Constant-Approximate and Constant-Strategyproof Two-Facility Location

Elijah Journey Fullerton[1][0009-0005-7254-5952], Zeyuan Hu[2][0000-0003-3036-2777], and C. Gregory Plaxton[2] ✉[0000-0001-7084-1612]

[1] Princeton University (ef0952@princeton.edu)
[2] University of Texas at Austin
(zeyuanhu@cs.utexas.edu, plaxton@cs.utexas.edu)

Abstract. We study deterministic mechanisms for the two-facility location problem. Given the reported locations of n agents on the real line, such a mechanism specifies where to build the two facilities. The single-facility variant of this problem admits a simple strategyproof mechanism that minimizes social cost. For two facilities, however, it is known that any strategyproof mechanism is $\Omega(n)$-approximate. We seek to circumvent this strong lower bound by relaxing the problem requirements. Following other work in the facility location literature, we consider a relaxed form of strategyproofness in which no agent can lie and improve their outcome by more than a constant factor. Because the aforementioned $\Omega(n)$ lower bound generalizes easily to constant-strategyproof mechanisms, we introduce a second relaxation: Allowing the facilities (but not the agents) to be located in the plane. Our first main result is a natural mechanism for this relaxation that is constant-approximate and constant-strategyproof. A characteristic of this mechanism is that a small change in the input profile can produce a large change in the solution. Motivated by this observation, and also by results in the facility reallocation literature, our second main result is a constant-approximate, constant-strategyproof, and Lipschitz continuous mechanism.

Keywords: Facility Location Problem · Mechanism Design

1 Introduction

Facility location is a canonical problem in algorithmic game theory and mechanism design. The typical problem formulation is that given a set of agent locations, a facility location mechanism determines a set of locations for the facilities. Each agent then has a service cost given by the distance between their location and their nearest facility. The social cost is defined as the sum of the agent service costs. It is desirable for a mechanism to be efficient, that is, to minimize the social cost (under the assumption of truthful reporting). To encourage truthful reporting, it is desirable for a mechanism to be strategyproof, which means that an agent cannot decrease their service cost by misreporting.

In this paper, we focus on deterministic mechanisms for the well-studied setting in which the agents are located on the real line. While our focus is on the design of mechanisms for locating two facilities, it is useful to briefly comment on the case of a single facility. As discussed by Procaccia and Tennenholtz [43], it is easy to design a single-facility location mechanism that is efficient and strategyproof; the idea is to place the facility at the median agent location when the number of agents is odd and at either the left or right middle agent location when the number of agents is even.[1]

When the number of facilities is expanded to two, strategyproofness starts to restrict the efficiency that a mechanism can provide. Lu et al. [36] show that any strategyproof mechanism has $\Omega(n)$ approximation ratio of the social cost.[2] In this paper, we seek to circumvent the $\Omega(n)$ lower bound by relaxing the problem requirements. The first relaxation is to target an approximate form of strategyproofness. In particular, we seek a mechanism that is constant-strategyproof [41] instead of strategyproof. A constant-strategyproof mechanism ensures that no agent can reduce their service distance by more than a multiplicative constant factor by misreporting. It is straightforward to generalize the $\Omega(n)$ result of Lu et al. [36] to show that the same $\Omega(n)$ lower bound holds for constant-strategyproof mechanisms. Accordingly, we introduce a second relaxation: We allow the facilities to be located in the plane. This is a departure from the prior work, where the agents and the facilities lie in the same metric space [4,17,24,32,42,47,48]. Intuitively, this relaxation allows us to more easily attain constant-strategyproofness at the expense of increasing our approximation ratio of optimal social cost.

1.1 Our Contributions

These relaxations enable our two main results. First, we present mechanism M_3, a natural mechanism that is constant-approximate (i.e., attains a constant approximation ratio of optimal social cost) and constant-strategyproof. Informally, the mechanism starts with an optimal solution on the line and then moves each facility away from the line by a distance inversely proportional to the number of agents it serves. Intuitively, lifting facilities off the real line weakens the ability of an agent to significantly reduce their own service cost through misreporting, while still keeping the social cost within a constant factor of optimal. This technique is reminiscent of 'money burning' ideas that limit manipulation at the cost of some efficiency to preserve truthful behavior. A characteristic of mechanism M_3 is that small changes in agent locations can cause unbounded changes in the locations of facilities. To address this characteristic, we present mechanism M_4,

[1] More general results of Moulin [40] establish that for agent locations on the real number line and single-peaked preferences, these are the only mechanisms for single-facility location that are efficient and strategyproof.

[2] The simple mechanism that locates the facilities at the leftmost and the rightmost agents achieves $O(n)$ approximation of the social cost, which indicates that the $\Omega(n)$ bound is asymptotically tight [43]. In fact, Fotakis and Tzamos [21] show that the foregoing mechanism achieves the best possible approximation ratio for the problem and is the only deterministic anonymous strategyproof mechanism with approximation ratio bounded by a function of n.

which is constant-approximate, constant-strategyproof, and constant-Lipschitz (i.e., Lipschitz continuous with a constant Lipschitz factor).

1.2 Related Work

Facility location is a central problem in mechanism design without money. There has been extensive research that studies facility location under various settings such as fully general metric spaces [18,38], strictly convex spaces [49], graph metrics [1,12,13,19,45], and probabilistic metrics [3,33]. We highlight results that are directly relevant to ours and defer the rest to a survey of the field [8].

Facility Location in Higher Dimensions One relevant line of investigation is the study of facility location in higher dimensions [4,17,24,32,42,47,48]. For example, Barberà et al. [4] generalize Moulin's single-facility result to higher dimension by selecting the facility's coordinate within each dimension independently in a “median-like" fashion. Sui et al. [48] generalize median mechanisms from [4,6,40] into percentile mechanisms by determining facility locations using percentiles of ordered agent location projections in each dimension. While these mechanisms are strategyproof, they show that for any dimension greater than one, no percentile mechanism attains an approximation ratio with respect to social cost bounded by a function of n.

Strategyproofness Relaxations In recent years, there has been growing interest in exploring approximate notions of strategyproofness in order to enable trade-offs with other desired properties [5,14,16,34,37,41,44,46,47]. One relaxation is to allow an agent to obtain a multiplicative constant factor gain through misreporting [22,29,41]. This relaxation closely parallels the incentive ratio in market design [9]. Oomine et al. [41] studied the obnoxious facility game [11] under this relaxation and exhibited a trade-off between strategyproofness and the approximation ratio. Specifically, with the multiplicative constant factor of λ, they obtained the approximation ratio of $1 + \frac{2}{\lambda}$. The multiplicative factor can be viewed as providing a bound on the maximum possible gain an agent can obtain by misreporting, which is broadly aligned with the goal of designing mechanisms with limited manipulability [10,16,28,44,46].

Facility Reallocation The setting of facility reallocation motivates the design of Lipschitz continuous mechanism [20,31]. The facility reallocation problem considers a sequence of agent preference profiles over time with the optimization goal being to minimize both the social cost at each time step and the facility reallocation distances across time steps. Lipschitz continuous mechanisms allow us to bound the facility reallocation distances as a function of the preference profile distances without any information about prior or future time steps.

2 Preliminaries

For any nonnegative integer k, we write $[k]$ as a shorthand for $\{1, \ldots, k\}$.

In this work, any instance of the facility location problem involves a set of two or more agents indexed starting at 1. A profile π is a vector of real values specifying the preferences of the agents on the real line in nondecreasing order. The nondecreasing requirement reflects the fact that all of the mechanisms we consider are anonymous. For any positive integer n, we define an n-profile as a profile for an n-agent instance. For any n-profile π, we define $\mu(\pi)$ as $\frac{1}{n}\sum_{i\in[n]} \pi_i$. A profile π is *trivial* if all of its components are equal; otherwise, it is *nontrivial.*

For any n-profile π and any i in $[n]$, we define $\text{interval}(\pi, i)$ as $(-\infty, \pi_2]$ if $i = 1$, as $[\pi_{n-1}, \infty)$ if $i = n$, and as $[\pi_{i-1}, \pi_{i+1}]$ otherwise. Observe that a real number x belongs to $\text{interval}(\pi, i)$ if and only if $(\pi_1, \ldots, \pi_{i-1}, x, \pi_{i+1}, \ldots, \pi_n)$ is an n-profile. For any n-profile π, we define $\Gamma(\pi)$ as the set of all n-profiles π^* such that $|\{i \in [n] \mid \pi_i \neq \pi^*_i\}| \leq 1$. For any n-profile π, any i in $[n]$, and any real number x, we define $\text{subst}(\pi, i, x)$ as the n-profile obtained from π by replacing π_i with x and rearranging the resulting components in nondecreasing order.

In this paper, we use the term "point" to refer to an element of $\mathbb{R}^2$, and we write p_1 (resp., p_2) to denote the x-coordinate (resp., y-coordinate) of a given point p. We say that a point p is *1-D* if $p_2 = 0$. For any real number x and any point p, we define $d(x, p)$ as $\|(x, 0) - p\|_1 = |x - p_1| + |p_2|$.[3]

A solution τ for a given profile π is an ordered pair (τ_1, τ_2) of points specifying the locations of the two facilities. We require that $\tau_1 \leq \tau_2$, where the comparison is performed lexicographically. We sometimes refer to τ_1 (resp., τ_2) as the left (resp., right) facility. We write $\tau_{i,j}$, where i and j belong to $\{1, 2\}$, to refer to component j of τ_i. We say that a solution τ is *1-D* if τ_1 and τ_2 are each 1-D.[4] For any real number x and any solution τ, we define $d(x, \tau)$ as $\min(d(x, \tau_1), d(x, \tau_2))$. In other words, $d(x, \tau)$ is equal to the distance between x and its nearest facility, which corresponds to the *individual cost function* of an agent with preference x. For any two solutions τ and τ^*, we define $\|\tau - \tau^*\|_1$ as $\sum_{\ell\in\{1,2\}} \|\tau_\ell - \tau^*_\ell\|_1$.

For any profile π and any solution τ, we define $C(\pi, \tau)$ as $\sum_{i\in[n]} d(\pi_i, \tau)$. In other words, we define the *social cost* of a profile π and solution τ as the sum of individual agent costs $d(\pi_i, \tau)$.

Lemma 2.1. *Let π and π^* be n-profiles and let τ be a solution. Then $|C(\pi^*, \tau) - C(\pi, \tau)|$ is at most $\|\pi^* - \pi\|_1$.*

Proof. For any i in $[n]$, $|d(\pi^*_i, \tau) - d(\pi_i, \tau)|$ is at most $|\pi^*_i - \pi_i|$. □

In our analysis, we sometimes want to identify the set of agents served by each facility under a given solution. For any profile π and any solution τ, we define $S_1(\pi, \tau)$ as the set of all indices i in $[n]$ such that $d(\pi_i, \tau_1) \leq d(\pi_i, \tau_2)$, and

[3] Even though we focus on the Manhattan distance in this paper, all of our results also hold for Euclidean distance up to constant factors.

[4] We adopt the convention that the symbol σ (rather than τ) is used to refer to a solution that is guaranteed to be 1-D.

we define $S_2(\pi, \tau)$ as $[n] \setminus S_1(\pi, \tau)$. Note that if an agent is equidistant from the two facilities, we consider it to be served by the left facility.

We define the left-median of $s \geq 1$ real numbers $x_1 \leq \cdots \leq x_s$ as $x_{\lceil s/2 \rceil}$. That is, if s is odd, then the left-median is the median, and if s is even, then the left-median is the lower-indexed of the two middle values. The following fact is attributable to Procaccia and Tennenholtz [43].

Fact 1 *For any profile π, the lexicographically first minimum-cost single-facility solution locates the facility at the left-median of π.*

For any n-profile π and any i in $[n-1]$, we define $\text{cand}(\pi, i)$ as the 1-D solution that locates the left facility at the left-median of the leftmost i agents (i.e., at $\pi_{\lceil i/2 \rceil}$), and the right facility at the left-median of the rightmost $n - i$ agents (i.e., at $\pi_{\lceil (n+i)/2 \rceil}$). We also define $\text{cand}(\pi, n)$ as the 1-D solution $((\pi_{\lceil n/2 \rceil}, 0), (\pi_{\lceil n/2 \rceil}, 0))$. Solution $\text{cand}(\pi, n)$ locates both facilities at the left-median of the entire set of n agents; note that the left facility is viewed as serving the entire set of n agents under this solution.

For any n-profile π, we define $\text{active}(\pi)$ as the set of all i in $[n]$ such that under solution $\text{cand}(\pi, i)$, the left facility serves exactly i agents. For any n-profile π, we define $\text{canonical}(\pi)$ as the lexicographically first minimum-cost solution for π if π is nontrivial, and as $\text{cand}(\pi, n)$ otherwise.

Lemma 2.2. *Let π be an n-profile and let σ denote* $\text{canonical}(\pi)$. *Then there is an index i in* $\text{active}(\pi)$ *such that* $\sigma = \text{cand}(\pi, i)$.

Proof. If π is trivial, it is clear that n belongs to $\text{active}(\pi)$ and $\sigma = \text{cand}(\pi, n)$. For the remainder of the proof, assume that π is nontrivial. Thus σ is the lexicographically first minimum-cost solution for π. Let i denote $|S_1(\pi, \sigma)|$. Since π is nontrivial and σ is minimum-cost, we find that $\sigma_{1,1} < \sigma_{2,1}$ and i belongs to $[n-1]$. Since σ is the lexicographically first minimum-cost solution for π, we deduce from Fact 1 that $\sigma_{1,1}$ is the left-median of $\pi_1, \ldots, \pi_i$ and $\sigma_{2,1}$ is the left-median of $\pi_{i+1}, \ldots, \pi_n$. Thus i belongs to $\text{active}(\pi)$ and $\sigma = \text{cand}(\pi, i)$. □

Using Lemma 2.2, we deduce that for any n-profile π, there is exactly one i in $[n]$ such that $\text{canonical}(\pi) = \text{cand}(\pi, i)$; we define this i as $\text{index}(\pi)$.

For any profile π, we define the *optimal social cost* $C(\pi)$ as $C(\pi, \sigma)$ and $S_\ell(\pi)$ as $S_\ell(\pi, \sigma)$ for all ℓ in $\{1, 2\}$, where σ denotes $\text{canonical}(\pi)$.

Lemma 2.3. *Let π and π^* be n-profiles. Then $|C(\pi^*) - C(\pi)| \leq \|\pi^* - \pi\|_1$.*

Proof. By symmetry, it is sufficient to establish that $C(\pi^*) - C(\pi) \leq \|\pi^* - \pi\|_1$. Let σ (resp., σ^*) denote the canonical solution for profile π (resp., π^*). We have $C(\pi^*) \leq C(\pi^*, \sigma) \leq C(\pi) + \|\pi^* - \pi\|_1$, where the second inequality follows from Lemma 2.1. □

A mechanism M maps any given profile to a solution. We say that a mechanism is 1-D if it only produces 1-D solutions.

For any mechanism M and any function $\alpha : \mathbb{N} \to \mathbb{R}^+$, we define the following terms:

1. M is $\alpha(n)$-*approximate* if $C(\pi, M(\pi)) \leq \alpha(n)C(\pi)$ for all n-profiles π;
2. M is $\alpha(n)$-*strategyproof* if $d(\pi_i, M(\pi)) \leq \alpha(n)d(\pi_i, M(\text{subst}(\pi, i, x)))$ for all n-profiles π, all i in $[n]$, and all real numbers x;
3. M is $\alpha(n)$-*Lipschitz* if $\|M(\pi^*) - M(\pi)\|_1 \leq \alpha(n)\|\pi^* - \pi\|_1$ for all n-profiles π and π^*.

As discussed in Section 1, our interest is in achieving constant-factor bounds (i.e., independent of n). We say a mechanism is constant-approximate (resp. constant-strategyproof, constant-Lipschitz) if the factor $\alpha(n)$ is $O(1)$. Note that when $\alpha(n) = 1$, $\alpha(n)$-strategyproof is strategyproof and when $\alpha(n)$ is a constant, $\alpha(n)$-strategyproof corresponds to constant-strategyproof defined in [41].

3 Facility Location in $\mathbb{R}$

3.1 Mechanism M_1

We define our first mechanism M_1 as follows: For any profile π, $M_1(\pi)$ is equal to canonical(π).

By construction, it is clear to see that this mechanism attains the optimal social cost and therefore is constant-approximate with an approximation ratio of 1. However, the following counter-example shows that this mechanism is neither constant-Lipschitz nor constant-strategyproof.

Lemma 3.1. *Mechanism M_1 is not constant-strategyproof and is not constant-Lipschitz.*

Proof. Let λ be an arbitrarily large positive real number, let z and ε be positive real numbers such that $z > 2\varepsilon\lambda$, let π denote the profile $(-z, 0, 0, z+\varepsilon)$, let σ denote the solution $M_1(\pi) = ((0,0),(z+\varepsilon, 0))$, let π^* denote the profile $(-z - 2\varepsilon, 0, 0, z+\varepsilon)$, and let σ^* denote the solution $M_1(\pi^*) = ((-z-2\varepsilon, 0), (0,0))$. Observe that only agent 1 changed their report (from π to π^*) and only changed it by 2ε. We have

$$d(\pi_1, \sigma) = z \geq 2\varepsilon\lambda = \lambda d(\pi_1^*, \sigma^*).$$

Hence M_1 is not constant-strategyproof. Furthermore,

$$\frac{\|\sigma - \sigma^*\|_1}{\|\pi - \pi^*\|_1} = \frac{2z + 3\varepsilon}{2\varepsilon},$$

which goes to infinity as $z \to \infty$. Hence M_1 is not constant-Lipschitz. □

The linear bound on the approximation ratio of optimal social cost for deterministic, strategyproof, mechanisms provided by [36] can trivially be extended to constant-strategyproof mechanisms to show that, in fact, no 1-D, constant-strategyproof, mechanism can attain a constant approximation ratio of optimal social cost. However, such a mechanism can be constant-Lipschitz as we show with the mechanism M_2 (Section 3.2). We use the following lemma in our analysis of M_2.

Lemma 3.2. *Let π be a profile and let σ denote $M_1(\pi)$. Then $|\{i \mid \pi_i \leq \sigma_{1,1}\}|$ is at least $\lceil |S_1(\pi)|/2 \rceil$ and $|\{i \mid \pi_i \geq \sigma_{2,1}\}|$ is at least $\lceil |S_2(\pi)|/2 \rceil$.*

Proof. Follows from Lemma 2.2. □

3.2 Mechanism M_2

Below we define our second mechanism M_2. Before doing so, we provide some useful definitions.

For any n-profile π, let $h_1(\pi)$ denote the unique real number x greater than or equal to π_1 such that $\sum_{i\in[n]} \max(0, x - \pi_i) = C(\pi)$, let $f_1(\pi)$ denote $\min(h_1(\pi), \mu(\pi))$. Similarly, let $h_2(\pi)$ denote the unique real number x less than or equal to π_n such that $\sum_{i\in[n]} \max(0, \pi_i - x) = C(\pi)$, and let $f_2(\pi)$ denote $\max(h_2(\pi), \mu(\pi))$. It is easy to argue that $f_1(\pi) = \mu(\pi)$ if and only if $f_2(\pi) = \mu(\pi)$.

The reason that we introduce $f_1(\pi)$ and $f_2(\pi)$ is because it is possible that $h_1(\pi) \geq h_2(\pi)$. For example, let π denote the profile $(-2, -1, 0, 0, 1, 1)$. It is straightforward to check that $h_1(\pi) = 0$ and $h_2(\pi) = -1/4$. However, for all ℓ belonging to $\{1, 2\}$, $f_\ell(\pi) = -1/6$.

We define mechanism M_2 as follows: For any profile π, $M_2(\pi)$ is the 1-D solution $((f_1(\pi), 0), (f_2(\pi), 0))$.

We now show that mechanism M_2 is constant-approximate and establish properties to be used in Section 4.2 to prove that mechanism M_4 is constant-Lipschitz. Lemma 4.10, which proves that mechanism M_4 is constant-Lipschitz, also establishes that mechanism M_2 is constant-Lipschitz.

We begin with the following lemma towards the goal of showing that mechanism M_2 is constant-approximate.

Lemma 3.3. *Let π be a profile and let σ denote $M_1(\pi)$. Then*

$$\sigma_{1,1} \leq f_1(\pi) \leq f_2(\pi) \leq \sigma_{2,1}.$$

Proof. Immediate from the definitions. □

The following lemma establishes that mechanism M_2 is constant-approximate.

Lemma 3.4. *Let π be an n-profile, let σ denote $M_2(\pi)$. Then $C(\pi, \sigma) \leq 3C(\pi)$.*

Proof. Let X denote $\{i \in [n] \mid \pi_i \leq f_1(\pi)\}$, Y denote $\{i \in [n] \mid \pi_i \geq f_2(\pi)\}$, and Z denote $[n] \setminus (X \cup Y)$. The definition of mechanism M_2 implies $\sum_{i\in X} d(\pi_i, \sigma) \leq C(\pi)$ and $\sum_{i\in Y} d(\pi_i, \sigma) \leq C(\pi)$. Lemma 3.3 implies $\sum_{i\in Z} d(\pi_i, \sigma) \leq C(\pi)$. The claim of the lemma follows. □

Again, given a trivial extension to the lower bound provided by [36], we know that mechanism M_2 cannot be constant-strategy proof because it is a 1-D mechanism that attains a constant approximation ratio of optimal social cost.

We now move on to establish properties of mechanism M_2 that are used for analyzing its 2-D generalization, mechanism M_4, in Section 4.2.

We begin with some useful definitions. For any profile π, let $\Delta(\pi)$ denote $f_2(\pi) - f_1(\pi)$, the distance between the two facilities on the x-axis. For any profile π such that $\Delta(\pi) > 0$ and any i in $[n]$, we define $w_2(\pi, i)$ as

$$\frac{\max(0, \min(f_2(\pi), \pi_i) - f_1(\pi))}{\Delta(\pi)}$$

and $w_1(\pi, i)$ as $1 - w_2(\pi, i)$. For any profile π such that $\Delta(\pi) > 0$ and any ℓ in $\{1, 2\}$, we define $w_\ell(\pi)$ as $\sum_{1 \leq i \leq n} w_\ell(\pi, i)$ and $\psi_\ell(\pi)$ as $C(\pi)/w_\ell(\pi)$.

Lemma 3.5. *For any ℓ in $\{1, 2\}$ and any profile π such that $\Delta(\pi) > 0$, we have*

$$w_\ell(\pi) \geq |S_\ell(\pi)|/2.$$

Proof. Immediate from Lemmas 3.2 and 3.3. □

Lemma 3.6. *Let π be a profile such that $\Delta(\pi) > 0$, let ℓ belong to $\{1, 2\}$, and let σ denote $M_2(\pi)$. Then*

$$w_\ell(\pi) \geq |S_\ell(\pi, \sigma)|/2.$$

Proof. Observe that each agent in $S_\ell(\pi, \sigma)$ contributes at least $1/2$ to $w_\ell(\pi)$. □

Lemma 3.7. *Let π be an n-profile such that $\Delta(\pi) > 0$ and let ℓ belong to $\{1, 2\}$. Then*

$$|S_\ell(\pi)| + C(\pi)/\Delta(\pi) \geq w_\ell(\pi).$$

Proof. Below we address the case $\ell = 2$. A similar argument holds for the case $\ell = 1$.

Let σ denote $M_1(\pi)$ and let x denote $(\sigma_{1,1} + \sigma_{2,1})/2$. We have

$$\begin{aligned} C(\pi) &\geq \sum_{i \in [n]: \pi_i \leq x} \max(0, \pi_i - \sigma_{1,1}) \\ &\geq \sum_{i \in [n]: \pi_i \leq x} \max(0, \pi_i - f_1(\pi)) \\ &\geq \sum_{i \in [n]: \pi_i \leq x} \max(0, \min(f_2(\pi), \pi_i) - f_1(\pi)) \\ &= \Delta(\pi) \sum_{i \in [n]: \pi_i \leq x} w_2(\pi, i). \end{aligned}$$

Moreover,

$$|S_2(\pi)| = |\{i \in [n] \mid \pi_i > x\}| \geq \sum_{i \in [n]: \pi_i > x} w_2(\pi, i).$$

Since $w_2(\pi) = \sum_{1 \leq i \leq n} w_2(\pi, i)$, we conclude that the claim of the lemma holds for $\ell = 2$. □

Lemma 3.8. *Let π be a profile such that $\Delta(\pi) > 0$, let ℓ belong to $\{1, 2\}$, and assume that $\Delta(\pi) \geq 2\psi_\ell(\pi)$. Then*

$$w_\ell(\pi) \leq 2|S_\ell(\pi)|.$$

Proof. The inequality $\Delta(\pi) \geq 2\psi_\ell(\pi)$ implies $C(\pi)/\Delta(\pi)$ is at most $w_\ell(\pi)/2$. Hence the claim follows from Lemma 3.7. □

4 Facility Location in $\mathbb{R}^2$

We now move on to defining 2-D mechanisms, both of which are constant-strategyproof and constant-approximate, one of which is also constant-Lipschitz.

4.1 Mechanism M_3

Consider the following 2-D generalization of mechanism M_1, which we refer to as mechanism M_3. Given a profile π, and letting σ denote $M_1(\pi)$, we define

$$M_3(\pi) = ((\sigma_{1,1}, C(\pi)/|S_1(\pi)|), (\sigma_{2,1}, C(\pi)/|S_2(\pi)|)).$$

In other words, M_3 is the generalization of M_1 in which each facility is vertically backed off of the x-axis by a distance equal to the optimal social cost divided by the number of agents served by that facility in M_1.[5]

The following lemma, which establishes that M_3 is constant-approximate, is straightforward to prove.

Lemma 4.1. *For any profile π, we have $C(\pi, M_3(\pi)) \leq 3C(\pi)$.*

We now present a sequence of lemmas to be used to establish that M_3 is constant-strategyproof.

In Lemma 4.2, Lemma 4.3, and Lemma 4.4, let ε denote $(2-\sqrt{3})/3$. For this choice of ε, one may verify that $(\frac{1}{3}-\varepsilon)(1-3\varepsilon) = 2\varepsilon$.

Lemma 4.2. *Let π be an n-profile, let i belong to $[n]$, let x be a real number, let π^* denote $\mathrm{subst}(\pi, i, x)$, let τ denote $M_3(\pi)$, let τ^* denote $M_3(\pi^*)$, let σ denote $((\tau_{1,1}, 0), (\tau_{2,1}, 0))$, let σ^* denote $((\tau^*_{1,1}, 0), (\tau^*_{2,1}, 0))$, and assume that $d(\pi_i, \sigma^*) \leq \varepsilon d(\pi_i, \tau)$. Then $C(\pi^*) \geq (\frac{1}{3}-\varepsilon)C(\pi, \tau)$.*

Proof. We have

$$\begin{aligned} C(\pi^*) &= C(\pi, \sigma^*) + d(x, \sigma^*) - d(\pi_i, \sigma^*) \\ &\geq C(\pi) + 0 - \varepsilon d(\pi_i, \tau) \\ &\geq \frac{C(\pi, \tau)}{3} - \varepsilon d(\pi_i, \tau) \\ &\geq \left(\frac{1}{3} - \varepsilon\right) C(\pi, \tau), \end{aligned}$$

where the second inequality follows from Lemma 4.1. □

Lemma 4.3. *Let π be an n-profile, let i belong to $[n]$, let x be a real number, let π^* denote $\mathrm{subst}(\pi, i, x)$, let τ denote $M_3(\pi)$, let τ^* denote $M_3(\pi^*)$, let σ^* denote $((\tau^*_{1,1}, 0), (\tau^*_{2,1}, 0))$, let ℓ belong to $\{1, 2\}$, and assume that $d(\pi_i, \sigma^*_\ell) \leq \varepsilon d(\pi_i, \tau)$. Then $\tau^*_{\ell,2} \geq \varepsilon d(\pi_i, \tau)$.*

[5] If π is trivial, then we set $\tau_{2,2} = 0$.

Proof. Let u be a bijection from $[n]$ to $[n]$ such that $\pi^*_{u(j)} = \pi_j$ for all j in $[n] \setminus \{i\}$ and $\pi^*_{u(i)} = x$, let J denote $\{j \in [n] \mid u(j) \in S_\ell(\pi^*, \sigma^*)\}$, let A denote $\{j \in J \mid d(\pi^*_{u(j)}, \sigma^*) \leq 2\varepsilon d(\pi_i, \tau)\}$, and let B denote $J \setminus A$.

For any j in A, we claim that $d(\pi_j, \tau) \geq (1 - 3\varepsilon) d(\pi_i, \tau)$. If $i = j$, the claim holds trivially. Suppose $i \neq j$. Since $|\sigma^*_\ell - \pi_i| = d(\pi_i, \sigma^*_\ell) \leq \varepsilon d(\pi_i, \tau)$ and $|\sigma^*_\ell - \pi^*_{u(j)}| = |\sigma^*_\ell - \pi_j| \leq 2\varepsilon d(\pi_i, \tau)$, the triangle inequality implies $|\pi_i - \pi_j| \leq 3\varepsilon d(\pi_i, \tau)$. Since $d(\pi_j, \tau)$ is at least $d(\pi_i, \tau) - |\pi_i - \pi_j|$, the claim follows.

The preceding claim implies $C(\pi, \tau) \geq \sum_{j \in A} d(\pi_j, \tau) \geq (1 - 3\varepsilon)|A| \cdot d(\pi_i, \tau)$. Since $d(\pi_i, \sigma^*) \leq d(\pi_i, \sigma^*_\ell) \leq \varepsilon d(\pi_i, \tau)$, Lemma 4.2 implies

$$C(\pi^*) \geq \left(\frac{1}{3} - \varepsilon\right)(1 - 3\varepsilon)|A| \cdot d(\pi_i, \tau).$$

The definition of B implies $d(\pi^*_{u(j)}, \sigma^*) \geq 2\varepsilon d(\pi_i, \tau)$ for all j in B and hence $C(\pi^*) \geq \sum_{j \in B} d(\pi^*_{u(j)}, \sigma^*) \geq 2\varepsilon |B| \cdot d(\pi_i, \tau)$. Thus

$$C(\pi^*) \geq \max\left[\left(\frac{1}{3} - \varepsilon\right)(1 - 3\varepsilon)|A|, 2\varepsilon|B|\right] d(\pi_i, \tau).$$

Recall that $(\frac{1}{3} - \varepsilon)(1 - 3\varepsilon) = 2\varepsilon$. Thus

$$C(\pi^*) \geq 2\varepsilon \max(|A|, |B|) d(\pi_i, \tau) \geq \varepsilon |J| \cdot d(\pi_i, \tau)$$

and hence $\tau^*_{\ell,2} = C(\pi^*)/|J| \geq \varepsilon d(\pi_i, \tau)$. □

Lemma 4.4. *Mechanism M_3 is $\frac{1}{\varepsilon}$-strategyproof.*

Proof. Let π be an n-profile, let i belong to $[n]$, let x be a real number, let π^* denote $\text{subst}(\pi, i, x)$, let τ denote $M_3(\pi)$, let τ^* denote $M_3(\pi^*)$, and let σ^* denote $((\tau^*_{1,1}, 0), (\tau^*_{2,1}, 0))$. We need to prove that $\varepsilon d(\pi_i, \tau) \leq d(\pi_i, \tau^*)$. Assume for the sake of contradiction that $d(\pi_i, \tau^*) < \varepsilon d(\pi_i, \tau)$ and let ℓ be an element of $\{1, 2\}$ such that i belongs to $S_\ell(\pi, \tau^*)$. Then

$$d(\pi_i, \sigma^*_\ell) \leq d(\pi_i, \tau^*_\ell) = d(\pi_i, \tau^*) < \varepsilon d(\pi_i, \tau)$$

and hence Lemma 4.3 implies $\tau^*_{\ell,2} \geq \varepsilon d(\pi_i, \tau)$. It follows that $d(\pi_i, \tau^*) \geq \varepsilon d(\pi_i, \tau)$, a contradiction. □

Theorem 1. *Mechanism M_3 is constant-approximate and constant-strategyproof.*

Proof. The theorem follows immediately from Lemma 4.1 and Lemma 4.4. □

The proof of Lemma 3.1 can be easily adapted to show that M_3 is neither strategyproof nor constant-Lipschitz.

4.2 Mechanism M_4

Our final mechanism, M_4, is a 2-D generalization of mechanism M_2. In order to define M_4, we first introduce the following definitions for any profile π and any ℓ in $\{1,2\}$: $\xi_\ell(\pi)$ denotes 0 if $\Delta(\pi) = 0$ and $\min(\Delta(\pi), 2\psi_\ell(\pi))$ otherwise; $\varphi_\ell(\pi)$ denotes $\max(8C(\pi)/n, \xi_\ell(\pi))$.

Given a profile π, and letting σ denote $M_2(\pi)$, we define $M_4(\pi)$ as the solution τ such that $\tau_{\ell,1} = \sigma_{\ell,1}$ and $\tau_{\ell,2} = \varphi_\ell(\pi)$ for all ℓ in $\{1,2\}$.

Lemma 4.5. *For any n-profile π, we have $C(\pi, M_4(\pi)) \le 15C(\pi)$.*

Proof. Let σ denote $M_2(\pi)$. Observe that

$$C(\pi, M_4(\pi)) \le C(\pi, \sigma) + \sum_{\ell \in \{1,2\}} \varphi_\ell(\pi)|S_\ell(\pi, \sigma)|. \tag{1}$$

Since $C(\pi, \sigma) \le 3C(\pi)$ by Lemma 3.4, it is sufficient to prove that the sum in Equation (1) is at most $12C(\pi)$. If $\Delta(\pi) = 0$, then $\varphi_\ell(\pi) = 8C(\pi)/n$ for all ℓ in $\{1,2\}$ and hence this sum is equal to $8C(\pi)$.

For the remainder of the proof, assume that $\Delta(\pi) > 0$. For any ℓ in $\{1,2\}$, we have

$$\begin{aligned}\varphi_\ell(\pi)|S_\ell(\pi,\sigma)| &\le \max(8C(\pi)/n, 2\psi_\ell(\pi))|S_\ell(\pi,\sigma)| \\ &\le \max(8C(\pi)/n, 4C(\pi)/|S_\ell(\pi,\sigma)|)|S_\ell(\pi,\sigma)| \\ &= \max(8C(\pi)|S_\ell(\pi,\sigma)|/n, 4C(\pi)),\end{aligned}$$

where the first inequality follows from the definition of $\varphi_\ell(\pi)$ and the second inequality follows from Lemma 3.6. Thus the sum in Equation (1) is at most $12C(\pi)$. □

We next define a core lemma which we will use to relate the constant-strategyproofness of mechanism M_3 to that of mechanism M_4.

Lemma 4.6. *Let M and M^* be mechanisms, let a_1 and a_2 be positive real numbers such that $a_1 d(x, M(\pi)) \le d(x, M^*(\pi)) \le a_2 d(x, M(\pi))$ for all profiles π and all real numbers x, and assume that M is λ-strategyproof. Then M^* is $\lambda a_2/a_1$-strategyproof.*

Proof. Let π be an n-profile, let i belong to $[n]$, let x be a real number, and let π^* denote $\mathrm{subst}(\pi, i, x)$. We have

$$\begin{aligned}d(\pi_i, M^*(\pi)) &\le a_2 \cdot d(\pi_i, M(\pi)) \\ &\le \lambda a_2 \cdot d(\pi_i, M(\pi^*)) \\ &\le (\lambda a_2/a_1) d(\pi_i, M^*(\pi^*)).\end{aligned}$$

□

We next show that mechanism M_4 is constant-strategyproof. The following simple fact is used in the proof of Lemma 4.7 below.

Fact 2 *Let p and q be two points such that p_2 and q_2 are positive and let z be a real number such that $|p_1 - q_1| \leq z$. Then*

$$\min\left(1, \frac{p_2}{q_2+z}\right) \leq \frac{d(x,p)}{d(x,q)} \leq \max\left(1, \frac{p_2+z}{q_2}\right)$$

for all real numbers x.

Lemma 4.7. *Mechanism M_4 is constant-strategyproof.*

Proof. Since mechanism M_3 is constant-strategyproof, Lemma 4.6 implies it is sufficient to exhibit positive real numbers a_1 and a_2 such that, for all profiles π and all real numbers x, $a_1 d(x, M_3(\pi)) \leq d(x, M_4(\pi)) \leq a_2 d(x, M_3(\pi))$.

Let π be an n-profile, let x be a real number, let σ denote $M_1(\pi)$, let σ^* denote $M_2(\pi)$, let τ denote $M_3(\pi)$, let τ^* denote $M_4(\pi)$, and let s_ℓ denote $|S_\ell(\pi)|$ for all ℓ in $\{1,2\}$. Thus $\tau_{\ell,2} = C(\pi)/s_\ell$.

If $C(\pi) = 0$, it is easy to see that $\tau^* = \tau$ and hence $d(x,\tau^*) = d(x,\tau)$. For the remainder of the proof, we assume that $C(\pi) > 0$. It follows that $\tau_{\ell,2}$, $\tau^*_{\ell,2}$, and s_ℓ are positive for all ℓ in $\{1,2\}$. We assume without loss of generality that $s_1 \geq n/2$; the case $s_2 \geq n/2$ is symmetric.

We claim that $\tau^*_{1,2} = 8C(\pi)/n$. If $\Delta(\pi) \leq 8C(\pi)/n$, the claim is immediate from the definition of mechanism M_4. Otherwise, $\Delta(\pi) > 0$ and since $s_1 \geq n/2$, Lemma 3.5 implies $w_1(\pi) \geq n/4$, which in turn implies $2\psi_1(\pi) \leq 8C(\pi)/n$, so the claim is once again immediate from the definition of mechanism M_4.

We claim that $|\sigma^*_{\ell,1} - \sigma_{\ell,1}| \leq 2C(\pi)/s_\ell$ for all ℓ in $\{1,2\}$. We prove this claim for the case $\ell = 1$; the case $\ell = 2$ is symmetric. By Lemma 3.3, we need to prove that $\sigma^*_{1,1} \leq z$ where z denotes $\sigma_{1,1} + 2C(\pi)/s_1$. By Lemma 3.2, there are at least $s_1/2$ agents at or to the left of $\sigma_{1,1}$. Thus $\sum_{i\in[n]:\pi_i\leq z} z - \pi_i$ is at least $C(\pi)$. Hence the definition of mechanism M_2 implies $\sigma^*_{1,1} \leq z$, as required.

Given that $n/2 \leq s_1 \leq n$ and using Fact 2 with p_2 equal to $8C(\pi)/n$, q_2 equal to $C(\pi)/s_1$, and z equal to $2C(\pi)/s_1$, we have

$$\begin{aligned}\frac{d(x,\tau_1^*)}{d(x,\tau_1)} &\leq \max\left(1, \max_{n/2\leq s_1\leq n} \frac{8C(\pi)/n + 2C(\pi)/s_1}{C(\pi)/s_1}\right)\\ &= \max(1, \max_{n/2\leq s_1\leq n} 8s_1/n + 2)\\ &= 10.\end{aligned}$$

We can use Fact 2 in a similar manner to derive a lower bound of 1 for the ratio $d(x,\tau_1^*)/d(x,\tau_1)$. In summary, we have

$$1 \leq d(x,\tau_1^*)/d(x,\tau_1) \leq 10. \quad (2)$$

Case 1: $\Delta(\pi) \geq 2\psi_2(\pi)$. Lemmas 3.5 and 3.8 imply $w_2(\pi)/2 \leq s_2 \leq 2w_2(\pi)$. Since $|\sigma^*_{2,1} - \sigma_{2,1}| \leq 2C(\pi)/s_2$, we deduce that $|\sigma^*_{2,1} - \sigma_{2,1}| \leq 4\psi_2(\pi)$. Below we use two subcases to establish that

$$1/3 \leq d(x,\tau_2^*)/d(x,\tau_2) \leq 12. \quad (3)$$

Equations (2) and (3) together imply $1/3 \leq d(x,\tau^*)/d(x,\tau) \leq 12$.

Case 1.1: $w_2(\pi) \geq n/4$. Thus $\tau^*_{2,2} = 8C(\pi)/n$ and $n/8 \leq s_2 \leq n/2$. Hence $2C(\pi)/n \leq \tau_{2,2} \leq 8C(\pi)/n$ and $|\sigma^*_{2,1} - \sigma_{2,1}| \leq 16C(\pi)/n$. Using Fact 2 with p_2 equal to $8C(\pi)/n$, q_2 between $2C(\pi)/n$ and $8C(\pi)/n$, and z equal to $16C(\pi)/n$, we find that Equation (3) holds.

Case 1.2: $w_2(\pi) \leq n/4$. Thus $\tau^*_{2,2} = 2\psi_2(\pi)$. Since $w_2(\pi)/2 \leq s_2 \leq 2w_2(\pi)$, we have $\psi_2(\pi)/2 \leq \tau_{2,2} \leq 2\psi_2(\pi)$. Using Fact 2 with p_2 equal to $2\psi_2(\pi)$, q_2 between $\psi_2(\pi)/2$ and $2\psi_2(\pi)$, and z equal to $4\psi_2(\pi)$, we again find that Equation (3) holds.

Case 2: $8C(\pi)/n \leq \Delta(\pi) \leq 2\psi_2(\pi)$. Thus $\tau^*_{2,2} = \Delta(\pi)$. Since $\tau^*_{1,2} = 8C(\pi)/n$, we have $d(x,\tau^*_1) \leq 2d(x,\tau^*_2)$. Thus

$$\begin{aligned} d(x,\tau^*) &= \min(d(x,\tau^*_1), d(x,\tau^*_2)) \\ &\geq \min(d(x,\tau^*_1), d(x,\tau^*_1)/2) \\ &= d(x,\tau^*_1)/2 \\ &\geq d(x,\tau_1)/2 \\ &\geq d(x,\tau)/2, \end{aligned}$$

where the second inequality follows from Equation (2).

Recall that $w_2(\pi) \geq s_2/2$. Thus the case condition implies $\tau^*_{2,2} \leq 4C(\pi)/s_2$. Recall that $|\sigma^*_{2,1} - \sigma_{2,1}| \leq 2C(\pi)/s_2$. Using Fact 2 with p_2 equal to $4C(\pi)/s_2$, q_2 equal to $C(\pi)/s_2$, and z equal to $2C(\pi)/s_2$, we find that

$$\frac{d(x,\tau^*_2)}{d(x,\tau_2)} \leq 6. \tag{4}$$

Thus

$$\begin{aligned} d(x,\tau^*) &= \min(d(x,\tau^*_1), d(x,\tau^*_2)) \\ &\leq \min(10d(x,\tau_1), 6d(x,\tau_2)) \\ &\leq 10\min(d(x,\tau_1), d(x,\tau_2)) \\ &= 10d(x,\tau), \end{aligned}$$

where the first inequality follows from Equations (2) and (4).

Case 3: $\Delta(\pi) \leq 8C(\pi)/n$. Thus $\tau^*_{2,2} = 8C(\pi)/n$ and since $\tau^*_{1,2} = 8C(\pi)/n$, we have $d(x,\tau^*_1) \leq 2d(x,\tau^*_2)$. As in Case 2, we deduce that $d(x,\tau^*) \geq d(x,\tau)/2$.

Recall that $|\sigma^*_{2,1} - \sigma_{2,1}| \leq 2C(\pi)/s_2$. Given that $1 \leq s_2 \leq n/2$ and using Fact 2 with p_2 equal to $8C(\pi)/n$, q_2 equal to $C(\pi)/s_2$, and z equal to $2C(\pi)/s_2$, we find that Equation (4) holds as in Case 2. The rest of the argument proceeds as in Case 2. □

The remainder of this section is devoted to establishing that mechanism M_4 is constant-Lipschitz. The plan is to argue Lipschitz-type bounds for the functions $f_\ell(\pi)$ and $\varphi_\ell(\pi)$ that define the four components of the solution produced by M_4. These functions are defined in terms of a number of simpler auxiliary functions. Since the functions that we are studying all map a given π in $\mathbb{R}^n$ to $\mathbb{R}$, a natural approach to deriving Lipschitz-type bounds is to study the magnitude of the gradient. Instead of directly reasoning about the gradient, we find it convenient to fix all components except component k of a given profile π, and then study the magnitude of the derivative as π_k varies over $I^* = \text{interval}(\pi, k)$.

A technical obstacle that we need to overcome is that most of the functions we study are not differentiable over all of I^*; rather, they are piecewise differentiable. In the full version of the paper [23], we first show how to break I^* into a finite set of pieces, each of which is a closed interval, such that all of the functions of interest are differentiable over the interior of each piece. We then use calculus to bound the magnitude of the derivative of each of these functions over the interior of each piece. Combining these bounds with some other basic results for establishing Lipschitz-type bounds, we obtain our main technical lemma, which states that the functions $f_\ell(\pi)$ and $\varphi_\ell(\pi)$ are constant-Lipschitz over each piece. Using this lemma, it is straightforward to establish that the functions $f_\ell(\pi)$ and $\varphi_\ell(\pi)$ are constant-Lipschitz over I^*; see the full version of the paper for the proof details. Below we use the latter result to establish Lemma 4.8, which provides a Lipschitz-type bound for mechanism M_4 when a profile π is changed to another profile in $\Gamma(\pi)$. We then use Lemma 4.8 to obtain the desired constant-Lipschitz bound for M_4 (see Lemma 4.10 below).

Lemma 4.8. *There exists a positive constant κ such that for any profile π and any profile π^* in $\Gamma(\pi)$, we have*

$$\|M_4(\pi^*) - M_4(\pi)\|_1 \leq \kappa \cdot \|\pi^* - \pi\|_1.$$

Proof. As discussed above, in the full version of the paper we establish that the functions $f_\ell(\pi)$ and $\varphi_\ell(\pi)$ are constant-Lipschitz over I^* for all ℓ in $\{1, 2\}$. The present lemma follows immediately. □

Lemma 4.9. *Let π and π^* be distinct n-profiles. Then there exists an i in $[n]$ such that $\pi_i \neq \pi_i^*$ and π_i^* belongs to* interval(π, i).

Proof. If there exists a j in $[n]$ such that $\pi_j < \pi_j^*$, then we can take i to be the maximum such j. Otherwise, we can take i to be the minimum j in $[n]$ such that $\pi_j > \pi_j^*$. □

We finally state our lemma that mechanism M_4 is constant-Lipschitz as follows.

Lemma 4.10. *Let κ denote the Lipschitz constant of Lemma 4.8. Then for any n-profiles π and π^*, we have*

$$\|M_4(\pi^*) - M_4(\pi)\|_1 \leq \kappa \cdot \|\pi^* - \pi\|_1.$$

Proof. Let π and π^* be n-profiles and let s denote $|\{i \in [n] \mid \pi_i \neq \pi_i^*\}|$. By applying Lemma 4.9 s times, we find that there exists a sequence of n-profiles

$\pi^{(0)}, \ldots, \pi^{(s)}$ such that $\pi^{(0)} = \pi$, $\pi^{(s)} = \pi^*$, and $\pi^{(i)}$ belongs to $\Gamma(\pi^{(i-1)})$ for all i in $[s]$. Thus

$$\begin{aligned}\|M_4(\pi^*) - M_4(\pi)\|_1 &\leq \sum_{i \in [s]} \|M_4(\pi^{(i)}) - M_4(\pi^{(i-1)})\|_1 \\ &\leq \kappa \sum_{i \in [s]} \|\pi^{(i)} - \pi^{(i-1)}\|_1 \\ &= \kappa \cdot \|\pi^* - \pi\|_1,\end{aligned}$$

where the second inequality follows from Lemma 4.8. □

Theorem 2. *Mechanism M_4 is constant-approximate, constant-strategyproof, and constant-Lipschitz.*

Proof. The theorem follows immediately from Lemma 4.5, Lemma 4.7, and Lemma 4.10. □

5 Concluding Remarks

Our work shows that constant-strategyproof and constant-Lipschitz mechanisms can provide good truthfulness and stability guarantees without suffering significant loss of efficiency. For example, consider the stark difference between a $\Theta(n)$ approximation ratio of optimal social cost for strategyproof mechanisms and a constant approximation ratio of optimal social cost for constant-strategyproof mechanisms. Our positive results suggest that these properties are deserving of further study for the facility location problem and other problems in mechanism design.

There are many interesting directions in which our work can be generalized. Most pressingly, we are eager to see whether our methods can be generalized to facility location in other metric spaces. In particular, we conjecture that it is possible to design constant-strategy proof mechanisms for facility location with both preferences and facilities in $\mathbb{R}^d$. Beyond this, there are other natural directions for generalizing our results such as allowing for more than two facilities or considering different cost functions for individual agent or social costs [15,30].

There are also numerous other properties which are of interest. For example, future work could integrate recent fairness results for facility location [26,27,39,50] with constant-strategyproof and constant-Lipschitz mechanisms. Additionally, our results implicitly reveal a tradeoff between the approximation ratio and the relaxation of strategyproofness, governed by the choice of vertical offset from the line. It would be interesting to more precisely characterize this tradeoff.

The integration of theoretical results with practical, real-world, considerations is another potentially fruitful direction for future research. Examples of work in this vein include [7] which studies the effective use on individual mobility pattern data for facility location and [2,25,35] which study "obvious" strategyproofness. Ultimately, facility location is a theoretical problem with deep practical motivations. Any work that attempts to narrow the gap between theory and implementation could be valuable both for embedding better theoretical guarantees into practical implementations and for informing which problems and relaxations are most deserving of theoretical study.

References

1. Alon, N., Feldman, M., Procaccia, A., Tennenholtz, M.: Strategyproof approximation of the minimax on networks. Mathematics of Operations Research **35**(3), 513–526 (2010)
2. Ashlagi, I., Gonczarowski, Y.A.: Stable matching mechanisms are not obviously strategy-proof. Journal of Economic Theory **177**, 405–425 (2018)
3. Auricchio, G., Zhang, J.: The k-facility location problem via optimal transport: A Bayesian study of the percentile mechanisms. In: Proceedings of the 17th International Symposium on Algorithmic Game Theory. pp. 147–164 (Sep 2024)
4. Barberà, S., Gul, F., Stacchetti, E.: Generalized median voter schemes and committees. Journal of Economic Theory **61**(2), 262–289 (1993)
5. Birrell, E., Pass, R.: Approximately strategy-proof voting. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence. pp. 67–72 (Jul 2011)
6. Black, D.: On the rationale of group decision-making. Journal of Political Economy **56**(1), 23–34 (1948)
7. Candogan, O., Feng, Y.: Mobility data in operations: Multi-location facility location problem. In: Proceedings of the 25th ACM Conference on Economics and Computation. p. 201 (Jul 2024)
8. Chan, H., Filos-Ratsikas, A., Li, B., Li, M., Wang, C.: Mechanism design for facility location problems: A survey. In: Proceedings of the 13th International Joint Conference on Artificial Intelligence. pp. 4356–4365 (Aug 2021)
9. Chen, N., Deng, X., Zhang, J.: How profitable are strategic behaviors in a market? In: Proceedings of the European Symposium on Algorithms. pp. 106–118 (Sep 2011)
10. Chen, P., Egesdal, M., Pycia, M., Yenmez, M.B.: Manipulability of stable mechanisms. American Economic Journal: Microeconomics **8**(2), 202–14 (2016)
11. Cheng, Y., Yu, W., Zhang, G.: Mechanisms for obnoxious facility game on a path. In: Proceedings of the 5th International Conference on Combinatorial Optimization and Applications. pp. 262–271 (Aug 2011)
12. Cheng, Y., Yu, W., Zhang, G.: Strategy-proof approximation mechanisms for an obnoxious facility game on networks. Theoretical Computer Science **497**, 154–163 (2013)
13. Church, R.L., Garfinkel, R.S.: Locating an obnoxious facility on a network. Transportation Science **12**(2), 107–118 (1978)
14. Dale, E., Fielding, J., Ramakrishnan, H., Sathyanarayanan, S., Weinberg, S.M.: Approximately strategyproof tournament rules with multiple prizes. In: Proceedings of the 23rd ACM Conference on Economics and Computation. pp. 1082–1100 (Jul 2022)
15. Deligkas, A., Lotfi, M., Voudouris, A.A.: Agent-constrained truthful facility location games. In: Proceedings of the 17th International Symposium on Algorithmic Game Theory. pp. 129–146 (Sep 2024)
16. Ding, K., Weinberg, S.M.: Approximately strategyproof tournament rules in the probabilistic setting. In: Proceedings of the 12th Innovations in Theoretical Computer Science Conference. pp. 14:1–14:20 (Jan 2021)

17. Escoffier, B., Gourvès, L., Nguyen, K.T., Pascual, F., Spanjaard, O.: Strategy-proof mechanisms for facility location games with many facilities. In: Proceedings of the 2nd International Conference on Algorithmic Decision Theory. pp. 67–81 (Oct 2011)
18. Fernandes, C.G., Meira, L.A.A., Miyazawa, F.K., Pedrosa, L.L.C.: A systematic approach to bound factor-revealing LPs and its application to the metric and squared metric facility location problems. Mathematical Programming **153**(2), 655–685 (2015)
19. Filimonov, A., Meir, R.: Strategyproof facility location mechanisms on discrete trees. Autonomous Agents and Multi-Agent Systems **37**(10) (2022)
20. Fotakis, D., Kavouras, L., Kostopanagiotis, P., Lazos, P., Skoulakis, S., Zarifis, N.: Reallocating multiple facilities on the line. Theoretical Computer Science **858**, 13–34 (2021)
21. Fotakis, D., Tzamos, C.: On the power of deterministic mechanisms for facility location games. ACM Transactions on Economics and Computation **2**(4) (2014)
22. Fukui, Y., Shurbevski, A., Nagamochi, H.: λ-group strategy-proof mechanisms for the obnoxious facility game in star networks. IEICE Transactions on Fundamentals **E102-A**(9), 1179–1186 (2019)
23. Fullerton, E.J., Hu, Z., Plaxton, C.G.: Constant-approximate and constant-strategyproof two-facility location (2025), https://arxiv.org/abs/2507.04485
24. Goel, S., Hann-Caruthers, W.: Optimality of the coordinate-wise median mechanism for strategyproof facility location in two dimensions. Social Choice Welfare **61**(1), 11–34 (2023)
25. Gonczarowski, Y.A., Heffetz, O., Thomas, C.: Strategyproofness-exposing mechanism descriptions. In: Proceedings of the 24th ACM Conference on Economics and Computation. p. 782 (Jul 2023)
26. Gupta, S., Moondra, J., Singh, M.: Which L_p norm is the fairest? Approximations for fair facility location across all "p". In: Proceedings of the 24th ACM Conference on Economics and Computation. p. 817 (Jul 2023)
27. Hossain, S., Micha, E., Shah, N.: The surprising power of hiding information in facility location. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence. pp. 2168–2175 (Feb 2020)
28. Hyafil, N., Boutilier, C.: Regret-based incremental partial revelation mechanisms. In: Proceedings of the 21st AAAI Conference on Artificial Intelligence. pp. 672–678 (Jul 2006)
29. Istrate, G., Bonchis, C.: Mechanism design with predictions for obnoxious facility location (2022), https://arxiv.org/abs/2212.09521
30. Kanellopoulos, P., Voudouris, A.A., Zhang, R.: Truthful two-facility location with candidate locations. Theoretical Computer Science **1024**, 114913 (2025)
31. de Keijzer, B., Wojtczak, D.: Facility reallocation on the line. Algorithmica **84**(10), 2898–2925 (2022)
32. Kim, K.H., Roush, F.W.: Nonmanipulability in two dimensions. Mathematical Social Sciences **8**(1), 29–43 (1984)
33. Klootwijk, S., Manthey, B.: Probabilistic analysis of facility location on random shortest path metrics. In: Proceedings of the 15th Conference on Computability in Europe. pp. 37–49 (Jul 2019)
34. Lee, D.T.: Efficient, private, and ϵ-strategyproof elicitation of tournament voting rules. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence. pp. 2026–2032 (Jul 2015)
35. Li, S.: Obviously strategy-proof mechanisms. American Economic Review **107**(11), 3257–3287 (2017)

36. 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. pp. 315–324 (Jun 2010)
37. Lubin, B., Parkes, D.C.: Approximate strategyproofness. Current Science **103**(9), 1021–1032 (2012)
38. Mahdian, M., Ye, Y., Zhang, J.: Approximation algorithms for metric facility location problems. SIAM Journal on Computing **36**(2), 411–432 (2006)
39. Micha, E., Shah, N.: Proportionally fair clustering revisited. In: Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (Jul 2020)
40. Moulin, H.: On strategy-proofness and single peakedness. Public Choice **35**(4), 437–455 (1980)
41. Oomine, M., Shurbevski, A., Nagamochi, H.: Parameterization of strategy-proof mechanisms in the obnoxious facility game. Journal of Graph Algorithms and Applications **21**(3), 247–263 (2017)
42. Peters, H., van der Stel, H., Storcken, T.: Pareto optimality, anonymity, and strategy-proofness in location problems. International Journal of Game Theory **21**, 221–235 (1992)
43. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. ACM Transactions on Economics and Computation **1**(4) (2013)
44. Schneider, J., Schvartzman, A., Weinberg, S.M.: Condorcet-consistent and approximately strategyproof tournament rules. In: Proceedings of the 8th Innovations in Theoretical Computer Science Conference. pp. 35:1–35:20 (Jan 2017)
45. Schummer, J., Vohra, R.V.: Strategy-proof location on a network. Journal of Economic Theory **104**(2), 405–428 (2002)
46. Schvartzman, A., Weinberg, S.M., Zlatin, E., Zuo, A.: Approximately strategyproof tournament rules: On large manipulating sets and cover-consistence. In: Proceedings of the 11th Innovations in Theoretical Computer Science Conference. pp. 3:1–3:25 (Jan 2020)
47. Sui, X., Boutilier, C.: Approximately strategy-proof mechanisms for (constrained) facility location. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems. pp. 605–613 (May 2015)
48. Sui, X., Boutilier, C., Sandholm, T.: Analysis and optimization of multi-dimensional percentile mechanisms. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence. pp. 367–374 (Aug 2013)
49. Tang, P., Yu, D., Zhao, S.: Characterization of group-strategyproof mechanisms for facility location in strictly convex space. In: Proceedings of the 21st ACM Conference on Economics and Computation. pp. 133–157 (Jul 2020)
50. Zhou, H., Li, M., Chan, H.: Strategyproof mechanisms for group-fair facility location problems. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence (Jul 2022)