Truthful Two-Facility Location with Candidate Locations

Panagiotis Kanellopoulos, Alexandros A. Voudouris, and Rongsen Zhang

School of Computer Science and Electronic Engineering, University of Essex, UK
{panagiotis.kanellopoulos,alexandros.voudouris,rz19109}@essex.ac.uk

Abstract. We study a truthful two-facility location problem in which a set of agents have private positions on the line of real numbers and known approval preferences over two facilities. Given the locations of the two facilities, the cost of an agent is the total distance from the facilities she approves. The goal is to decide where to place the facilities from a given finite set of candidate locations so as to (a) approximately optimize desired social objectives, and (b) incentivize the agents to truthfully report their private positions. We focus on the class of deterministic strategyproof mechanisms and pinpoint the ones with the best possible approximation ratio in terms of the social cost (i.e., the total cost of the agents) and the max cost. In particular, for the social cost, we show a tight bound of $1 + \sqrt{2}$ when the preferences of the agents are homogeneous (i.e., all agents approve both facilities), and a tight bound of 3 when the preferences might be heterogeneous. For the max cost, we show tight bounds of 2 and 3 for homogeneous and heterogeneous preferences, respectively.

Keywords: Facility location · Mechanism design · Approximation.

1 Introduction

In the well-studied *truthful single-facility location problem*, there is a set of agents with private positions on the line of real numbers, and a facility (such as a park or a school) that is to be located somewhere on the line. Given such a location, its distance from the position of an agent is interpreted as the individual cost that the agent would suffer by having to travel to the facility in order to be serviced by it. The goal is to determine the facility location so that the agents are given the right incentives to *truthfully* report their positions (that is, not being able to affect the outcome to decrease their cost), and, at the same time, a social function of the individual agent costs (such as the total or the maximum cost) is (approximately) optimized. Since the work of Procaccia and Tennenholtz [19], who were the first to consider facility location problems through the prism of *approximate mechanism design without money*, research on this topic has flourished and a large number of more complex variants of the problem have been introduced and analyzed; see the survey of [4] for an overview.

The original work of Procaccia and Tennenholtz [19] focused on a *continuous* model, where the facility is allowed to be placed at any point of the line, and showed tight bounds on the approximation ratio of deterministic and randomized *strategyproof* mechanisms in terms of the *social cost* (total individual cost of the agents) and the *max cost* (maximum individual cost among all agents). The *discrete* model, where the facility can be placed only at a given set of candidate points of the line, has also been studied, most notably by Feldman et al. [9]. They observed that this setting is equivalent to voting on a line, and the strategyproofness constraint leads to deterministic mechanisms that make decisions using only the ordinal preferences of agents over the candidate points; an assumption typically made in the distortion literature [3]. Many other truthful single-facility location models have been studied under different assumptions, such as that the location space is more general than a line [1, 18, 14], that the facility location must be decided in a distributed way [11, 10], or that the facility is obnoxious and the agents aim to be as far from it as possible [5].

Aiming to locate multiple facilities is a natural generalization. Most of the work in this direction has focused on the fundamental case of two facilities, under several different assumptions about the types of the facilities, the preferences of the agents for them, the agent-related information that is public or private, and whether the setting is continuous or discrete, aiming to capture different applications. In the *homogeneous* case, where the facilities are assumed to be of the same type and serve the same purpose (for example, they might correspond to two parks), with only a few exceptions, the typical assumption is that each agent cares about one of them, such as the facility placed closest to her position (i.e., the agent would like to have a shortest path to a facility) [19, 17, 12], or the facility farthest from her position (i.e., the agent would like to be in a short radius from both facilities) [25]. More generally, the facilities might be of different types and serve a different purpose (for example, they might correspond to a park and a school), in which case the agents might have different, *heterogeneous* preferences over them. For example, some agents might be interested in both facilities, some agents might be interested in only one of them, or some agents might consider one facility to be useful and the other to be obnoxious [8, 26, 16, 2, 6, 20, 15, 13].

In this paper, we consider the case of two facilities that can be placed only at candidate locations. For each agent, we assume that her position on the line is private, she has approval preferences over the facilities which are publicly known, and her individual cost is given by the total distance from the facilities she approves (rather than the distance from the closest or farthest such facility). To the best of our knowledge, these assumptions have not been considered in combination before in the literature. We provide more details below.

1.1 Our Contribution

We study a truthful two-facility location problem, in which there is a set of agents with *known approval preferences* (0 or 1) over two facilities $\{F_1, F_2\}$, so that each agent approves at least one facility; we can safely ignore agents that approve neither facility. The agents have *private positions* on the line of real

Table 1: An overview of the tight bounds that we show in this paper on the approximation ratio of deterministic strategyproof mechanisms for the different combinations of social objectives functions (social cost and max cost) and agent preferences (homogeneous or heterogeneous).

	Social cost	Max cost
Homogeneous	$1+\sqrt{2}$	2
Heterogeneous	3	3

numbers, and the facilities can only be placed at *different* locations chosen from a given set of *candidate* locations. Once the facilities have both been placed, the individual cost of each agent is the *total* distance from the facilities she approves.

Our goal is to design mechanisms that take as input the positions reported by the agents, and, using also the available information about the preferences of the agents, decide where to place the two facilities, so that (a) a social objective function is (approximately) optimized, and (b) the agents are incentivized to truthfully report their positions. As in previous work, we consider the well-known *social cost* (the total individual cost of the agents) and the *max cost* (the maximum individual cost over all agents) as our social objective functions. We also treat separately the class of instances in which the preferences of the agents are *homogeneous* (that is, when all agents approve both facilities) and the general class of instances in which the preferences might be *heterogeneous*. For all possible combinations of objectives and types of preferences, we design *deterministic* strategyproof mechanisms that achieve the best possible approximation ratio. An overview of our results is given in Table 1. Due to space constraints, the proofs of some of our statements have been omitted.

We start by focusing on the social cost in Section 3. For homogeneous instances, we show a tight bound of $1+\sqrt{2}$ by considering a natural mechanism, which, for a parameter $\alpha \in (0, 1/2)$, places the first facility at the candidate location closest to the αn-leftmost agent, and then places the second facility at the available candidate location that is closest to the $(1-\alpha)n$-leftmost agent, where n is the number of agents. We show that the desired bound is achieved for $\alpha = \sqrt{2}-1$. For general instances, we first observe that no strategyproof mechanism can achieve an approximation ratio better than 3; this follows from the fact that instances in which all agents approve just one of the facilities lead to a version of the problem that is essentially equivalent to single-facility location (for which 3 is the best possible bound [9]). We then show a matching upper bound of 3 by considering a mechanism that places each facility at the available candidate location closest to the median agent among those that approve it. To decide the order in which the facilities are placed, we first perform a voting step that allows the agents that approve each facility to decide if they prefer the closest or second-closest candidate location to the respective median agent; this is necessary since just blindly choosing the order of placing the facilities leads to a mechanism with a rather large approximation ratio.

In Section 4, we turn our attention to the max cost objective. For homogeneous instances, we show a best possible bound of 2 by considering a simple mechanism that places the facilities at the available candidate locations closest to the leftmost and rightmost agents. For general instances, we show a best possible bound of 3 by considering two mechanisms that work under different approval profiles. The first mechanism is a direct extension of the mechanism used for homogeneous instances and works when there is at least one agent that approves both facilities. The second mechanism works when each agent approves just one facility (that is, some agents approve only the first facility, while the others approve only the second facility). This case is harder to tackle since the straightforward variant of the mechanism for homogeneous instances does not work well. The bottleneck lies in that, after placing the first facility at the candidate location closest to one of the agents that approve it (such as the leftmost), we then need to dynamically decide whether the second facility can be placed closer to the leftmost or rightmost among the agents that approve it, or neither of them. This again is done by a voting-like procedure that is used to decide the order of the agents that approve the second facility relative to the two candidate locations that are closest to where the first facility has been placed.

1.2 Related Work

For an overview of the many different truthful facility location problems that have been considered in the literature, we refer the reader to the survey of [4]. Here, we briefly discuss papers that are mostly related to our work.

The truthful two-facility location problem was first considered in the original work of Procaccia and Tennenholtz [19], where the agents were assumed to have homogeneous preferences over the two facilities, their individual costs were equal to the distance from the closest facility to their positions, and the facilities were allowed to be placed anywhere on the line, even at the same point. For deterministic mechanisms, Procaccia and Tennenholtz showed a constant lower bound and a linear upper bound on the approximation ratio in terms of the social cost, and a tight bound of 3 in terms of the maximum cost. They also showed how randomization can lead to further improvements. Lu et al. [17] improved the lower bound for the social cost and deterministic mechanisms to an asymptotically linear one, before Fotakis and Tzamos [12] finally showed that the exact bound for this case is $n - 2$.

Sui and Boutilier [21] were among the first to consider truthful facility location problems with candidate locations (referred to as constrained facility location), with a focus on achieving approximate strategyproofness by bounding the incentives of the agents to manipulate; for multiple facilities, they considered only the homogeneous case where each agent's individual cost is the distance to the closest facility. As already mentioned, Feldman et al. [9] considered a candidate selection problem with a fixed set of candidates, a model which translates into a single-facility location problem where the facility can only be placed at a location from a given set of discrete candidate locations. They focused on the

social cost objective and, among other results, proved that the MEDIAN mechanism that places the facility at the location closest to the position reported by the median agent, achieves an upper bound of 3; they also showed that this is the best possible bound among deterministic mechanisms.

Serafino and Ventre [20] considered a slightly different discrete facility location problem, where agents occupy nodes on a line graph and have approval preferences over two facilities that can only be placed at different nodes of the line. In contrast to our work here, where we assume that the agents have private positions and public preferences, Serafino and Ventre assumed that the positions are known and the preferences unknown. They showed several bounds on the approximation ratio for deterministic and randomized strategyproof mechanisms for the social cost and the max cost. Some of their results for deterministic mechanisms were improved by Kanellopoulos et al. [15].

Tang et al. [22] considered a setting with agents that have homogeneous preferences over two facilities, which can be placed at locations chosen from a set of candidate ones, allowing the facilities to be placed even at the same location. The positions of the agents are assumed to be private information and an agent's individual cost is defined as her distance from the closest facility. They proved an upper bound of $2n-3$ for the social cost objective and a tight bound of 3 for the maximum cost. They also considered the case of a single facility and the max cost objective, for which they showed a bound of 3 (extending the work of Feldman et al. [9] who only focused on the social cost). Walsh [23] considered a similar setting, where one or more facilities can only be placed at different subintervals of the line, and showed bounds on the approximation ratio of strategyproof mechanisms for many social objective functions, beyond the classic ones. Zhao et al. [25] studied a slightly different setting, in which the agents have known approval preferences over the two facilities and their individual costs are defined as their distance from the farthest facility among the ones they approve. For homogeneous agent preferences, they showed a tight bound of 3 for both the social cost and the max cost objectives, while for general, heterogeneous preferences, they showed an upper bound of $2n+1$ for the social cost, and an upper bound of 9 for the maximum cost.

Xu et al. [24] considered a setting where two facilities must be located so that there is a minimum distance between them. They showed results for two types of individual costs. The first one is, as in our case, the total distance (assuming that the facilities play a different role, and thus the agents are interested in both of them) and showed that, for any minimum distance requirement, the optimal solution for the social cost or the maximum cost can be attained by a strategyproof mechanism. The second one is the minimum distance (assuming that the facilities are of the same type, and thus the agents are interested only in their closest one), and showed that the approximation ratio of strategyproof mechanisms is unbounded. They also considered the case where the facility is obnoxious and showed a bound that depends on the minimum distance parameter. Duan et al. [7] later generalized the minimum distance setting by allowing for private fractional preferences over the two facilities.

2 Preliminaries

We consider the two-facility location problem with candidate locations. An instance I of this problem consists of a set N of $n \geq 2$ agents and two facilities $\{F_1, F_2\}$. Each agent $i \in N$ has a *private position* $x_i \in \mathbb{R}$ on the line of real numbers, and a *known approval preference* $p_{ij} \in \{0, 1\}$ for each $j \in [2]$, indicating whether she approves facility F_j ($p_{ij} = 1$) or not ($p_{ij} = 0$), such that $p_{i1} + p_{i2} \geq 1$. There is also a set of $m \geq 2$ candidate locations C where the facilities can be located. To be concise, we denote an instance using the tuple $I = (\mathbf{x}, \mathbf{p}, C)$, where $\mathbf{x} = (x_i)_{i \in N}$ is the *position profile* of all agent positions, and $\mathbf{p} = (p_{ij})_{i \in N, j \in [2]}$ is the *preference profile* of all agent approval preferences.

A *feasible solution* (or, simply, *solution*) is a pair $\mathbf{c} = (c_1, c_2) \in C^2$ of candidate locations with $c_1 \neq c_2$, where the two facilities can be placed; that is, for each $j \in [2]$, F_j is placed at c_j. A *mechanism* M takes as input an instance I of the problem and outputs a feasible solution $M(I)$. Our goal is to design mechanisms so that (a) some objective social function is (approximately) optimized, and (b) the agents truthfully report their private positions.

The *individual cost* of an agent $i \in N$ for a solution $\mathbf{c}$ is her total distance from the locations of the facilities she approves:

$$\text{cost}_i(\mathbf{c}|I) = \sum_{j \in [2]} p_{ij} \cdot d(x_i, c_j),$$

where $d(x, y) = |x - y|$ denotes the distance between any two points x and y on the line. Since the line is a special metric space, the distances satisfy the *triangle inequality*, which states that $d(x, y) \leq d(x, z) + d(z, y)$ for any three points x, y and z on the line, with the equality being true when $z \in [x, y]$. We consider the following two natural objective social functions that have been considered extensively within the truthful facility location literature:

- The *social cost* of a solution $\mathbf{c}$ is the total individual cost of the agents:

$$\text{SC}(c|I) = \sum_{i \in N} \text{cost}_i(\mathbf{c}).$$

- The *max cost* of a solution $\mathbf{c}$ is the maximum individual cost over all agents:

$$\text{MC}(c|I) = \max_{i \in N} \text{cost}_i(\mathbf{c}).$$

The *approximation ratio* of a mechanism M in terms of an objective social function $f \in \{\text{SC}, \text{MC}\}$ is the worst-case ratio (over all possible instances) of the f-value of the solution computed by the mechanism over the minimum possible f-value over all possible solutions:

$$\sup_I \frac{f(M(I)|I)}{\min_{\mathbf{c} \in C^2} f(c|I)}.$$

A mechanism is said to be *strategyproof* if the solution $M(I)$ it returns when given as input any instance $I = (\mathbf{x}, \mathbf{p}, C)$ is such that there is no agent i with incentive to misreport a position $x'_i \neq x_i$ to decrease her individual cost, that is,

$$\text{cost}_i(M(I)|I) \leq \text{cost}_i(M((x'_i, \mathbf{x}_{-i}), \mathbf{p}, C)|I),$$

where $(x'_i, \mathbf{x}_{-i})$ is the position profile obtained by $\mathbf{x}$ when only agent i reports a different position x'_i.

Finally, let us introduce some further notation and terminology that will be useful. For each $j \in [2]$, we denote by N_j the set of agents that approve facility F_j, i.e., $i \in N_j$ if $p_{ij} = 1$. Any agent that approves both facilities belongs to the intersection $N_1 \cap N_2$. We pay particular attention to the subclass of instances in which all agents approve both facilities, i.e., instances with $N = N_1 \cap N_2$. Since all agents have the same preference for both facilities, we will refer to such instances as *homogeneous*; the remaining instances are part of the general class of all instances. We will also denote by m_j, ℓ_j, and r_j the (leftmost) median, leftmost, and rightmost, respectively, agent in N_j. In addition, for any agent i we denote by $t(i)$ and $s(i)$ the closest and the second closest, respectively, candidate location to i.

3 Social Cost

3.1 Homogeneous Instances

We start by considering homogeneous instances, in which all agents approve both facilities, and show a tight bound of $1 + \sqrt{2}$ on the approximation ratio of strategyproof mechanisms in terms of the social cost.

Before we show the best possible approximation ratio, let us first discuss for a bit a straightforward idea. Recall that for the single-facility location problem, Feldman et al. [9] previously showed that the best possible approximation ratio of 3 is achieved by the MEDIAN mechanism, which places the facility at the candidate location closest to the position reported by the median agent m. We can generalize this mechanism by simply placing the two facilities at the two candidate locations that are closest to the position reported by m; that is, F_1 is placed at $w_1 = t(m)$ and F_2 is placed at $w_2 = s(m)$. It is not hard to show that this mechanism still achieves an approximation ratio of at most 3, even when there are two facilities, but cannot do better.

Theorem 1. *For homogeneous instances,* MEDIAN *achieves an approximation ratio of at most* 3*, and this is tight.*

Proof. Let $\mathbf{o} = (o_1, o_2)$ be an optimal solution. Since the position of the median agent minimizes the total distance of all agents, we have that

$$\sum_{i \in N} d(i, m) \leq \sum_{i \in N} d(i, x)$$

for any point x of the line (including o_1 and o_2), and thus

$$2\sum_{i\in N} d(i,m) \leq \sum_{i\in N} d(i,o_1) + \sum_{i\in N} d(i,o_2) = \mathrm{SC}(\mathbf{o}).$$

Also, since $t(m)$ and $s(m)$ are the two closest candidate locations to m, we have that $d(m,t(m)) \leq d(m,x)$ for any candidate location x, and there exists $o \in \{o_1,o_2\}$ such that $d(m,s(m)) \leq d(m,o)$; let $\tilde{o} \in \{o_1,o_2\} \setminus \{o\}$. Therefore, using these facts and the triangle inequality, we obtain

$$\begin{aligned}
\mathrm{SC}(\mathbf{w}) &= \sum_{i\in N} \Big(d(i,t(m)) + d(i,s(m)) \Big) \\
&\leq 2\sum_{i\in N} d(i,m) + \sum_{i\in N} d(m,t(m)) + \sum_{i\in N} d(m,s(m)) \\
&\leq \mathrm{SC}(\mathbf{o}) + \sum_{i\in N} d(m,\tilde{o}) + \sum_{i\in N} d(m,o) \\
&\leq \mathrm{SC}(\mathbf{o}) + 2\sum_{i\in N} d(i,m) + \sum_{i\in N} d(i,\tilde{o}) + \sum_{i\in N} d(i,o) \\
&\leq 3 \cdot \mathrm{SC}(\mathbf{o}).
\end{aligned}$$

The analysis of the mechanism is essentially tight due to the following instance: There are four candidate locations at 0, ε, $1-\varepsilon$, and 1, for some infinitesimal $\varepsilon > 0$. There are also two agents positioned at $1/2-\varepsilon$ and 1, respectively. Let the first agent be the median one (in case the second agent is the median, there is a symmetric instance). Then, the two facilities are placed at 0 and ε for a social cost of approximately $3/2$, whereas the optimal solution is to place the facilities at $1-\varepsilon$ and 1 for a social cost of approximately $1/2$, leading to a lower bound of nearly 3. □

To improve upon the bound of the MEDIAN mechanism, we focus on a family of parameterized mechanisms, which, for a parameter $\alpha \in (0,1/2)$, place one facility at the candidate location closest to the position reported by the αn-leftmost agent, and the other facility at the available candidate location closest to the position reported by the $(1-\alpha)n$-leftmost agent[1]. We refer to such mechanisms as α-STATISTIC; see Mechanism 1 for a description using pseudocode.

Theorem 2. *For any $\alpha \in (0,1/2)$, α-STATISTIC is strategyproof.*

Proof. It is not hard to see that no agent has an incentive to misreport as either the outcome remains the same, or at least one of the facilities is moved farther away from the agent. For example, let x_i and x_j be the positions of agents i

[1] Formally, it would be the $\lceil \alpha n \rceil$-leftmost and the $\lceil (1-\alpha)n \rceil$-leftmost agent, respectively, and we require that $\lceil \alpha n \rceil < \lceil (1-\alpha)n \rceil$. This can be guaranteed by creating an identical number of copies for each agent and running the mechanism on the modified instance; the approximation ratio for the modified instance is exactly the same as for the original instance. We omit the ceilings to make the exposition clearer.

Mechanism 1: α-Statistic

Input: Reported positions of agents;
Output: Facility locations $\mathbf{w} = (w_1, w_2)$;
$i \leftarrow \alpha n$-leftmost agent;
$j \leftarrow (1-\alpha)n$-leftmost agent;
$w_1 \leftarrow t(i)$;
if $t(j)$ is available **then**
  $w_2 \leftarrow t(j)$;
else
  $w_2 \leftarrow s(j)$;

and j, respectively, and consider an agent z whose true position x_z is such that $x_z \leq x_i$. Then, reporting another position $x'_z < x_i$ does not change the outcome at all, while reporting a position $x'_z > x_i$ may move either both facilities (in case $x'_z > x_j$) or facility F_1 (when $x'_z < x_j$) farther away from x_z; in any case, agent z has no incentive to misreport. □

Theorem 3. *For homogeneous instances,* $\left(\sqrt{2}-1\right)$-Statistic *achieves an approximation ratio of at most* $1+\sqrt{2}$.

We next show that the bound of $1+\sqrt{2}$ is the best possible that any strategyproof mechanism can achieve for homogeneous instances.

Theorem 4. *For homogeneous instances, the approximation ratio of any strategyproof mechanism is at least* $1+\sqrt{2}-\delta$, *for any* $\delta > 0$.

Proof. Let $\varepsilon > 0$ be an infinitesimal. To show the theorem, we will consider instances with four candidate locations, two in the ε-neighborhood of 0 (for example, $-\varepsilon$ and ε) and two in the ε-neighborhood of 2 (for example, $2-\varepsilon$ and $2+\varepsilon$). To simplify the calculations in the remainder of the proof, we will assume that there can be candidate locations at the same point of the line, so that we have two candidate locations at 0 and two at 2.

First, consider the following generic instance I with the aforementioned candidate locations: There is at least one agent at 0, at least one agent at 2, while each remaining agent is arbitrarily located at a location from $\{0, 1-\varepsilon, 1+\varepsilon, 2\}$. We make the following observation: Any solution returned by a strategyproof mechanism when given as input I must also be returned when given as input any of the following two instances:

- J_1: Same as I with the difference that an agent j_1 has been moved from 0 to $1-\varepsilon$.
- J_2: Same as I with the difference that an agent j_2 has been moved from 2 to $1+\varepsilon$.

Suppose towards a contradiction that this is not true for J_1; similar arguments can be used for J_2. We consider the following cases:

- Both facilities are placed at 2 in I. If this is not done in J_1, then j_1 can misreport her position as $1-\varepsilon$ in I so that the instance becomes J_1 and at least one facility moves at her true position 0.
- Both facilities are placed at 0 in I. If this is not done in J_1, then j_1 can misreport her position as 0 in J_1 so that the instance becomes I and both facilities move to 0 which is closer to her true position $1-\varepsilon$, a contradiction.
- One facility is placed at 0 and the other is placed at 2 in I. Observe that it cannot be the case that both facilities are placed at 0 in J_1 since that would mean that j_1 can misreport her position in I as $1-\varepsilon$ so that the instance becomes J_1 and both facilities move to her true position 0. So, the only possibility of having a different solution in I and J_1 is that both facilities are placed at 2 in J_1. But then, j_1 can misreport her true position as 0 in J_1 so that the instance becomes I and one of the facilities moves to 0 which is closer to her true position.

Hence, the same solution must be computed by the mechanism when given I or J_1 as input.

Now, consider an arbitrary strategyproof mechanism and let $\alpha = \sqrt{2}-1$; recall that α is such that $\frac{1+\alpha}{1-\alpha} = \frac{1}{\alpha} = 1+\sqrt{2}$. Let I_1 be the following instance with the aforementioned candidate locations: αn agents are at 0 and $(1-\alpha)n$ agents are at 2. We consider the following cases depending on the solution returned by the mechanism when given I_1 as input:

Case 1: The mechanism places both facilities at 0. We consider the sequence of instances obtained by moving one by one the αn agents that are positioned at 0 in I_1 to $1-\varepsilon$. By the observation above, the mechanism must return the same solution for any two consecutive instances of this sequence (essentially, the first one is of type I and the second one is of type J_1), which means that the mechanism must eventually return the same solution for all of them. Therefore, the mechanism must place both facilities at 0 in the last instance of this sequence, where αn agents are at $1-\varepsilon$ and the remaining $(1-\alpha)n$ agents are at 2. This solution has social cost $2\alpha n + 4(1-\alpha)n = 2(2-\alpha)n$. However, the solution that places both facilities at 2 has social cost $2\alpha n$, leading to an approximation ratio of $\frac{2}{\alpha} - 1 > 1+\sqrt{2}$.

Case 2: The mechanism places both facilities at 2. Similarly to Case 1 above, we now consider the sequence of instances obtained by moving one by one the $(1-\alpha)n$ agents that are positioned at 2 in I_1 to $1+\varepsilon$. Again, by the observation above, the mechanism must return the same solution for any two consecutive instances of this sequence (the first one is of type I and the second one is of type J_2), which means that the mechanism must eventually return the same solution for all of them. Therefore, the mechanism must place both facilities at 2 in the last instance of this sequence, where αn agents are at 0 and the remaining $(1-\alpha)n$ agents are at $1+\varepsilon$. This solution has social cost $4\alpha n + 2(1-\alpha)n = 2(1+\alpha)n$. However, the solution that places both facilities at 0 has social cost $2(1-\alpha)n$, leading to an approximation ratio of $\frac{1+\alpha}{1-\alpha} = 1+\sqrt{2}$.

Case 3: The mechanism places one facility at 0 and the other at 2. We consider the same sequence of instances as in Case 1. This results in that the mechanism must place one facility at 0 and the other at 2 when given as input the instance where αn agents are at $1-\varepsilon$ while the remaining $(1-\alpha)n$ agents are at 2. This solution has social cost $2\alpha n + 2(1-\alpha)n = 2n$. However, the solution that places both facilities at 2 has social cost $2\alpha n$, leading to an approximation ratio of $\frac{1}{\alpha} = 1 + \sqrt{2}$. □

3.2 General Instances

Having completely resolved the homogeneous case, we now turn our attention to general, possibly heterogeneous instances. It is not hard to observe that the problem is more general than the single-facility location problem studied by Feldman et al. [9]; indeed, there are instances in which all agents approve just one facility and the other facility does not affect their social cost nor the approximation ratio. Consequently, we cannot hope to achieve an approximation ratio better than 3.

Theorem 5. *For general instances, the approximation ratio of any strategyproof mechanism is at least* $3-\delta$*, for any* $\delta > 0$.

Since there is an adaptation of the MEDIAN mechanism that achieves an approximation ratio of 3 for homogeneous instances (see Theorem 1), one might wonder if there is a variant that can do so for heterogeneous instances as well. In particular, the natural extension of MEDIAN is to place F_1 at the candidate location closest to the (leftmost) median agent m_1 of N_1, and F_2 at the available candidate location closest to the (leftmost) median agent m_2 of N_2. While this seems like a good idea at first glance, the following example shows that it fails to achieve the desired approximation ratio bound.

Example 1. Consider an instance with two candidate locations at 0 and 2. For some $x \geq 1$, there are $2x+1$ agents that approve only F_1 such that $x+1$ of them are located at $1-\varepsilon$ and the other x are located at 2. There are also $2x+1$ agents that approve only F_2 and are all located at 0. According to the definition of the mechanism, F_1 is placed at 0 (which is the candidate location closest to the median agent in N_1), and then F_2 is placed at 2 as 0 is now occupied and 2 is available. This solution has social cost approximately $(x+2x)+4x = 7x$, whereas the solution that places F_1 at 2 and F_2 at 0 has social cost approximately x, leading to an approximation ratio of nearly 7.

The issue with aforementioned variant of the MEDIAN mechanism is the order in which it decides to place the facilities. If it were to place F_2 first and F_1 second then it would have made the optimal choice in the example. However, there is a symmetric example that would again lead to a lower bound of approximately 7. So, the mechanism needs to be able to dynamically determine the order in which it places F_1 and F_2. This brings us to the following idea: We will again place each facility one after the other at the closest candidate location to the

Mechanism 2: PROPORTIONAL-MAJORITY-MEDIAN

Input: Reported positions of agents;
Output: Facility locations $\mathbf{w} = (w_1, w_2)$;
for $j \in [2]$ **do**
  $S_j \leftarrow$ set of agents in N_j (weakly) closer to $t(m_j)$ than to $s(m_j)$;
if $|S_1|/|N_1| \geq |S_2|/|N_2|$ **then**
  $j \leftarrow 1$;
else
  $j \leftarrow 2$;
$w_j \leftarrow t(m_j)$;
if $t(m_{3-j})$ is available **then**
  $w_{3-j} \leftarrow t(m_{3-j})$;
else
  $w_{3-j} \leftarrow s(m_{3-j})$;

median among the agents that approve it. However, the facility that is placed first (and thus has priority in case the median agents of N_1 and N_2 are closer to the same candidate location) is determined by the proportion of agents that approve it who are closer to the top choice of the median agent rather than her second choice; ties are broken in favor of the facility that is approved by most agents, which is assumed to be F_1 without loss of generality. We refer to this mechanism as PROPORTIONAL-MAJORITY-MEDIAN; see Mechanism 2 for a more formal description.

We first show that this mechanism is strategyproof and then bound its approximation ratio. It is not hard to observe that this must be true as the mechanism is a composition of two simple strategyproof mechanisms (median plus majority voting).

Theorem 6. PROPORTIONAL-MAJORITY-MEDIAN *is strategyproof.*

Proof. Clearly, if $t(m_1) \neq t(m_2)$ then no agent has incentive to deviate as then the facilities are placed at $t(m_1)$ and $t(m_2)$ independently of whether $|S_j|/|N_j| \geq |S_{3-j}|/|N_{3-j}|$ or not for $j \in [2]$. So, it suffices to consider the case where $t(m_1) = t(m_2)$ and $|S_j|/|N_j| \geq |S_{3-j}|/|N_{3-j}|$ for some $j \in [2]$, leading to $w_j = t(m_j)$ and $w_{3-j} = s(m_{3-j})$, solving ties in favor of F_1.

- m_j and any agent $i \in N_j$ that is closer to $t(m_j)$ than to $s(m_j)$ have no incentive to deviate as $t(m_j)$ is the best choice for them.
- Any agent $i \in N_j$ that is closer to $s(m_j)$ than to $t(m_j)$ has no incentive to deviate, as going closer to $t(m_j)$ can only increase the ratio $|S_j|/|N_j|$ and cannot change the outcome.
- m_{3-j} and any agent $i \in N_{3-j}$ that is closer to $t(m_{3-j})$ than to $s(m_{3-j})$ have no incentive to deviate as moving closer to $s(m_{3-j})$ would decrease the ratio $|S_{3-j}|/|N_{3-j}|$ and would not change the outcome.
- Any agent $i \in N_{3-j}$ that is closer to $s(m_{3-j})$ than to $t(m_{3-j})$ has no incentive to deviate as $s(m_j)$ is the best choice for them.

Mechanism 3: LEFTMOST-PRIORITY

Input: Reported positions of agents;
Output: Facility locations $\mathbf{w} = (w_1, w_2)$;
$\ell_{12} \leftarrow$ leftmost agent in $N_1 \cap N_2$;
$r_{12} \leftarrow$ rightmost agent in $N_1 \cap N_2$;
$w_1 \leftarrow t(\ell_{12})$;
if $t(r_{12})$ is available **then**
  $w_2 \leftarrow t(r_{12})$;
else
  $w_2 \leftarrow s(r_{12})$;

So, the mechanism is strategyproof. □

Theorem 7. PROPORTIONAL-MAJORITY-MEDIAN *achieves an approximation ratio of at most* 3.

4 Max Cost

In this section, we turn our attention to the max cost objective for which we again show tight bounds on the approximation ratio of deterministic strategyproof mechanisms. In particular, we show a bound of 2 for homogeneous instances and a bound of 3 for general instances.

4.1 Homogeneous Instances

For homogeneous instances, we consider a simple mechanism that places F_1 at the candidate location closest to the leftmost agent ℓ_{12}, and then F_2 at the available candidate location closest to the rightmost agent r_{12}. We refer to this mechanism as LEFTMOST-PRIORITY; see Mechanism 3. It is not hard to show that this mechanism is strategyproof and that it achieves a best possible approximation ratio of 2 among deterministic strategyproof mechanisms.

Theorem 8. *For homogeneous instances,* LEFTMOST-PRIORITY *is strategyproof.*

Proof. Consider any agent i; recall that i approves both facilities. To affect the outcome, agent i would have to either report a position that lies at the left of ℓ_{12} or at the right of r_{12}. Changing the leftmost or rightmost positions, however, can only lead to placing the facilities at locations farther away from i and, hence, i has no incentive to misreport. □

Theorem 9. *For homogeneous instances,* LEFTMOST-PRIORITY *achieves an approximation ratio of at most* 2.

Proof. Let i be the agent that determines the maximum cost of the mechanism, i.e., $\mathrm{MC}(\mathbf{w}) = d(i, w_1) + d(i, w_2)$. Clearly, $i \in \{\ell_{12}, r_{12}\}$; let $j \in \{\ell_{12}, r_{12}\} \setminus \{i\}$.

Let $\mathbf{o} = (o_1, o_2)$ be an optimal solution. We consider the following two cases: If $w_1 = t(\ell_{12})$ and $w_2 = t(r_{12})$, then by the triangle inequality and the definition of $t(\cdot)$, we have

$$\begin{aligned}\mathrm{MC}(\mathbf{w}) &= d(i, t(i)) + d(i, t(j)) \\ &\leq d(i, t(i)) + d(i, o_2) + d(j, o_2) + d(j, t(j)) \\ &\leq d(i, o_1) + d(i, o_2) + d(j, o_2) + d(j, o_1) \leq 2 \cdot \mathrm{MC}(\mathbf{o}).\end{aligned}$$

If $w_1 = t(\ell_{12})$ and $w_2 = s(r_{12})$, then there exists $x \in \{o_1, o_2\}$ such that $d(r_{12}, w_2) \leq d(r_{12}, x)$. Let $y \in \{o_1, o_2\} \setminus \{x\}$.

- If $i = \ell_{12}$, then

$$\begin{aligned}\mathrm{MC}(\mathbf{w}) &= d(\ell_{12}, w_1) + d(\ell_{12}, w_2) \\ &\leq d(\ell_{12}, w_1) + d(\ell_{12}, y) + d(r_{12}, y) + d(r_{12}, w_2) \\ &\leq d(\ell_{12}, x) + d(\ell_{12}, y) + d(r_{12}, y) + d(r_{12}, x) \leq 2 \cdot \mathrm{MC}(\mathbf{o}).\end{aligned}$$

- If $i = r_{12}$, then

$$\begin{aligned}\mathrm{MC}(\mathbf{w}) &= d(r_{12}, w_1) + d(r_{12}, w_2) \\ &\leq d(r_{12}, y) + d(\ell_{12}, y) + d(\ell_{12}, w_1) + d(r_{12}, w_2) \\ &\leq d(r_{12}, y) + d(\ell_{12}, y) + d(\ell_{12}, x) + d(r_{12}, x) \leq 2 \cdot \mathrm{MC}(\mathbf{o}).\end{aligned}$$

Therefore, the approximation ratio is at most 2 in any case. □

Theorem 10. *For homogeneous instances, the approximation ratio of any deterministic strategyproof mechanism is at least* $2 - \delta$*, for any* $\delta > 0$*.*

4.2 General Instances

To tackle the general case, we split the set of instances into those with at least one agent that approves both facilities, and those in which each agent approves just one facility. It is not hard to observe that Leftmost-Priority can still be applied when there is at least one agent that approves both facilities, and is strategyproof since its decision is fully determined by the leftmost and rightmost such agents. In addition, Leftmost-Priority can be shown to achieve an approximation ratio of at most 3.

Theorem 11. *For instances with at least one agent that approves both facilities, the approximation ratio of* Leftmost-Priority *is at most* 3*.*

Proof. We consider cases depending on the preference of the agent i that determines the max cost of the mechanism.

(Case 1) The max cost is determined by an agent $i \in N_1 \setminus N_2$. Then, by the triangle inequality and since $d(\ell_{12}, w_1) \leq d(\ell_{12}, o_2)$, we have

$$\begin{aligned}\mathrm{MC}(\mathbf{w}) = d(i, w_1) &\leq d(i, o_1) + d(\ell_{12}, o_1) + d(\ell_{12}, w_1) \\ &\leq d(i, o_1) + d(\ell_{12}, o_1) + d(\ell_{12}, o_2) \\ &\leq 2 \cdot \mathrm{MC}(\mathbf{o}).\end{aligned}$$

(Case 2) The max cost is determined by an agent $i \in N_2 \setminus N_1$. Since w_2 is either $t(r_{12})$ or $s(r_{12})$, there exists $x \in \{o_1, o_2\}$ such that $d(r_{12}, w_2) \leq d(r_{12}, x) \leq \mathrm{MC}(\mathbf{o})$. Hence, by the triangle inequality, we have

$$\mathrm{MC}(\mathbf{w}) = d(i, w_2) \leq d(i, o_2) + d(r_{12}, o_2) + d(r_{12}, w_2) \leq 3 \cdot \mathrm{MC}(\mathbf{o}).$$

(Case 3) The max cost is determined by an agent $i \in N_1 \cap N_2$. Then, following the proof of Theorem 9 for homogeneous instances, we can show an upper bound of 2. □

Next, we focus on instances with only agents that have singleton preferences, that is, $N_1 \cap N_2 = \varnothing$. There are multiple ways for LEFTMOST-PRIORITY to work in such instances, for example, by placing F_1 at the closest candidate location to the leftmost agent ℓ_1 of N_1, and F_2 at the closest available candidate location to an agent of N_2, such as the leftmost agent ℓ_2 or the rightmost agent r_2. While such mechanisms are clearly strategyproof, it is not hard to observe that they cannot achieve a good enough approximation ratio.

Example 2. If we place F_1 at $t(\ell_1)$ and F_2 at $t(r_2)$ or $s(r_2)$ depending on availability, then consider the following instance: There are three candidate locations at 0, 2, and 6. There is an agent ℓ_1 that approves F_1 at $1 + \varepsilon$, an agent ℓ_2 that approves F_2 at 1, and another agent r_2 that approves F_2 at $3 + \varepsilon$, for some infinitesimal $\varepsilon > 0$. So, we place F_1 at 2 and F_2 at 6 for a max cost of 5 (determined by ℓ_2). On the other hand, we could place F_1 at 0 and F_2 at 2 for a max cost of approximately 1, leading to an approximation ratio of 5. Clearly, if we chose ℓ_2 instead of r_2 to determine the location of F_2, there is a symmetric instance leading again to the same lower bound.

The above example illustrates that it is not always a good idea to choose a priori ℓ_2 or r_2 to determine where to place F_2, especially when the closest candidate location to them might not be available after placing F_1. Instead, we need to carefully decide whether ℓ_2 or r_2 or neither of them is the best one to choose where to place F_2. We make this decision as follows: We "ask" ℓ_2 and r_2 to "vote" over two candidate locations; the candidate location L that is the closest at the left of $t(\ell_1)$ (where F_1 is placed) and the candidate location R that is the closest at the right of $t(\ell_1)$. If ℓ_2 and r_2 "agree", then they are both on the same side of the midpoint of the interval defined by L and R, and thus depending on whether they are on the left side (agree on L) or the right side (agree on R), we allow r_2 or ℓ_2, respectively, to make the choice of where to place F_2. If they "disagree", they are on different sides of the interval's midpoint, so neither ℓ_2 nor r_2 should make a choice of where to place F_2; in this case, the closest of L and R to $t(\ell_1)$ is a good candidate location to place F_2. This idea is formalized in Mechanism 4, which we call VOTE-FOR-PRIORITY.

Theorem 12. *For instances with only agents that have singleton preferences,* VOTE-FOR-PRIORITY *is strategyproof and achieves an approximation ratio of at most* 3*.*

Mechanism 4: VOTE-FOR-PRIORITY

Input: Reported positions of agents;
Output: Facility locations $\mathbf{w} = (w_1, w_2)$;
$w_1 \leftarrow t(\ell_1)$;
$L \leftarrow$ closest candidate location at the left of w_1;
$R \leftarrow$ closest candidate location at the right of w_1;
// **(case 1)** ℓ_2 and r_2 agree that L is closer: r_2 gets to choose
if ℓ_2 and r_2 are both closer to L than to R **then**
  if $t(r_2)$ is available **then**
    $w_2 \leftarrow t(r_2)$;
  else
    $w_2 \leftarrow s(r_2)$;
// **(case 2)** ℓ_2 and r_2 agree that R is closer: ℓ_2 gets to choose
else if ℓ_2 and r_2 are both closer to R than to L **then**
  if $t(\ell_2)$ is available **then**
    $w_2 \leftarrow t(\ell_2)$;
  else
    $w_2 \leftarrow s(\ell_2)$;
// **(case 3)** ℓ_2 and r_2 disagree: choose the closest of L and R to w_1
else
  $w_2 \leftarrow \arg\min_{x \in \{L,R\}} \{|w_1 - x|\}$;

Based on the above discussion, we have the following statement.

Corollary 1. *For general instances, there is a deterministic strategyproof mechanism with approximation ratio at most* 3*.*

Proof. Since the preferences of the agents are known, we can determine whether there is at least one agent that approves both facilities. If this is indeed the case, we run LEFTMOST-PRIORITY. Otherwise, we run VOTE-FOR-PRIORITY. The desired properties of the mechanism follow by the fact that both mechanisms are strategyproof and achieve approximation ratio at most 3 in their respective cases. □

The bound of 3 is in fact the best possible among all deterministic strategyproof mechanisms. As argued at the beginning of Section 3.2, instances in which all agents approve one of the facilities are equivalent to having just this one facility to place. Consequently, by the work of Tang et al. [22], we cannot hope to achieve an approximation ratio better than 3.

Theorem 13. *For general instances, the approximation ratio of any deterministic strategyproof mechanism is at least* $3 - \delta$*, for any* $\delta > 0$*.*

5 Conclusion

In this paper we studied a truthful two-facility location problem with candidate locations and showed tight bounds on the best possible approximation ratio of

deterministic strategyproof mechanisms in terms of the social cost and the max cost. An obvious question that our work leaves open is the design of randomized strategyproof mechanisms with improved approximation guarantees. There are also multiple ways to extend our model. One such way is to change the assumption about what type of information is public or private, and instead of considering the case where the positions are private and the preferences are known as we did in this paper, focus on the case where the positions are public and the preferences are private (which is a generalization of the models studied in [20, 15]). Other ways of extending our model include settings with more than just two facilities to place, and more general metric spaces than just the line.

References

1. Alon, N., Feldman, M., Procaccia, A.D., Tennenholtz, M.: Strategyproof approximation of the minimax on networks. Mathematics of Operations Research **35**(3), 513–526 (2010)
2. Anastasiadis, E., Deligkas, A.: Heterogeneous facility location games. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS). pp. 623–631 (2018)
3. Anshelevich, E., Filos-Ratsikas, A., Shah, N., Voudouris, A.A.: Distortion in social choice problems: The first 15 years and beyond. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI). pp. 4294–4301 (2021)
4. Chan, H., Filos-Ratsikas, A., Li, B., Li, M., Wang, C.: Mechanism design for facility location problems: A survey. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI). pp. 4356–4365 (2021)
5. Cheng, Y., Yua, W., Zhang, G.: Strategy-proof approximation mechanisms for an obnoxious facility game on networks. Theoretical Computer Science **497**, 154–163 (2013)
6. Deligkas, A., Filos-Ratsikas, A., Voudouris, A.A.: Heterogeneous facility location with limited resources. Games and Economic Behavior **139**, 200–215 (2023)
7. Duan, L., Gong, Z., Li, M., Wang, C., Wu, X.: Facility location with fractional preferences and minimum distance. In: Proceedings of the 27th International Conference on Computing and Combinatorics (COCOON). pp. 499–511 (2021)
8. Feigenbaum, I., Sethuraman, J.: Strategyproof mechanisms for one-dimensional hybrid and obnoxious facility location models. In: AAAI Workshop on Incentive and Trust in E-Communities. vol. WS-15-08 (2015)
9. Feldman, M., Fiat, A., Golomb, I.: On voting and facility location. In: Proceedings of the 2016 ACM Conference on Economics and Computation (EC). pp. 269–286 (2016)
10. Filos-Ratsikas, A., Kanellopoulos, P., Voudouris, A.A., Zhang, R.: Settling the distortion of distributed facility location. In: Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 2152–2160 (2023)
11. Filos-Ratsikas, A., Voudouris, A.A.: Approximate mechanism design for distributed facility location. In: Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT). pp. 49–63 (2021)
12. Fotakis, D., Tzamos, C.: Strategyproof facility location for concave cost functions. Algorithmica **76**(1), 143–167 (2016)

13. Gai, L., Liand, M., Wang, C.: Obnoxious facility location games with candidate locations. In: Proceedings of the 16th International Conference on Algorithmic Aspects in Information and Management (AAIM). pp. 96–105 (2022)
14. Goel, S., Hann-Caruthers, W.: Coordinate-wise median: Not bad, not bad, pretty good. CoRR **abs/2007.00903** (2020)
15. Kanellopoulos, P., Voudouris, A.A., Zhang, R.: On discrete truthful heterogeneous two-facility location. SIAM Journal on Discrete Mathematics **37**, 779–799 (2023)
16. Li, M., Lu, P., Yao, Y., Zhang, J.: Strategyproof mechanism for two heterogeneous facilities with constant approximation ratio. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI). pp. 238–245 (2020)
17. Lu, P., Sun, X., Wang, Y., Zhu, Z.A.: Asymptotically optimal strategy-proof mechanisms for two-facility games. In: Proceedings of the 11th ACM Conference on Electronic Commerce (EC). pp. 315–324 (2010)
18. Meir, R.: Strategyproof facility location for three agents on a circle. In: Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT). vol. 11801, pp. 18–33 (2019)
19. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. ACM Transactions on Economics and Computation **1**(4), 18:1–18:26 (2013)
20. Serafino, P., Ventre, C.: Heterogeneous facility location without money. Theoretical Computer Science **636**, 27–46 (2016)
21. Sui, X., Boutilier, C.: Approximately strategy-proof mechanisms for (constrained) facility location. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 605–613 (2015)
22. Tang, Z., Wang, C., Zhang, M., Zhao, Y.: Mechanism design for facility location games with candidate locations. In: Proceedings of the 14th International Conference on Combinatorial Optimization and Applications (COCOA). pp. 440–452 (2020)
23. Walsh, T.: Strategy proof mechanisms for facility location at limited locations. In: Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence (PRICAI). pp. 113–124 (2021)
24. Xu, X., Li, B., Li, M., Duan, L.: Two-facility location games with minimum distance requirement. Journal of Artificial Intelligence Research **70**, 719–756 (2021)
25. Zhao, Q., Liu, W., Nong, Q., Fang, Q.: Constrained heterogeneous facility location games with max-variant cost. Journal on Combinatorial Optimization **45**(3), 90 (2023)
26. Zou, S., Li, M.: Facility location games with dual preference. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 615–623 (2015)