} \delta(i, x)$ for any location x , and denote by o_d the location that minimizes this sum (so, o_d is the median agent of d). By definition, $\text{cost}(w) = \text{cost}_{d^*}(w)$ and $\text{cost}_d(o) \leq \text{cost}(o)$ for every district d . We will show the claim for $o < w$.

By the definition of d^* and the triangle inequality, we have

$$\begin{aligned} \text{cost}(w) &= \sum_{i \in N_{d^*}} \delta(i, w) \\ &\leq \sum_{i \in N_{d^*}} \delta(i, o) + \sum_{i \in N_{d^*}} \delta(o, w) \leq \text{cost}(o) + \lambda\delta(o, w). \end{aligned}$$

Let $S = \{i \in N_{d_w} : x_i \geq w\}$ be the set of agents that are positioned at the right of (or exactly at) w in d_w . By the definition of w , $|S| \geq \frac{1}{\sqrt{2}}\lambda$. Since $o < w$, we have

$$\begin{aligned} \text{cost}_{d_w}(o) &\geq |S| \cdot \delta(w, o) \geq \frac{1}{\sqrt{2}}\lambda \cdot \delta(w, o) \\ &\Leftrightarrow \lambda\delta(w, o) \leq \sqrt{2} \cdot \text{cost}_{d_w}(o) \leq \sqrt{2} \cdot \text{cost}(o). \end{aligned}$$

By combining everything together, we obtain a bound of $1 + \sqrt{2}$. \square

7 OPEN PROBLEMS

In this paper we settled the distortion of unrestricted and strategyproof mechanisms for the distributed single-facility location problem in terms of social objectives that are combinations of sum and max. There are several interesting directions for future work, such as to extend our work to more general metric spaces or to define further meaningful objectives and study similar questions about efficiency and strategyproofness. Beyond the single-facility location problem that we studied here, one could consider settings with more facilities and agents that have heterogeneous preferences over the facilities.

REFERENCES

- [1] Ben Abramowitz, Elliot Anshelevich, and Wennan Zhu. 2019. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*. 3–16.
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2021. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence* 296 (2021), 103488.
- [3] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. Don't Roll the Dice, Ask Twice: The Two-Query Distortion of Matching Problems and Beyond. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- [4] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. A Few Queries Go a Long Way: Information-Distortion Tradeoffs in Matching. *Journal of Artificial Intelligence Research* 74 (2022).
- [5] Eleftherios Anastasiadis and Argyrios Deligkas. 2018. Heterogeneous facility location games. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 623–631.
- [6] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2015. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 777–783.
- [7] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 4294–4301.
- [8] Elliot Anshelevich, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. The distortion of distributed metric social choice. *Artificial Intelligence* 308 (2022), 103713.
- [9] Elliot Anshelevich and John Postl. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.
- [10] Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 376–382.
- [11] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227 (2015), 190–213.
- [12] Qingpeng Cai, Aris Filos-Ratsikas, and Pingzhong Tang. 2016. Facility Location with Minimax Envy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 137–143.
- [13] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.
- [14] Ioannis Caragiannis, Nisarg Shah, and Alexandros A. Voudouris. 2022. The Metric Distortion of Multiwinner Voting. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4900–4907.
- [15] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. 2021. Mechanism Design for Facility Location Problem: A Survey. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 1–17.
- [16] Moses Charikar and Prasanna Ramakrishnan. 2022. Metric Distortion Bounds for Randomized Social Choice. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2986–3004.
- [17] Argyrios Deligkas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. Heterogeneous Facility Location with Limited Resources. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4966–4974.
- [18] Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. 2022. Optimized Distortion and Proportional Fairness in Voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 563–600.
- [19] Edith Elkind, Minming Li, and Houyu Zhou. 2022. Facility Location With Approval Preferences: Strategyproofness and Fairness. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 391–399.
- [20] Itai Feigenbaum and Jay Sethuraman. 2015. Strategyproof Mechanisms for One-Dimensional Hybrid and Obnoxious Facility Location Models. In *AAAI Workshops*.
- [21] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*. 269–286.
- [22] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th Symposium of Algorithmic Game Theory (SAGT)*. 1–12.
- [23] Aris Filos-Ratsikas, Minming Li, Jie Zhang, and Qiang Zhang. 2017. Facility location with double-peaked preferences. *Autonomous Agents and Multi Agent Systems* 31, 6 (2017), 1209–1235.
- [24] Aris Filos-Ratsikas, Evi Micha, and Alexandros A. Voudouris. 2020. The distortion of distributed voting. *Artificial Intelligence* 286 (2020), 103343.
- [25] Aris Filos-Ratsikas and Alexandros A. Voudouris. 2021. Approximate Mechanism Design for Distributed Facility Location. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT)*. 49–63.
- [26] Dimitris Fotakis and Christos Tzamos. 2016. Strategyproof facility location for concave cost functions. *Algorithmica* 76, 1 (2016), 143–167.
- [27] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 1427–1438.
- [28] Panagiotis Kanellopoulos, Alexandros A. Voudouris, and Rongsen Zhang. 2022. On Discrete Truthful Heterogeneous Two-Facility Location. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 342–348.
- [29] David Kempe. 2020. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2087–2094.
- [30] Fatih Erdem Kizilkaya and David Kempe. 2022. Plurality Veto: A Simple Voting Rule Achieving Optimal Metric Distortion. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 349–355.
- [31] Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*. 315–324.
- [32] Pinyan Lu, Yajun Wang, and Yuan Zhou. 2009. Tighter bounds for facility games. In *Proceedings of the 5th International Workshop Internet and Network Economics (WINE)*. 137–148.
- [33] Thomas Ma, Vijay Menon, and Kate Larson. 2021. Improving Welfare in One-Sided Matchings using Simple Threshold Queries. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 321–327.
- [34] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. 2019. Efficient and thrifty voting by any means necessary. In *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*. 7178–7189.
- [35] Debmalya Mandal, Nisarg Shah, and David P. Woodruff. 2020. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 795–813.
- [36] Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents (CIA)*. 317–331.
- [37] Ariel D. Procaccia and Moshe Tennenholtz. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation* 1, 4 (2013), 18:1–18:26.
- [38] Paolo Serafino and Carmine Ventre. 2016. Heterogeneous facility location without money. *Theoretical Computer Science* 636 (2016), 27–46.
- [39] Alexandros A. Voudouris. 2023. Tight Distortion Bounds for Distributed Single-Winner Metric Voting on a Line. *CoRR* abs/2301.01742 (2023).
- [40] Xiping Xu, Bo Li, Minming Li, and Lingjie Duan. 2021. Two-facility Location Games with Minimum Distance Requirement. *Journal of Artificial Intelligence Research* 70 (2021), 719–756.
- [41] Houyu Zhou, Minming Li, and Hau Chan. 2022. Strategyproof Mechanisms for Group-Fair Facility Location Problems. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 613–619.