



Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/gebHeterogeneous facility location with limited resources [☆]Argyrios Deligkas ^a, Aris Filos-Ratsikas ^b, Alexandros A. Voudouris ^{c,*}^a Department of Computer Science, Royal Holloway University of London, UK^b School of Informatics, University of Edinburgh, UK^c School of Computer Science and Electronic Engineering, University of Essex, UK

ARTICLE INFO

Article history:

Received 28 March 2022

Available online 8 March 2023

Keywords:

Facility location

Heterogeneous preferences

Voting

ABSTRACT

We initiate the study of the heterogeneous facility location problem with limited resources. We mainly focus on the fundamental case where a set of agents are positioned in the line segment $[0, 1]$ and have approval preferences over two available facilities. A mechanism takes as input the positions and the preferences of the agents, and chooses to locate a *single* facility based on this information. We study mechanisms that aim to maximize the social welfare (the total utility the agents derive from facilities they approve), under the constraint of incentivizing the agents to truthfully report their positions and preferences. We consider three different settings depending on the level of agent-related information that is public or private. For each setting, we design deterministic and randomized strategyproof mechanisms that achieve a good approximation of the optimal social welfare, and complement these with nearly-tight impossibility results.

© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The truthful facility location problem is one of the most prominent paradigms in environments with strategic participants, and it was in fact the prototypical problem used by Procaccia and Tennenholtz (2013) to put forward their very successful research agenda of *approximate mechanism design without money* about a decade ago. Since then, the problem has been extensively studied in the literature of theoretical computer science and artificial intelligence, with a plethora of interesting variants emerging over the years. Among those, one particularly meaningful variant, which captures several important scenarios, is that of *heterogeneous facility location*, introduced by Feigenbaum and Sethuraman (2015) and studied notably by Serafino and Ventre (2015, 2016), Anastasiadis and Deligkas (2018), Fong et al. (2018), Chen et al. (2020) and Li et al. (2020a). In this setting, there are multiple facilities, and each of them plays a different role – for example, a library and a basketball court. Consequently, the preferences of the agents for the possible outcomes do not only depend on the *location* of the facility (as in the original model of Procaccia and Tennenholtz (2013)), but also on the *type* of the facility. As a result, the mechanism design problem now becomes far more challenging.¹

While the literature on heterogeneous facility location is quite rich by this point, there is a fundamental setting that has surprisingly eluded previous investigations. In particular, all previous works have considered the case of multiple (pre-

[☆] A preliminary version of this paper appears in Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), 2022.

* Corresponding author.

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

¹ In particular, the preference domain is no longer *single-peaked*, and therefore maximizing the happiness of the agents cannot be achieved by simple median mechanisms.

Table 1

Overview of our results for deterministic and randomized strategyproof mechanisms. The lower bound $4/3$ (marked with \star) in the known-preferences setting holds only for the class of Random-Median mechanisms defined in Section 4. For the general and known-preference settings, the bound of 2 also holds for the more general case where we can choose k out of $m \geq 2$ facilities, for appropriate values of k and m .

	Deterministic	Randomized
General	2	(1, 2]
Known-preferences	2	[$4/3^\star$, $4/3$]
Known-positions	[13/11, 2]	(1, $3/2$]

dominantly two) facilities which *all* have to be located, based on the positions and the preferences of the agents. However, in many real-world applications, resources are *limited*, and therefore a decision has to be made about which *subset* of the facilities should be build and where. For instance, the governing body might have sufficient funds to build only one of two options, either a library or a basketball court. The decision must be made based on the preferences of the agents over the two facilities, but also on their positions, in a way that incentivizes the agents to reveal all their private information truthfully; this is clearly a challenging mechanism design problem.

1.1. Our setting

We initiate the study of the heterogeneous facility location problem with *limited resources*. We focus on the most fundamental case where there are two facilities, and only one of them must be located somewhere in the line segment $[0, 1]$. In particular, there is a set of agents, each of whom is associated with a *position* in $[0, 1]$ and an *approval preference* over the facilities. An agent may approve one of the two facilities or both, and obtains positive *utility*² only if a facility that she approves is built; otherwise, she has zero utility irrespectively of her position.

Our goal is to design *strategyproof mechanisms* that choose and locate a single facility, so as to maximize the *social welfare* (the total utility of the agents) and incentivize the agents to truthfully report their private information. We study the following three settings depending on the level of information about the positions and the preferences of the agents that is assumed to be public or private.

- *General* setting: Both the positions and the preferences are private information of the agents.
- *Known-preferences* setting: The positions are private information of the agents, whereas the preferences are public information.
- *Known-positions* setting: The preferences are private information of the agents, whereas the positions are public information.

We measure the performance of a strategyproof mechanism by its *approximation ratio*, defined as the worst-case ratio over all instances of the problem between the maximum possible social welfare and the social welfare achieved by the mechanism. For each of the aforementioned settings, we derive upper and lower bounds on the achievable approximation ratio of strategyproof mechanisms. An overview of our results can be found in Table 1.

1.2. Discussion of our results

We start our investigation by studying deterministic mechanisms in the general setting, where we show that a simple group-strategyproof mechanism, which we call MIDDLE, achieves an approximation ratio of 2 (Theorem 3.1); the same guarantee extends to the other two settings we consider. We complement this result by showing a lower bound of 2 on the approximation ratio of *any* deterministic strategyproof mechanism, even when the preferences of the agents are assumed to be known (Theorem 4.1). Combining these two results, we completely resolve the problem of identifying the best possible deterministic strategyproof mechanism for both the general and the known-preferences settings. For the known-positions setting, we show that there is no deterministic strategyproof mechanism with approximation ratio better than $13/11$ (Theorem 5.1).

We also consider randomized mechanisms, and provide improved approximation guarantees for both the known-preferences and the known-positions settings. More specifically, for the known-preferences setting we derive a novel universally group-strategyproof mechanism, termed MIRROR, which achieves an approximation ratio of $4/3$ (Theorem 4.4). This mechanism is in fact a member of a larger class of universally group-strategyproof mechanisms, and as we prove, it

² We remark that in several facility location settings (e.g., see (Procaccia and Tennenholtz, 2013; Lu et al., 2009, 2010)), the agents are associated with costs instead of utilities. In the literature of heterogeneous facility problems however, the setting is commonly defined in terms of utilities, as there is no meaningful way of assigning a cost to undesirable outcomes, such as a facility which the agent does not approve.

is the best possible mechanism in this class (Theorem 4.5). For the known-positions setting, we prove that a variant of the well-known RANDOM DICTATORSHIP mechanism, equipped with a carefully chosen tie-breaking rule for the agents that approve both facilities, is a universally group-strategyproof mechanism (Theorem 5.2) and achieves an approximation ratio of $3/2$ (Theorem 5.6).

Finally, we make initial progress in more general settings with $m \geq 2$ facilities, from which we can choose to locate $k < m$. We show that an adaptation of MIDDLE still has an approximation ratio of 2 in the general setting, it is group-strategyproof for $k = 1$, but it is only strategyproof for $k \geq 2$ (Theorem 6.1 and Lemma 6.2). We complement this result by showing that it is impossible to do better in some cases, even when the preferences of the agents are known (Theorem 6.3).

1.3. Related work

As we mentioned earlier, the literature on truthful facility location is long and extensive; here, we discuss only those works that are most closely related to our setting. The fundamental difference between our work and virtually all of the papers on heterogeneous facility location is that they consider settings with two facilities, where *both* facilities have to be built, and the utility/cost of an agent is calculated with respect to the closest or the farthest among the two.

In particular, Chen et al. (2020) consider a setting in which agents have approval preferences over the facilities, similarly to what we do here, and for which the positions of the agents are known. Li et al. (2020a) consider a more general metric setting along the lines of Chen et al. (2020), and design a deterministic mechanism which improves upon the result of Chen et al. (2020) when the metric is a line. Fong et al. (2018) consider a setting in which the agents have fractional preferences in $(0, 1)$; similarly to us, besides studying the general setting, they also consider restricted settings with known preferences or known positions. Serafino and Ventre (2015, 2016); Kanellopoulos et al. (2021) consider a discrete setting, where the agents are positioned on the nodes of a graph, and the facilities must be located on different nodes. Feigenbaum and Sethuraman (2015) were the first to study heterogeneous facility location, by presenting a “hybrid” model combining the standard facility location problem with the obnoxious facility location problem (Cheng et al., 2011, 2013). This setting was extended by Anastasiadis and Deligkas (2018), who allowed agents to be indifferent between whether a facility would be built or not. Xu et al. (2021) study a setting where the goal is to locate two facilities under the constraint that the distance between the locations of the facilities is at least larger than a predefined bound.

Li et al. (2020b) study a conceptually similar but fundamentally different facility location problem under budget constraints. In their setting, the facilities are *strategic* and need to be compensated monetarily in order for them to be built; the goal is to maximize an aggregate objective given that the total payment is below a predefined budget. Besides these works, there is long literature of (homogeneous) facility location, studying different objectives (Alon et al., 2010; Cai et al., 2016; Feigenbaum et al., 2013; Feldman and Wilf, 2013), multiple facilities (Escoffier et al., 2011; Fotakis and Tzamos, 2013; Lu et al., 2009, 2010), different domains (Schummer and Vohra, 2002; Tang et al., 2020; Sui et al., 2013; Sui and Boutilier, 2015), different cost functions (Filos-Ratsikas et al., 2015; Fotakis and Tzamos, 2016), and several interesting variants (Golomb and Tzamos, 2017; Kyropoulou et al., 2019; Zhang and Li, 2014; Filos-Ratsikas and Voudouris, 2021; Anshelevich et al., 2021b). We refer the reader to the recent survey of Chan et al. (2021) for a detailed overview. We also refer the reader to the survey of Anshelevich et al. (2021a) for an overview of the literature on distortion, which has been applied for analyzing facility location settings.

Finally, it is instructive to explain how our setting fits within the framework of general social choice. The original facility location setting of Procaccia and Tennenholtz (2013) can be seen as a voting setting in which there is a continuum of alternatives on the real line, and the agents have single-peaked preferences over them, given by the distance functions defined by their most-preferred positions. In our heterogeneous facility location setting, there is also another “voting component” in which agents cast approval votes over a different set of alternatives. In that sense, one can view our setting as a social choice scenario over two dimensions, one for the type and one for the position of the facility. On the “type axis”, the 1-out-of- m setting can be seen as a standard single-winner voting setting, whereas the k -out-of- m is a *multiwinner election* setting, e.g., see (Caragiannis et al., 2017, 2022; Faliszewski et al., 2017). Very recently, Elkind et al. (2022) also considered such a social choice scenario, which is however quite different from what we do in this paper. Besides asking different types of questions, one crucial conceptual difference is that, in the model of Elkind et al. (2022), the approval preferences of the agents over the facilities depend on the locations of the facilities.

2. Preliminaries

We consider a facility location setting with a set N of n agents and two facilities; we will discuss extensions to settings with more than two facilities in Section 6. Every agent $i \in N$ has a position $x_i \in [0, 1]$; let $\mathbf{x} = (x_i)_{i \in N}$ be the position profile consisting of the positions of all agents. Furthermore, every agent $i \in N$ also has an approval preference (or, simply, preference) $\mathbf{t}_i = \{0, 1\}^2$ over the two facilities, where $t_{ij} = 1$ denotes that the agent approves facility $j \in \{1, 2\}$ and $t_{ij} = 0$ denotes that she does not approve facility j ; let $\mathbf{t} = (\mathbf{t}_i)_{i \in N}$ be the preference profile consisting of the preferences of all agents. Let $I = (\mathbf{x}, \mathbf{t})$ denote an instance of this setting.

Given an instance $I = (\mathbf{x}, \mathbf{t})$, our goal is to choose and locate a single facility so as to optimize some objective function that depends on both the distances of the agents from the facility location and on whether they approve the chosen facility. In particular, if facility $j \in \{1, 2\}$ is chosen to be located at $y \in [0, 1]$, the utility of every agent $i \in N$ is defined to be

$u_i(j, y|I) = t_{ij} \cdot (1 - d(x_i, y))$, where $d(x_i, y) = |x_i - y|$ is the distance between x_i and y . Then, the *social welfare* is the sum of the utilities of all agents:

$$W(j, y|I) := \sum_{i \in N} u_i(j, y|I).$$

We denote the *optimal* social welfare for instance I as $W^*(I) := \max_{(j,y)} W(j, y|I)$.

A *mechanism* M takes as input an instance $I = (\mathbf{x}, \mathbf{t})$ consisting of the position and preference profiles of the agents, and outputs an *outcome* $M(I) = (j_M, y_M)$ consisting of a facility $j_M \in \{1, 2\}$ that is to be placed at $y_M \in [0, 1]$. The *approximation ratio* $\rho(M)$ of M is defined as the worst-case ratio (over all possible instances) between the optimal social welfare and the social welfare of the outcome chosen by the mechanism, that is,

$$\rho(M) = \sup_I \frac{W^*(I)}{W(M(I)|I)}.$$

A mechanism is *strategyproof* if it is in the best interest of every agent to report their true position and preferences, irrespectively of the reports of the other agents. Formally, a mechanism M is strategyproof if, for every pair of instances $I = (\mathbf{x}, \mathbf{t})$ and $I' = ((x'_i, \mathbf{x}_{-i}), (\mathbf{t}'_i, \mathbf{t}_{-i}))$ in which only a single agent i misreports a different position and preferences, it holds that

$$u_i(M(I)|I) \geq u_i(M(I')|I).$$

Besides mechanisms that deterministically select a facility and its location, we will also study *randomized* mechanisms, which choose the outcome according to probability distributions. In particular, a randomized mechanism locates each facility $j \in \{1, 2\}$ at $y \in [0, 1]$ with some probability $p_j(y)$ such that $\sum_{j \in \{1,2\}} \int_0^1 p_j(y) dy = 1$. Denoting by $\mathbf{p} = (p_1, p_2)$ the probability distribution (for both facilities) used by the mechanism, the *expected utility* of every agent $i \in N$ is computed as

$$u_i(\mathbf{p}|I) = \sum_{j \in \{1,2\}} t_{ij} \cdot \int_0^1 (1 - |x_i - y|) \cdot p_j(y) dy.$$

A randomized mechanism is *strategyproof in expectation* if no agent can increase her *expected utility* by misreporting. Also, we say that a randomized mechanism is *universally strategyproof* if it is a probability distribution over deterministic strategyproof mechanisms. Clearly, a universally strategyproof mechanism is strategyproof in expectation, but the converse is *not* necessarily true.

We will also discuss about mechanisms that are resilient to misreports by coalitions of agents. In particular, a mechanism is *group-strategyproof* if no coalition of agents can simultaneously misreport such that the utility of *every* agent in the coalition strictly increases. A mechanism is *strongly group-strategyproof* if no coalition of agents can simultaneously misreport such that the utility of at least one agent in the coalition strictly increases and the utility of the other agents does not decrease.

We are interested in mechanisms that satisfy strategyproofness properties (like the ones discussed above) and at the same time achieve an as low as possible approximation ratio (that is, an approximation ratio as close as possible to 1). In our technical analysis in the upcoming sections, we will distinguish between the following settings:

- In the *general setting*, the agents can misreport both their positions and preferences.
- In the *known-preferences setting*, the preferences of the agents are assumed to be known and the agents can misreport only their positions.
- In the *known-positions setting*, the positions of the agents are assumed to be known and the agents can misreport only their preferences.

Observe that positive results (i.e., (group-)strategyproof mechanisms with proven approximation guarantees) for the general setting are also positive results for the known-preferences and known-positions settings. Moreover, negative results (i.e., lower bounds on the approximation of (group-)strategyproof mechanisms) for the restricted settings are also negative results for the general setting. Finally, results (positive or negative) for one of the two restricted settings do not imply anything for the other restricted setting.

3. General setting

We start the presentation of our technical results by focusing on the general setting; recall that in this setting the agents can misreport both their positions and their preferences. Due to the structure of the problem, which combines voting (based on the preferences of the agents) and facility location (based on the positions of the agents), it is natural to wonder whether simple adaptations of the median mechanism (which is known to be strategyproof and optimal for the original

Mechanism 1: MIDDLE.

for each facility $j \in \{1, 2\}$ do
 | Count the number n_j of agents that approve each facility j
 Locate the most-preferred facility at $1/2$



Fig. 1. The two instances used in the proof of Theorem 4.1, which differ only on the position of an agent with preferences (1, 0) marked in blue.

single-facility location problem) lead to good solutions. For example, we could define mechanisms that locate the majority-winner facility (breaking ties in a consistent way) at the median among the agents that approve it, or at the overall median agent. Unfortunately, it is not hard to observe that the first mechanism is *not* strategyproof, while the second one has an approximation ratio that is linear in the number of agents.

Luckily, there is an even simpler deterministic mechanism that is group-strategyproof and achieves an approximation of at most 2 in the general setting. We call this mechanism MIDDLE; see Mechanism 1. In the next section, we will further show that this mechanism is best possible among all deterministic strategyproof mechanisms in terms of approximation, even when the preferences of the agents are known.

Theorem 3.1. MIDDLE is group-strategyproof and has an approximation ratio of at most 2.

Proof. Consider any instance $I = (\mathbf{x}, \mathbf{t})$. To show that the mechanism is group-strategyproof, first observe that the positions of the agents are not taken into account when deciding which facility to locate and where. Hence, no agent has a reason to misreport her position. It remains to argue that there exists no group of agents who can all strictly increase their utility by misreporting their preferences. To this end, assume that facility $j \in \{1, 2\}$ is chosen to be placed at $1/2$. Observe that the utility of any agent that approves j is maximized subject to the constraint that the chosen facility is always placed at $1/2$. Hence, such agents would not have incentive to participate in a misreporting coalition. Moreover, the count n_j of facility j would only increase if any group of agents that truly disapprove facility j , misreport that they approve it. Hence, the outcome would not change in such a case, thus proving that is indeed group-strategyproof.

We now focus on the approximation ratio of the mechanism. Let w be the facility chosen by the mechanism, and let o be the optimal facility. We make the following simple observations:

- Since the facility is placed at $1/2$, every agent i that approves w has utility at least $1/2$.
- By the definition of the mechanism, since w is the majority winner, we have that $n_w \geq n_o$.
- Since the maximum utility of any agent is 1, we have that $W^*(I) \leq n_o$.

Putting all of these together, we have:

$$W(\text{MIDDLE}(I)|I) \geq \frac{1}{2}n_w \geq \frac{1}{2}n_o \geq \frac{1}{2}W^*(I),$$

and the bound on the approximation ratio follows. □

4. Known-preferences setting

Here, we focus on the known-preferences setting, where we assume that the agents can only strategize over their positions. Our first result is a lower bound of 2 on the approximation ratio of any strategyproof deterministic mechanism, thus proving that MIDDLE (the mechanism presented in the previous section) is best possible for the general and the known-preferences settings.

Theorem 4.1. In the known-preferences setting, there is no deterministic strategyproof mechanism with approximation ratio better than $2 - \delta$, for any $\delta > 0$.

Proof. Consider an arbitrary deterministic strategyproof mechanism and the following instance I depicted in Fig. 1. There are four agents, two with preferences (0, 1) and two with preferences (1, 0). One agent of each type is positioned at some

Mechanism 2: PROPORTIONAL.

for each facility $j \in \{1, 2\}$ do
 | Count the number n_j of agents that approve each facility j
 for each facility $j \in \{1, 2\}$ do
 | Choose facility j with probability $\frac{n_j}{n_1+n_2}$
 Locate the chosen facility at the median among the agents that approve it

$\varepsilon \in (0, 1/2)$ and the other is positioned at 1. Without loss of generality, we can assume that the mechanism chooses to locate facility 2.

Now consider a second instance I' that is obtained from I when only the agent i with preference $(1, 0)$ that is positioned at ε is moved to 1. Since the mechanism is strategyproof, it must choose to locate facility 2 in instance I' as well; otherwise, agent i would prefer to misreport her position in instance I as 1, thus leading to instance I' and the selection of facility 1, which would increase her utility from 0 to positive. However, the welfare from locating facility 2 in instance I' is at most $1 + \varepsilon$ (no matter where it is located), whereas the optimal welfare is equal to 2, achieved when facility 1 is located at 1. The bound on the approximation ratio follows by selecting ε to be arbitrarily small. \square

Next, we turn our attention to randomization and consider the class of *Random-Median* mechanisms.

Random-median mechanisms Every Random-Median mechanism operates by first randomly choosing one of the facilities based on the preferences of the agents (the mechanism chooses one of the facilities with probability $\alpha \in [0, 1]$ and the other with probability $1 - \alpha$), which is then located at the median among the agents that approve it. So, different choices of the probability distribution according to which the facility is chosen lead to different Random-Median mechanisms. It is not hard to observe that all such mechanisms are universally group-strategyproof.

Lemma 4.2. *Every Random-Median mechanism M is universally group-strategyproof.*

Proof. The lemma follows directly by the following two facts: (1) The choice of the facility to be located is made only based on the preferences of the agents, which are assumed to be known, and thus cannot be manipulated. (2) Given the facility, the location is chosen to be the position of the median agent among the ones that approve it, which is known to be a strongly group-strategyproof mechanism. \square

Probably the simplest Random-Median mechanism one may come up with to select every facility with probability proportional to the number of agents that approve it. As such, we call this mechanism PROPORTIONAL; see Mechanism 2. By exploiting the definition of this probability, we can show that PROPORTIONAL has an approximation of $(1 + \sqrt{3})/2 \approx 1.366$, thus significantly improving upon the bound of 2 achieved by deterministic mechanisms.

Theorem 4.3. *In the known-preferences setting, PROPORTIONAL is universally group-strategyproof and has an approximation ratio $(1 + \sqrt{3})/2 \approx 1.366$.*

Proof. Since the mechanism is Random-Median, it is universally group-strategyproof due to Lemma 4.2. To bound the approximation ratio, let W_j be the welfare of the agents that approve facility j when it is chosen (and located at the median of those agents). Without loss of generality, assume that $W_1 \geq W_2$. We also have that $W_1 \leq n_1$ since the maximum possible utility of any agent is 1. Furthermore, we have that $W_2 \geq n_2/2$. To see why this is the case, consider the agents that approve facility 2 in pairs, where one is on the left of the median (among those that approve facility 2) and the other is on the right of the median, and observe that the total utility of this pair of agents is at least 1. Since there are $n_2/2$ such pairs, the claim follows. The approximation ratio is

$$\begin{aligned} \rho(\text{PROPORTIONAL}) &= \frac{W_1}{\frac{n_1}{n_1+n_2} W_1 + \frac{n_2}{n_1+n_2} W_2} = \frac{1}{\frac{n_1}{n_1+n_2} + \frac{n_2}{n_1+n_2} \cdot \frac{W_2}{W_1}} \\ &\leq \frac{1}{\frac{n_1}{n_1+n_2} + \frac{n_2}{n_1+n_2} \cdot \frac{n_2/2}{n_1}} = \frac{2n_1^2 + 2n_1n_2}{2n_1^2 + n_2^2}. \end{aligned}$$

Now, let $y = n_1/n_2$ and observe that, since $n_1 \geq W_1 \geq W_2 \geq n_2/2$, it must be that $y \geq 1/2$. By dividing the last expression above by n_2^2 , we obtain:

$$\rho(\text{PROPORTIONAL}) \leq \frac{2(n_1/n_2)^2 + 2(n_1/n_2)}{2(n_1/n_2)^2 + 1} = \frac{2y^2 + 2y}{2y^2 + 1}.$$

Mechanism 3: MIRROR.

```

for each facility  $j \in \{1, 2\}$  do
  | Count the number  $n_j$  of agents that approve each facility  $j$ 
for each facility  $j \in \{1, 2\}$  do
  | if  $n_j \geq n_{3-j}$  then
  | |  $\alpha := \frac{3n_j - 2n_{3-j}}{4n_j - 2n_{3-j}}$ 
  | | Choose facility  $j$  with probability  $\alpha$  and facility  $3 - j$  with probability  $1 - \alpha$ 
  Locate the chosen facility at the median among the agents that approve it
    
```

Hence, in order to bound the approximation ratio of the mechanism it suffices to maximize the function $\frac{2y^2+2y}{2y^2+1}$ subject to the constraint $y \geq 1/2$. It is not hard to observe that the maximum value is $(1 + \sqrt{3})/2 \approx 1.366$ for $y^* = 1/(\sqrt{3} - 1)$, thus proving the upper bound on the approximation ratio. \square

We can further improve upon the bound of PROPORTIONAL, by optimizing over the probability that Random-Median mechanisms choose each facility. In particular, we define the slightly more involved MIRROR mechanism, which uses a probability distribution that is a piecewise function of the numbers of agents that approve the two facilities; see Mechanism 3. Following along the lines of the proof of Theorem 4.3, we can show that MIRROR has an approximation ratio of 4/3.

Theorem 4.4. *In the known-preferences setting, MIRROR is universally group-strategyproof and has an approximation ratio of 4/3.*

Proof. Since the mechanism is Random-Median, it is universally group-strategyproof due to Lemma 4.2. To bound the approximation ratio, let W_j be the welfare of the agents that approve facility j when it is chosen (and located at the median of those agents). Observe that for any facility $j \in \{1, 2\}$ it holds that $W_j \leq n_j$ since the maximum possible utility of any agent is 1, and $W_j \geq n_j/2$ following the same reasoning as in the proof of Theorem 4.3. Due to symmetry, we can assume without loss of generality that $n_1 \geq n_2$, in which case facility 1 is chosen with probability $\alpha := \frac{3n_1 - 2n_2}{4n_1 - 2n_2}$ and facility 2 is chosen with probability $1 - \alpha$. We distinguish between the following two cases:

- $W_1 \geq W_2$. Then, the approximation ratio is:

$$\begin{aligned} \rho(\text{MIRROR}) &= \frac{W_1}{\alpha \cdot W_1 + (1 - \alpha) \cdot W_2} = \frac{1}{\frac{3n_1 - 2n_2}{4n_1 - 2n_2} + \frac{n_1}{4n_1 - 2n_2} \cdot \frac{W_2}{W_1}} \\ &\leq \frac{1}{\frac{3n_1 - 2n_2}{4n_1 - 2n_2} + \frac{n_1}{4n_1 - 2n_2} \cdot \frac{n_2/2}{n_1}} = \frac{4}{3}. \end{aligned}$$

- $W_2 > W_1$. We have:

$$\begin{aligned} \rho(\text{MIRROR}) &= \frac{W_2}{\alpha \cdot W_1 + (1 - \alpha) \cdot W_2} = \frac{1}{\frac{3n_1 - 2n_2}{4n_1 - 2n_2} \cdot \frac{W_1}{W_2} + \frac{n_1}{4n_1 - 2n_2}} \\ &\leq \frac{1}{\frac{3n_1 - 2n_2}{4n_1 - 2n_2} \cdot \frac{n_1/2}{n_2} + \frac{n_1}{4n_1 - 2n_2}} = \frac{2n_2(4n_1 - 2n_2)}{3n_1^2}. \end{aligned}$$

It is not hard to observe that this last expression is maximized to 4/3 when $n_1 = n_2$.

Hence, in any case the approximation ratio of the mechanism is at most 4/3. \square

We conclude this section by showing that MIRROR is best possible among all Random-Median mechanisms in terms of approximation.

Theorem 4.5. *In the known-preferences setting, the approximation ratio of any Random-Median mechanism is at least $4/3 - \delta$, for any $\delta > 0$.*

Proof. Consider an arbitrary Random-Median mechanism and the following instance I (which was also used in the proof of Theorem 4.1). There are four agents, two with preferences $(0, 1)$ and two with preferences $(1, 0)$. One agent of each type is positioned at $\varepsilon \in (0, 1/2)$ while the other is positioned at $1 - \varepsilon$. Due to symmetry, we can assume that the agents located at ε are the medians for the two facilities. The mechanism randomly chooses one of the facilities. Without loss of generality, we can assume that it chooses facility 1 with some probability $p \leq 1/2$ and facility 2 with probability $1 - p$. Hence, facility 1 is located at ε with probability p .



Fig. 2. The two instances used in the proof of Theorem 5.1, which differ only on the preference of the agent positioned at 1/6 marked in blue.

Now consider the instance I' that is obtained from I by moving the agent i at $1 - \varepsilon$ with preference $(1, 0)$ to ε . We claim that the mechanism must choose facility 1 with probability $p' \leq p \leq 1/2$ in this new instance. Suppose that this is not the case and the mechanism chooses facility 1 with probability $p' > p$ in I' . Since the expected utility of agent i is $p \cdot 2\varepsilon$ in I , she would have incentive to misreport her position as ε so that facility 1 is chosen with probability p' and her expected utility is increased to $p' \cdot 2\varepsilon$. This contradicts the fact that the mechanism is strategyproof in expectation (since it is universally strongly group-strategyproof).

In instance I' , the maximum welfare we can achieve by placing facility 1 is 2 (when it is placed at ε), and by placing facility 2 is $1 + \varepsilon$ (no matter where it is located). Hence, since the maximum possible expected social welfare achieved by the mechanism is $2p' + (1 + \varepsilon)(1 - p') \leq \frac{3+\varepsilon}{2}$, the approximation ratio of the mechanism is at least $\frac{4}{3+\varepsilon} \geq 4/3 - \delta$, for any $\delta > \frac{4\varepsilon}{9+3\varepsilon}$. \square

5. Known-positions setting

We now turn our attention to the known-positions setting, in which the positions of the agents are fixed, and thus the agents can misreport only their preferences. Our first result is a lower bound of 13/11 on the approximation ratio of any deterministic strategyproof mechanism.

Theorem 5.1. *In the known-positions setting, there is no deterministic strategyproof mechanism with approximation ratio smaller than 13/11.*

Proof. Suppose towards a contradiction that there exists a deterministic strategyproof mechanism that has an approximation ratio strictly smaller than 13/11, and consider the following instance I depicted in Fig. 2. There is an agent with preferences $(0, 1)$ positioned at 0, an agent with preferences $(1, 1)$ positioned at 1/6, an agent with preferences $(1, 1)$ positioned at 5/6, and an agent with preferences $(1, 0)$ positioned at 1. Due to the symmetry of the instance, without loss of generality, we can assume that the mechanism chooses to place facility 1 at some location $y \in [0, 1]$.

If $y \leq 1/2$, the social welfare achieved by the mechanism is

$$\left(1 - \left|y - \frac{1}{6}\right|\right) + \left(1 - \left(\frac{5}{6} - y\right)\right) + \left(1 - (1 - y)\right) = \frac{7}{6} + 2y - \left|y - \frac{1}{6}\right| \leq \frac{11}{6}.$$

Since the optimal social welfare is 13/6 (achieved by placing either facility 1 at 5/6 or facility 2 at 1/6), the approximation ratio of the mechanism is then 13/11, contradicting the assumption that it is strictly smaller than 13/11. Consequently, it must be $y > 1/2$.

Now consider the instance I' that is obtained from I by changing the preference of the agent at 1/6 to $(0, 1)$. In I' , the maximum possible social welfare one can hope to achieve by placing facility 1 is 11/6 (when it is located anywhere in the interval $[5/6, 1]$ is 11/6) and by placing facility 2 is 13/6 (when it is located at 1/6). Hence, to have an approximation ratio strictly smaller than 13/11, the mechanism must choose to locate facility 2 in I' at some position $z \in [0, 1]$. Similarly to instance I , we can show that it must be $z < 1/2$ as otherwise the approximation ratio of the mechanism would be at least 13/11. Hence, the agent positioned at 1/6 with preferences $(1, 1)$ in instance I has incentive to misreport her preferences as $(0, 1)$ so that the facility is located closer to her, and thus increase her utility. However, this contradicts the assumption that the mechanism is strategyproof, and the theorem thus follows. \square

Unfortunately, we have been unable to design any deterministic mechanism with approximation ratio strictly smaller than the bound of 2 achieved by MIDDLE (see Section 3), and leave it as a challenging open question. Instead, we continue by considering randomized mechanisms. Observe that the instances from Fig. 2 show that there is no randomized strategyproof mechanism with approximation ratio 1 for the known-positions setting. This is because no matter how the mechanism locates the facilities on instance I , it has to locate facility 2 on Instance I' , thus the agent positioned at 1/6 has incentives to misreport her preferences.

Next, we analyze a particular version of the well-known RANDOM DICTATORSHIP (RD) mechanism. Our version picks each agent uniformly at random and locates her favorite facility at her position, breaking ties in favor of the optimal facility; see Mechanism 4. We will show that, in the known-positions setting, this mechanism is universally group-strategyproof and has an approximation ratio of 3/2.

Theorem 5.2. *In the known-positions setting, RD is universally group-strategyproof.*

Mechanism 4: RANDOMIZED DICTATORSHIP (RD).

```

for each agent  $i \in N$  do
  | Pick agent  $i$  with probability  $1/n$ 
if the chosen agent  $i^*$  approves a single facility  $j$  then
  | Locate facility  $j$  at  $x_{i^*}$ 
else
  | Locate the optimal facility at  $x_{i^*}$ 
    
```

Proof. Clearly, RD is a uniform probability distribution over the set of all possible deterministic mechanisms each of which treats a different agent as dictator, and locates one of the facilities that the dictator approves at her position. The lemma follows since the agents cannot misreport their positions and no agent has any incentive to lie about which facilities she approves when she is chosen as a dictator; also note that any coalition of agents that does not include the dictator cannot change the outcome of the mechanism. \square

To show that the approximation ratio of RD is $3/2$, we will exploit a series of structural properties of worst-case instances, in which the approximation ratio of the mechanism is maximized. In particular, we will show that there exists a worst-case instance that is characterized by the following three properties:

- There are no agents that approve both facilities (Lemma 5.3).
- Every agent that approves the non-optimal facility is positioned at 0 or 1 (Lemma 5.4).
- Every agent that approves the optimal facility is positioned at 0 or some median position $x \in [0, 1]$ or 1 (Lemma 5.5).

To show these properties, we will start with an arbitrary instance and gradually change the preferences and the positions of the agents in a specific order so that the aforementioned properties are satisfied. Every change we make leads to a transformed instance in which the approximation ratio of RD does not decrease. It then suffices to define a worst-case instance satisfying these properties and bound the approximation ratio of the mechanism for this instance. This will be a function of a handful of variables representing the number of agents that are positioned at $\{0, x, 1\}$ and approve one of the two facilities.

The first property is quite easy to observe, given the definition of the mechanism.

Lemma 5.3. *In a worst-case instance, all agents approve one facility.*

Proof. Consider an arbitrary instance in which there is a set of agents that approve both facilities. If the mechanism chooses any such agent, then the optimal facility is placed at the position of this agent. By transforming the preference of the agent so that she only approves the optimal facility, the optimal welfare remains unaffected, whereas the welfare achieved by the mechanism can only decrease. This is because the agent does not contribute to the welfare gained when the dictator is an agent who approves only the non-optimal facility. \square

Next, we show the second property.

Lemma 5.4. *In a worst-case instance, every agent that approves a non-optimal facility is positioned at 0 or 1.*

Proof. Consider an arbitrary instance in which each agent approves a single facility (from Lemma 5.3). Without loss of generality, assume that the optimal facility is 1, and let S be the set of agents that approve the non-optimal facility 2. We order the agents in S in terms of their positions, such that $x_1 \leq x_2 \leq \dots \leq x_{|S|}$. We partition S into two sets L and R , such that L consists of the first $\lceil |S|/2 \rceil$ agents in S and $R = S \setminus L$. Observe that if the number of agents in S is odd, the unique median agent in S is included in L , whereas if the number of agents in S is even, one of the two medians in S is included in L , while the other is included in R . To prove the lemma, we claim that moving the agents in L to 0 from left to right, and the agents in R to 1 from right to left, leads to a sequence of instances such that the approximation ratio of RD does not decrease between consecutive instances. See Fig. 3 for an example with $|S| = 3$. Due to symmetry, it suffices to prove this claim only for instances obtained by moving the agents in L .

Let m_L denote the median agent that is included in L and observe that facility 1 will remain optimal after moving any agent $\ell \in L \setminus \{m_L\}$. In particular, since the agents in $L \setminus \{\ell\}$ are not moved, their contribution to the maximum welfare achieved from facility 2, i.e. when it is placed at the position of m_L , remains the same. On the other hand, the contribution of agent ℓ (who is moved) cannot increase as her distance from the median agent(s) in S either remains the same (if $x_\ell = 0$) or increases (if $x_\ell > 0$). Once all the agents in $L \setminus \{m_L\}$ have been moved to 0, we can also show that moving agent m_L to 0 cannot increase the maximum possible welfare from facility 2. This follows directly by the discussion below, where we show that the expected welfare of RD does not increase each time we move an agent in L , and thus the approximation ratio does not decrease.

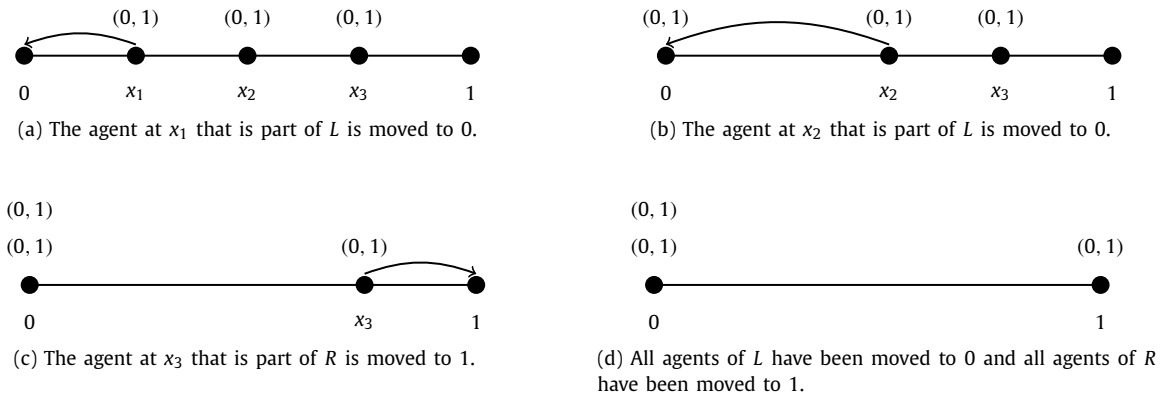


Fig. 3. An example of the transformations made in the proof of Lemma 5.4 with $|S| = 3$; agents with preferences different than $(0, 1)$ are not depicted. The first two agents (located at x_1 and x_2) are part of L and are moved to 0, while the last agent (located at x_3) is part of R and is moved to 1.

Suppose it is time to move agent $\ell \leq |L|$ to 0, that is, it holds that $x_i = 0$ for every $i < \ell$. The expected welfare of RD can be partitioned into the contribution W_S of agents in S (that approve only facility 2) and the contribution $W_{\bar{S}}$ of the remaining agents (that approve only facility 1). Clearly, changing the position of any agent in S does not affect $W_{\bar{S}}$. We can write W_S as

$$\begin{aligned}
 W_S &= \frac{1}{n} \sum_{i \in S} \sum_{j \in S} (1 - d(x_i, x_j)) \\
 &= \frac{1}{n} \left(\sum_{i \in S \setminus \{\ell\}} \sum_{j \in S \setminus \{\ell\}} (1 - d(x_i, x_j)) + 2 \sum_{i \in S \setminus \{\ell\}} (1 - d(x_i, x_\ell)) + 1 \right).
 \end{aligned}$$

After moving agent ℓ to $x'_\ell = 0$, the first double sum in the above expression and the constant 1 will remain unaffected. So, we will focus on the second sum, and show that it cannot increase. We define the set $S_{<\ell}$ of agents in S before ℓ and the set $S_{>\ell}$ of agents in S after ℓ . By the definition of the set L , it holds that $|S_{<\ell}| \leq |S_{>\ell}|$, with the equality holding only in the case where the number of agents in S is odd and $\ell = m_L$. Moreover, by the definition of agent ℓ , it holds that $x_i = 0$ for every $i \in S_{<\ell}$. We now have that

$$\begin{aligned}
 \sum_{i \in S \setminus \{\ell\}} (1 - d(x_i, x_\ell)) &= \sum_{i \in S_{<\ell}} (1 - d(x_i, x_\ell)) + \sum_{i \in S_{>\ell}} (1 - d(x_i, x_\ell)) \\
 &= \sum_{i \in S_{<\ell}} (1 - (x_\ell - x_i)) + \sum_{i \in S_{>\ell}} (1 - (x_i - x_\ell)) \\
 &= \sum_{i \in S_{<\ell}} (1 - (0 - x_i)) + \sum_{i \in S_{>\ell}} (1 - (x_i - 0)) + |S_{>\ell}| - |S_{<\ell}| \\
 &\geq \sum_{i \in S_{<\ell}} (1 - (x'_\ell - x_i)) + \sum_{i \in S_{>\ell}} (1 - (x_i - x'_\ell)) \\
 &= \sum_{i \in S \setminus \{\ell\}} (1 - d(x_i, x'_\ell)),
 \end{aligned}$$

as desired. This last sequence of equalities and inequalities also shows that moving m_L to 0 cannot increase the maximum possible social welfare from facility 2, as claimed above. \square

The third and last property follows by a proof similar to the one of the lemma above.

Lemma 5.5. *Let $x \in [0, 1]$ be the position of the median agent among those that approve the optimal facility in a worst-case instance. Then, every agent that approves the optimal facility is positioned at 0 or x or 1.*

Proof. Consider an arbitrary instance in which each agent approves a single facility (from Lemma 5.3), and all agents approving the non-optimal facility are positioned at 0 or 1 (from Lemma 5.4). Without loss of generality, assume that the optimal facility is 1, and let S be the set of agents that approve it. We will gradually transform this instance into one that satisfies the conditions of the lemma by appropriately moving the agents, such that each time we move an agent the approximation ratio of RD does not decrease. Let L and R be the sets of the non-median agents in S who are at the left

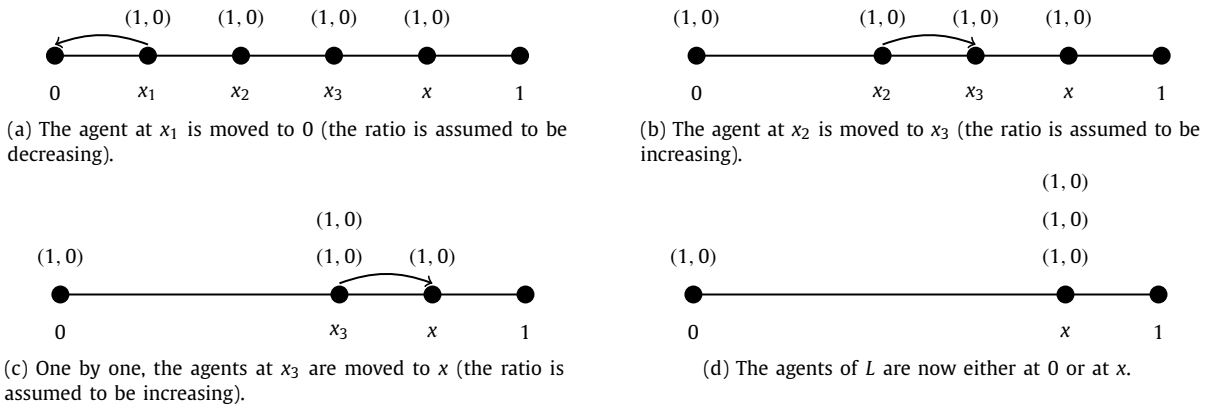


Fig. 4. An example of the transformations made in the proof of Lemma 5.5 with $|L| = 3$; agents of R or with preferences different than $(1, 0)$ are not depicted. When an agent is considered, she is either moved to 0 if the ratio is decreasing (a) or to the location on her right if the ratio is increasing (b, c).

and at the right of the median agent(s) in S , respectively. Due to symmetry, it suffices to focus only on L and show that the agents therein can be moved either to 0 or x . We will also assume that facility 1 remains optimal throughout the whole process; observe that this is without loss of generality since if facility 2 becomes the optimal one at some point of the process, then by replicating the procedure used in the proof of Lemma 5.4, we can move all agents in S to 0 or 1, thus obtaining the desired structure.

Suppose the agents in S are ordered in terms of their positions, such that $x_1 \leq x_2 \leq \dots \leq x_{|S|}$. Let ℓ be the left-most agent in L who is not positioned at 0 or x , that is, $x_\ell < x$ and $x_i = 0$ for every $i < \ell$. We can write the optimal social welfare as

$$W^* = \sum_{i \in S} (1 - d(x_i, x)) = \sum_{i \in S \setminus \{\ell\}} (1 - d(x_i, x)) + 1 - (x - x_\ell) = A + x_\ell.$$

Let W_S be the contribution of the agents in S to the expected welfare achieved by RD, and let $W_{\bar{S}}$ denote the contribution of the remaining agents who approve facility 2. We have

$$\begin{aligned} W_S &= \frac{1}{n} \sum_{i \in S} \sum_{j \in S} (1 - d(x_i, x_j)) \\ &= \frac{1}{n} \left(\sum_{i \in S \setminus \{\ell\}} \sum_{j \in S \setminus \{\ell\}} (1 - d(x_i, x_j)) + 2 \sum_{i \in S \setminus \{\ell\}} (1 - d(x_i, x_\ell)) + 1 \right). \end{aligned}$$

Let $S_{\leq \ell}$ be the set of agents in S different than ℓ with position at most x_ℓ , and let $S_{> \ell}$ be the set of agents in S with position strictly larger than ℓ . By the definition of ℓ (who belongs to L and is the left-most agent that is not positioned at 0 or x), it holds that $|S_{\leq \ell}| < |S_{> \ell}|$. Therefore,

$$\begin{aligned} W_S &= \frac{1}{n} \left(\sum_{i \in S \setminus \{\ell\}} \sum_{j \in S \setminus \{\ell\}} (1 - d(x_i, x_j)) + 2 \sum_{i \in S_{\leq \ell}} (1 - d(x_i, x_\ell)) + 2 \sum_{i \in S_{> \ell}} (1 - d(x_i, x_\ell)) + 1 \right) \\ &= \frac{1}{n} \left(\sum_{i \in S \setminus \{\ell\}} \sum_{j \in S \setminus \{\ell\}} (1 - d(x_i, x_j)) + 2 \sum_{i \in S_{\leq \ell}} (1 + x_i) + 2 \sum_{i \in S_{> \ell}} (1 - x_i) + 1 + 2x_\ell(|S_{> \ell}| - |S_{\leq \ell}|) \right) \\ &= B + \frac{2}{n} x_\ell (|S_{> \ell}| - |S_{\leq \ell}|). \end{aligned}$$

Hence, the approximation ratio of RD is

$$\rho(RD) = \frac{A + x_\ell}{W_{\bar{S}} + B + \frac{2}{n} x_\ell (|S_{> \ell}| - |S_{\leq \ell}|)}.$$

Now observe that, since $|S_{> \ell}| - |S_{\leq \ell}| > 0$, the ratio is a monotonic function of x_ℓ . If it is decreasing, we can move agent ℓ to 0, and if it is increasing, we can move agent ℓ to the position $y \in (x_\ell, x)$ of the first agent that lies strictly to the right of ℓ . Therefore, we can always move the left-most agent in L who is not positioned at 0 or x , to either 0 or strictly to the right. At some point, this procedure will lead all agents in L to be positioned either at 0 or x , and symmetrically, all agents in R to be positioned either at x or 1. See Fig. 4 for an example with $|L| = 3$. \square

We are now ready to prove the bound on the approximation ratio of RD.

Theorem 5.6. *The approximation ratio of RD is 3/2.*

Proof. Consider a worst-case instance I in which the optimal facility is 1. From Lemmas 5.3, 5.4 and 5.5, we can assume that there are $n = \alpha_0 + \alpha_x + \alpha_1 + \beta_0 + \beta_1$ agents in total, such that:

- α_0 agents approve only facility 1 and are positioned at 0;
- α_x agents approve only facility 1 and are positioned at some $x \in [0, 1]$;
- α_1 agents approve only facility 1 and are positioned at 1;
- β_0 agents approve only facility 2 and are positioned at 0;
- β_1 agents approve only facility 2 and are positioned at 1.

Since x is the position of the median agent among those that approve facility 1, we have that $\alpha_0 + \alpha_x \geq \alpha_1$ and $\alpha_x + \alpha_1 \geq \alpha_0$. We also have that $\beta_0 \geq \beta_1$. The optimal social welfare is achieved by placing facility 1 at position x , and is equal to

$$W^*(I) = \alpha_0 \cdot (1 - x) + \alpha_x + \alpha_1 \cdot x = \alpha_0 + \alpha_x + (\alpha_1 - \alpha_0)x.$$

Since facility 2 is not optimal, we have that $W^* \geq \beta_0 \geq \beta_1$. The expected welfare of RD is

$$\begin{aligned} W(\text{RD}(I)) &= \frac{1}{n} \left(\alpha_0(\alpha_0 + \alpha_x(1 - x)) + \alpha_x(\alpha_0(1 - x) + \alpha_x + \alpha_1x) + \alpha_1(\alpha_x x + \alpha_1) + \beta_0^2 + \beta_1^2 \right) \\ &= \frac{1}{n} \left((\alpha_0 + \alpha_x)^2 + \alpha_1^2 + 2\alpha_x(\alpha_1 - \alpha_0)x + \beta_0^2 + \beta_1^2 \right). \end{aligned}$$

Hence, by replacing $n = \alpha_0 + \alpha_x + \alpha_1 + \beta_0 + \beta_1$, the approximation ratio of RD can be written as the following function of x :

$$\rho(\text{RD}) = (\alpha_0 + \alpha_x + \alpha_1 + \beta_0 + \beta_1) \cdot \frac{\alpha_0 + \alpha_x + (\alpha_1 - \alpha_0)x}{(\alpha_0 + \alpha_x)^2 + \alpha_1^2 + \beta_0^2 + \beta_1^2 + 2\alpha_x(\alpha_1 - \alpha_0)x}.$$

It is now not hard to see that, since the factors of x in the numerator and the denominator are either both positive or negative, the ratio is a monotonic function in terms of x , and it thus attains its maximum value for $x = 0$ or $x = 1$.

Therefore, we can further simplify the worst-case instance I and assume that there are $n = \alpha_0 + \alpha_1 + \beta_0 + \beta_1$ agents in total, such that:

- α_0 agents approve only facility 1 and are positioned at 0;
- α_1 agents approve only facility 1 and are positioned at 1;
- β_0 agents approve only facility 2 and are positioned at 0;
- β_1 agents approve only facility 2 and are positioned at 1;

Without loss of generality, we assume that $\alpha_0 \geq \max\{\alpha_1, \beta_0, \beta_1\}$, hence the optimal welfare is $W^* = \alpha_0$, achieved by placing facility 1 at 0. The expected social welfare of RD is

$$W(\text{RD}(I)) = \frac{\alpha_0^2 + \alpha_1^2 + \beta_0^2 + \beta_1^2}{\alpha_0 + \alpha_1 + \beta_0 + \beta_1},$$

and thus the approximation ratio is

$$\rho(\text{RD}) = \frac{\alpha_0(\alpha_0 + \alpha_1 + \beta_0 + \beta_1)}{\alpha_0^2 + \alpha_1^2 + \beta_0^2 + \beta_1^2}.$$

By nullifying the partial derivatives of this function in terms of $\alpha_0, \alpha_1, \beta_0$ and β_1 we obtain a system of four equations, whose solution shows that the ratio is maximized to 3/2 when $\alpha_0 = 3\alpha_1$ and $\alpha_1 = \beta_0 = \beta_1$. □

6. Extension to choosing k out of m facilities

So far we have exclusively focused on the fundamental case where there are two facilities and one of them must be located. In this section, we define and make initial progress for the natural generalization when there are $m \geq 2$ different facilities from which we can choose to locate $k < m$. There are of course several different ways in which one can define the utility of an agent; here we will consider the case where the utility of an agent is the sum of utilities that she derives from the facilities that are located among the ones she approves. Formally, let $I = (\mathbf{x}, \mathbf{t}, m, k)$ be an instance with position profile \mathbf{x} , preference profile \mathbf{t} , and $m \geq 2$ facilities out of which we must choose and locate $k \geq 1$. Let S a subset of k facilities that are chosen to be located, and denote by $y_j \in [0, 1]$ the location of facility $j \in S$; let $\mathbf{y} = (y_j)_{j \in S}$. The utility of an agent is then defined as:

Mechanism 5: (k, m) -MIDDLE.

for each facility $j \in \{1, \dots, m\}$ do
 | Count the number n_j of agents that approve each facility j
 Locate the k most-preferred facilities at $1/2$, breaking ties arbitrarily

$$u_i(S, \mathbf{y}|I) = \sum_{j \in S} t_{ij} \cdot (1 - d(x_i, y_j))$$

For this case, a straightforward adaptation of MIDDLE satisfies strategyproofness constraints and has an approximation ratio of at most 2. In addition, by extending the proof of Theorem 4.1, we can show that for particular values of m and k this is the best-possible approximation among deterministic mechanisms, even when the preferences of the agents are assumed to be known. We provide the corresponding theorems below.

Theorem 6.1. *The (k, m) -MIDDLE mechanism is group-strategyproof when $k = 1$ and strategyproof when $k \geq 2$, and has an approximation ratio of at most 2 for any $k \geq 1$.*

Proof. For $k = 1$, the proof that the mechanism is group-strategyproof and has approximation ratio of at most 2 follows directly by Theorem 3.1, since the utility of every agent is defined by a single facility which is located at $1/2$. Now consider any instance $I = (\mathbf{x}, \mathbf{t}, m, k)$ with $m > k \geq 2$. To show that the mechanism is strategyproof, first observe that since it does not take into account the positions of the agents when deciding which subset of facilities to locate and where, the agents have no incentive to misreport their positions. In addition, no agent has any incentive to unilaterally misreport her preferences since any such misreport can only increase the count of facilities she does not approve, and thus her utility cannot increase.

To show that the mechanism has approximation ratio at most 2, let S be the subset of k facilities chosen by the mechanism, and also let O be the optimal subset of k facilities. We make the following simple observations.

- For every $j \in S$, it holds that $n_j = \sum_{i \in N} t_{ij}$.
- Since every $j \in S$ is placed at $1/2$, it holds that $1 - d(x_i, y_j) \geq 1/2$ for every agent i that approves j .
- By the definition of the mechanism, we have that $n_j \geq n_o$ for every $j \in S \setminus O$ and $o \in O \setminus S$. So, we can match each $j \in S \setminus O$ to a different facility $j(o) \in O \setminus S$ such that $n_j \geq n_{j(o)}$.
- We have that $W^*(I) \leq \sum_{o \in O} n_o$.

Putting everything together, we have that

$$\begin{aligned} W((k, m)\text{-MIDDLE}(I)|I) &= \sum_{i \in N} \sum_{j \in S} t_{ij} \cdot \left(1 - d\left(x_i, \frac{1}{2}\right)\right) \\ &\geq \frac{1}{2} \sum_{j \in S} \sum_{i \in N} t_{ij} \\ &= \frac{1}{2} \sum_{j \in S} n_j \\ &= \frac{1}{2} \sum_{j \in S \cap O} n_j + \frac{1}{2} \sum_{j \in S \setminus O} n_j \\ &\geq \frac{1}{2} \sum_{j \in S \cap O} n_j + \frac{1}{2} \sum_{j \in S \setminus O} n_{j(o)} \\ &= \frac{1}{2} \sum_{o \in O} n_o \geq \frac{1}{2} W^*(I), \end{aligned}$$

and the bound on the approximation ratio follows. \square

For completeness, we present a simple instance showing that (k, m) -MIDDLE is not group-strategyproof when $k \geq 2$.

Lemma 6.2. *(k, m) -MIDDLE is not group-strategyproof when $k \geq 2$.*

Proof. Consider an instance with $m \geq 4$ facilities, from which we must choose and locate $k \in \{2, \dots, m - 2\}$. There are $n = m$ agents such that every agent approves a different facility; specifically, agent $i \in \{1, \dots, m\}$ approves only facility i . Since $n_i = 1$ for every facility i , the (k, m) -Middle mechanism can choose any set S of k facilities. Then, every agent



Fig. 5. The two instances used in the proof of Lemma 7.1 to show that RD is not strategyproof in the general setting. In instance I , RD chooses facility 1 whenever an agent with preferences $(1, 1)$ is chosen as the dictator. This gives the agent with preferences $(1, 0)$ marked in blue incentive to misreport her position, thus leading to instance I' , where the tie is broken in favor of facility 2.

approving a facility $j \notin S$ obtains zero utility. Since $k \leq m - 2$, every pair of agents (i, j) approving facilities $i, j \notin S$ have incentive to form a coalition and change their preferences so that they both approve i and j . Such a group misreport would lead to $n_i = n_j = 2$, and result in both i and j being part of any set of k facilities chosen by the mechanism, thus showing that the agents have successfully manipulated the mechanism. \square

By appropriately extending the proof of Theorem 4.1, we can show that, there are values of m and k (in particular, $m \geq k(k + 1)$), such that the approximation ratio of any deterministic mechanism is at least 2, even when the preferences of the agents are known. As a result, (k, m) -MIDDLE is the best possible strategyproof deterministic mechanism in terms of approximation in the general and in the known-preferences settings.

Theorem 6.3. *The approximation ratio of every deterministic strategyproof mechanism that locates k out of m facilities, when $m \geq k(k + 1)$, is at least $2 - \delta$ for any $\delta > 0$, even when the preferences of the agents are known.*

Proof. Let I be an instance with $m = k(k + 1)$ facilities from which we want to choose and locate k , and $\lambda > 1$ some parameter. We partition the facilities into $k + 1$ sets S_1, \dots, S_{k+1} of size k each. For each $\ell \in \{1, \dots, k + 1\}$, there are 2λ agents that approve the facilities of set S_ℓ and nothing else, such that λ of them are positioned at $\varepsilon > 0$ while the remaining λ are positioned at 1.

Consider an arbitrary deterministic strategyproof mechanism. Since the mechanism can only choose k facilities when given I as input and the facilities are partitioned into $k + 1$ sets of size k , there exists an $\ell^* \in \{1, \dots, k + 1\}$ such that all of the k facilities of S_{ℓ^*} are not chosen by the mechanism. For each facility chosen by the mechanism, there are λ pairs of agents (one at ε and one at 1) each of which contributes total utility at most $1 + \varepsilon$ to the social welfare of the mechanism. Hence, the social welfare of the mechanism is at most $\lambda k(1 + \varepsilon)$.

One by one, we move the λ agents that approve the facilities of S_{ℓ^*} and are positioned at ε to 1. In the instance we obtain after moving one such agent, the mechanism must still not choose any of the facilities of S_{ℓ^*} as otherwise the agent that moved would have increased her utility from zero to positive, contradicting the strategyproofness of the mechanism. In the last instance, all the 2λ agents that approve the k facilities of S_{ℓ^*} are positioned at 1, whereas all other agents are as in I . Choosing the facilities of S_{ℓ^*} leads to a social welfare of $2\lambda k$, and thus the optimal social welfare is at least this much. Since the social welfare of the mechanism is still at most $\lambda k(1 + \varepsilon)$, by choosing ε to be arbitrarily small, we obtain a lower bound of at least $2 - \delta$ for any $\delta > 0$ on the approximation ratio of the mechanism. \square

7. Conclusion and open problems

There are several interesting problems that either remain open or arise from our work. The first natural direction is to tighten our results for deterministic and randomized mechanisms for the different settings we have considered. For deterministic mechanisms, while the general and the known-preferences settings are resolved by our work, it would still be quite interesting to close the gap between $13/11$ and 2 for the known-positions setting. For randomized mechanisms, the most intriguing open question is whether there exists such a mechanism with approximation ratio significantly smaller than 2 in the general setting. An obvious candidate is the RD mechanism that we presented in the context of the known-positions setting. Unfortunately, the particular variant of RD is no longer strategyproof when both the positions and the preferences of the agents are private, as shown by the following lemma.

Lemma 7.1. *RD is not strategyproof in the general setting.*

Proof. Let $n \geq 4$, and consider the following instance I depicted in the left part of Fig. 5. There is an agent with preferences $(1, 0)$ positioned at 0, $n - 2$ agents with preferences $(1, 1)$ positioned at $1/2$, and one agent with preferences $(0, 1)$ positioned at 1. Since the maximum possible social welfare by placing facility 1 is the same as the maximum possible social welfare by placing facility 2, we can without loss of generality assume that the optimal facility is 1, and thus when an agent with preferences $(1, 1)$ is randomly chosen, facility 1 is placed at her position; to avoid ties altogether, we could move the $n - 2$ agents that are now positioned at $1/2$ to $1/2 - \varepsilon$, for some arbitrarily small $\varepsilon > 0$. Observe that the expected utility of the agent i that has preferences $(0, 1)$ positioned at 1 is $\frac{1}{n}$, as she only gets some utility when she is chosen.

Now consider the instance I' that is obtained from I by moving agent i from 1 to $1/2$; this instance is depicted in the right part of Fig. 5. Since the optimal facility in I' is 2, when an agent with preferences $(1, 1)$ is randomly chosen as the dictator, facility 2 is placed at her position. Therefore, agent i has incentive to deviate to $1/2$ and change instance I to I' in order for her expected utility to become $(1 - \frac{1}{n}) \cdot \frac{1}{2} = \frac{n-1}{2n} > \frac{1}{n}$, thus proving that RD is not strategyproof when the agents can misreport their positions. \square

Intuitively, the reason that makes RD manipulable in the general setting is the tie-breaking rule that we use for the agents who approve both facilities; recall that such ties are broken in favor of the optimal facility. Breaking ties in this way is crucial for our characterization of the worst-case instances in Section 5, but can be exploited by agents who are allowed to misreport their positions. Designing a variant of RD that is strategyproof in the general setting is straightforward, for example, by breaking ties between the two facilities equiprobably. Importantly however, the aforementioned characterization no longer holds in that case, which makes the analysis much more challenging. As a matter of fact, we can show that for any variant that uses a fixed probabilistic tie-breaking rule, the approximation ratio is *strictly larger* than $3/2$! In the following theorem, we show this for the version of RD that breaks ties by locating facility 1 with probability p and facility 2 with probability $1 - p$; we refer to this mechanism as p -RD.

Lemma 7.2. p -RD has approximation ratio at least 1.518, for every fixed $p \in [0, 1]$.

Proof. Without loss of generality assume that $p \in [0, 1/2]$, and consider the following instance. There are 30 agents positioned at 0: 15 of them have preferences $(1, 1)$ and the remaining 15 have preferences $(1, 0)$. Furthermore, there are 20 agents positioned at 1: 10 of them have preferences $(1, 0)$ and the remaining 10 have preferences $(0, 1)$. Observe that the optimal solution locates facility 1 at 0 and achieves social welfare 30. The expected social welfare of p -RD is $\frac{(3+p) \cdot 15^2 + 2 \cdot 10^2}{50}$. Hence, the approximation ratio of p -RD is $\frac{1500}{(3+p) \cdot 15^2 + 200}$, which is larger than $\frac{120}{79} > 1.518$ for every $p \in [0, 1/2]$. \square

Finding the exact approximation ratio of p -RD is an intriguing open question. Perhaps more interestingly, one can define yet another strategyproof variant of RD, whose approximation ratio is not ruled out by Lemma 7.2. For example, we can count how many agents approve each facility and break ties *proportionally* to those numbers. It is quite easy to observe that the RD mechanism using the proportional tie-breaking rule is strategyproof for the general setting, and is a promising candidate for achieving a better approximation ratio. To this end, we state the following conjecture.

Conjecture 7.3. For the general setting, the RD mechanism with the proportional tie-breaking rule has approximation ratio $3/2$.

Besides strengthening our results, there are several meaningful extensions of our model that could be the subject of future work. For the k out of m facilities setting, while we have made an important first step, there is still significant work to be done, particularly in the known-positions setting. One could also consider several different variants of our basis model. For example, the agents may have fractional preferences rather than approval preferences (that is, each agent may assign weights in $[0, 1]$ to the facilities, instead of weights in $\{0, 1\}$). Other possible variants include settings in which some facilities are *obnoxious* (Mei et al., 2018; Feigenbaum and Sethuraman, 2015), meaning that agents would like to be far from them if they are built, and discrete settings in which the facilities can only be built at predefined locations on the line (e.g., see (Dokow et al., 2012; Feldman et al., 2016; Serafino and Ventre, 2015, 2016)). Finally, an interesting generalization of our problem is when every facility comes at a different cost, and the objective is to maximize the social welfare by choosing and locating k facilities under the constraint that their accumulated costs is below a predefined budget. This latter setting is directly motivated by *participatory budgeting*, which has recently drawn the attention of the computational social choice community (Benade et al., 2021; Aziz and Shah, 2020).

Declaration of competing interest

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

Data availability

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

References

- Alon, Noga, Feldman, Michal, Procaccia, Ariel D., Tennenholtz, Moshe, 2010. Strategyproof approximation of the minimax on networks. *Math. Oper. Res.* 35 (3), 513–526.
- Anastasiadis, Eleftherios, Deligkas, Argyrios, 2018. Heterogeneous facility location games. In: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 623–631.

- Anshelevich, Elliot, Filos-Ratsikas, Aris, Shah, Nisarg, Voudouris, Alexandros A., 2021a. 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.
- Anshelevich, Elliot, Filos-Ratsikas, Aris, Voudouris, Alexandros A., 2021b. The distortion of distributed metric social choice. In: Proceedings of the 17th International Conference on Web and Internet Economics (WINE), pp. 467–485.
- Aziz, Haris, Shah, Nisarg, 2020. Participatory budgeting: models and approaches. CoRR. arXiv:2003.00606 [abs].
- Benade, Gerdus, Nath, Swaprava, Procaccia, Ariel D., Shah, Nisarg, 2021. Preference elicitation for participatory budgeting. *Manag. Sci.* 67, 2813–2827.
- Cai, Qingpeng, Filos-Ratsikas, Aris, Tang, Pingzhong, 2016. Facility location with minimax envy. In: Kambhampati, Subbarao (Ed.), Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), pp. 137–143.
- Caragiannis, Ioannis, Nath, Swaprava, Procaccia, Ariel D., Shah, Nisarg, 2017. Subset selection via implicit utilitarian voting. *J. Artif. Intell. Res.* 58, 123–152.
- Caragiannis, Ioannis, Shah, Nisarg, Voudouris, Alexandros A., 2022. The metric distortion of multiwinner voting. In: Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI).
- Chan, Hau, Filos-Ratsikas, Aris, Li, Bo, Li, Minming, Wang, Chenhao, 2021. Mechanism design for facility location problem: a survey. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), pp. 1–17.
- Chen, Zhihuai, Fong, Ken C.K., Li, Minming, Wang, Kai, Yuan, Hongning, Zhang, Yong, 2020. Facility location games with optional preference. *Theor. Comput. Sci.* 847, 185–197.
- Cheng, Yukun, Yu, Wei, Zhang, Guochuan, 2011. Mechanisms for obnoxious facility game on a path. In: *Combinatorial Optimization and Applications*. Springer, pp. 262–271.
- Cheng, Yukun, Yu, Wei, Zhang, Guochuan, 2013. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theor. Comput. Sci.* 497, 154–163.
- Dokow, Elad, Feldman, Michal, Meir, Reshef, Nehama, Ilan, 2012. Mechanism design on discrete lines and cycles. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 423–440.
- Elkind, Edith, Li, Minming, Zhou, Houyu, 2022. Facility location with approval preferences: strategyproofness and fairness. In: Proceedings of the 2022 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).
- Escoffier, Bruno, Gourves, Laurent, Thang, Nguyen Kim, Pascual, Fanny, Spanjaard, Olivier, 2011. Strategy-proof mechanisms for facility location games with many facilities. In: *Algorithmic Decision Theory*. Springer, pp. 67–81.
- Faliszewski, Piotr, Skowron, Piotr, Slinko, Arkadii, Talmon, Nimrod, 2017. Multiwinner voting: a new challenge for social choice theory. *Trends in Computational Social Choice* 74 (2017), 27–47.
- Feigenbaum, Itai, Sethuraman, Jay, 2015. Strategyproof mechanisms for one-dimensional hybrid and obnoxious facility location models. In: Proceedings of the 2015 AAAI Workshop on Incentive and Trust in E-Communities, vol. WS-15-08.
- Feigenbaum, Itai, Sethuraman, Jay, Ye, Chun, 2013. Approximately optimal mechanisms for strategyproof facility location: minimizing L_p norm of costs. CoRR. arXiv:1305.2446 [abs].
- Feldman, Michal, Wilf, Yoav, 2013. Strategyproof facility location and the least squares objective. In: Proceedings of the 14th ACM Conference on Electronic Commerce (EC), pp. 873–890.
- Feldman, Michal, Fiat, Amos, Golomb, Iddan, 2016. On voting and facility location. In: Proceedings 2016 ACM Conference on Economics and Computation (EC), pp. 269–286.
- Filos-Ratsikas, Aris, Voudouris, Alexandros A., 2021. Approximate mechanism design for distributed facility location. In: Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT), pp. 49–63.
- Filos-Ratsikas, Aris, Li, Minming, Zhang, Jie, Zhang, Qiang, 2015. Facility location with double-peaked preferences. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), pp. 893–899.
- Fong, Chi Kit Ken, Li, Minming, Lu, Pinyan, Todo, Taiki, Yokoo, Makoto, 2018. Facility location games with fractional preferences. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), pp. 1039–1046.
- Fotakis, Dimitris, Tzamos, Christos, 2013. Winner-imposing strategyproof mechanisms for multiple facility location games. *Theor. Comput. Sci.* 472, 90–103.
- Fotakis, Dimitris, Tzamos, Christos, 2016. Strategyproof facility location for concave cost functions. *Algorithmica* 76 (1), 143–167.
- Golomb, Iddan, Tzamos, Christos, 2017. Truthful facility location with additive errors. CoRR. arXiv:1701.00529 [abs].
- Kanellopoulos, Panagiotis, Voudouris, Alexandros A., Zhang, Rongsen, 2021. On discrete truthful heterogeneous two-facility location. CoRR. arXiv:2109.04234 [abs].
- Kyropoulou, Maria, Ventre, Carmine, Zhang, Xiaomeng, 2019. Mechanism design for constrained heterogeneous facility location. In: Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT), pp. 63–76.
- Li, Minming, Lu, Pinyan, Yao, Yuhao, Zhang, Jialin, 2020a. 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.
- Li, Minming, Wang, Chenhao, Zhang, Mengqi, 2020b. Budgeted facility location games with strategic facilities. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), pp. 400–406.
- Lu, Pinyan, Wang, Yajun, Zhou, Yuan, 2009. Tighter bounds for facility games. In: Proceedings of the 5th International Workshop on Internet and Network Economics (WINE), pp. 137–148.
- Lu, Pinyan, Sun, Xiaorui, Wang, Yajun, Zhu, Zeyuan Allen, 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In: Proceedings of the 11th ACM Conference on Electronic Commerce (EC), pp. 315–324.
- Mei, Lili, Ye, Deshi, Zhang, Yong, 2018. Approximation strategy-proof mechanisms for obnoxious facility location on a line. *J. Comb. Optim.* 36 (2), 549–571.
- Procaccia, Ariel D., Tennenholtz, Moshe, 2013. Approximate mechanism design without money. *ACM Trans. Econ. Comput.* 1 (4), 18:1–18:26.
- Schummer, James, Vohra, Rakesh V., 2002. Strategy-proof location on a network. *J. Econ. Theory* 104 (2), 405–428.
- Serafino, Paolo, Ventre, Carmine, 2015. Truthful mechanisms without money for non-utilitarian heterogeneous facility location. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), pp. 1029–1035.
- Serafino, Paolo, Ventre, Carmine, 2016. Heterogeneous facility location without money. *Theor. Comput. Sci.* 636, 27–46.
- Sui, Xin, Boutilier, Craig, 2015. Approximately strategy-proof mechanisms for (constrained) facility location. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 605–613.
- Sui, Xin, Boutilier, Craig, Sandholm, Tuomas, 2013. Analysis and optimization of multi-dimensional percentile mechanisms. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), pp. 367–374.
- Tang, Pingzhong, Yu, Dingli, Zhao, Shengyu, 2020. Characterization of group-strategyproof mechanisms for facility location in strictly convex space. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC).
- Xu, Xinping, Li, Bo, Li, Minming, Duan, Lingjie, 2021. Two-facility location games with minimum distance requirement. *J. Artif. Intell. Res.* 70, 719–756.
- Zhang, Qiang, Li, Minming, 2014. Strategyproof mechanism design for facility location games with weighted agents on a line. *J. Comb. Optim.* 28 (4), 756–773.