Learning-Augmented Mechanism Design for Facility Location under L_p-Norm Social Costs

Hau Chan[1] **Jianan Lin**[2] **Chenhao Wang**[3,4]*
[1]University of Nebraska-Lincoln, Lincoln, NE, USA
[2]Rensselaer Polytechnic Institute, Troy, NY, USA
[3]Beijing Normal University, Zhuhai, Guangdong, China
[4]Beijing Normal-Hong Kong Baptist University, Zhuhai, Guangdong, China
hchan3@unl.edu hcm6755@gmail.com chenhwang@bnu.edu.cn

Abstract

We consider learning-augmented mechanism design for facility location in the standard setting of locating a single facility on the line. Leveraging the given (imperfect) prediction on the optimal facility location, our goal is to design strategyproof (SP) mechanisms that truthfully elicit agent location preferences and determine facility locations that approximately minimize the L_p-norm social cost for $p \in (1, +\infty)$, achieving high *consistency* (i.e., approximation ratios when the prediction is correct) and *robustness* (i.e., approximation ratios when the prediction is incorrect). For deterministic SP mechanisms, we propose a family of generalized median mechanisms parameterized by a fraction $t \in (0, 1)$ of phantom points placed at the predicted location π. We show that this mechanism achieves $\left(1 + \left(\frac{1-t}{1+t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$-consistency and $\left(1 + \left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$-robustness and that no deterministic SP mechanism with the same consistency can obtain a better robustness. For randomized SP mechanisms, we obtain an upper bound by mixing our deterministic mechanism with centroid random dictatorship (CRD), and also provide a general lower bound on consistency–robustness trade-offs. The resulting finite-n upper bound applies to every $p \in (1, +\infty)$ and unifies the previously separate squared-cost and two-agent cases.

1 Introduction

Mechanism design for facility location has continued to attract significant research interest in the economics, operations research, and computer science communities since its introduction and systematic study decades ago [24, 25, 7]. It has a wide range of applications, including urban planning (e.g., determining locations of public facilities based on community preferences) [7] and voting (e.g., selecting a position to represent voter preferences) [24].

In the standard mechanism design setting of facility location [24, 25], a social planner seeks to locate a facility (e.g., a library, park, or school) within a region (e.g., alongside a street) to best serve a set of agents, in which each agent has a location preference for their desirable location of the facility, by minimizing a given cost objective defined on some aggregation of the agents' costs. These agents' costs are often defined as distances between agents' preferred locations and the facility location. Unfortunately, the social planner does not know and is required to elicit the agent location preferences. Furthermore, agents can misreport preferences to manipulate facility locations to minimize their own costs due to the interdependence between the cost objective, agent costs, and preferences. As a result, the typical research goal is to design approximately optimal strategyproof mechanisms that elicit

*Corresponding author. chenhwang@bnu.edu.cn

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

agent location preferences truthfully (i.e., incentivize agents to report true preferences) and determine facility locations optimizing the given cost objective approximately.

To guide the design of strategyproof (SP) mechanisms, Moulin's seminal study [24] established generalized median voter schemes as a canonical class of strategyproof rules on single-peaked domains. For the same setting, Procaccia and Tennenholtz [25] initiated the design of approximately optimal SP mechanisms that minimize the social cost and maximum cost objectives, which are defined as the total/sum (L_1-norm) and maximum (L_∞-norm) distances between the facility location and the agent preferred locations, respectively. While they showed that the median mechanism (i.e., locating the facility at the median of the preferred locations) is SP and minimizes the social cost objective optimally, they also showed that no SP mechanism can approximate the maximum cost objective better than 2 and that the leftmost mechanism (i.e., locating the facility at the leftmost of the preferred locations) is SP and approximates the maximum cost objective within a (multiplicative) factor of 2. Consequently, the work of Procaccia and Tennenholtz [25] has led to designing approximately optimal SP mechanisms for various studies of mechanism design for facility location, including the consideration of alternative cost objectives (e.g., L_2-norm [18], L_p-norm [17, 19], and fairness [29]), other regions/metric spaces (e.g., graphs and higher dimensions [2, 23, 20]), and agent preferences (e.g., obnoxious [13, 28, 10], hybrid [16]). We refer readers to a survey [7] of other studies and extensions.

Learning-Augmented Mechanism Design for Facility Location. In recent years, there has been a growing body of work that incorporates "predictions" into mechanism design for facility location [1, 27, 6, 21, 15, 12, 14, 4, 8, 3]. In this context, predictions can be viewed as outputs generated by machine learning models or as advice from domain experts based on historical data on optimal facility locations or agent profiles aiming to optimize the given cost objective.

Despite these predictions being imperfect, ranging from perfectly correct to completely wrong, the work of [1, 27], who first initiated the learning-augmented mechanism design for facility location, demonstrated the possibility of designing SP mechanisms that obtain high *consistency*, the approximation ratio when the prediction is perfect, and *robustness*, the worst-case ratio when the prediction is arbitrarily inaccurate, to overcome existing impossibility results of existing studies. For instance, the work of [27] examined predictions on the optimal facility location in the above mentioned standard setting under the maximum cost objective and designed an SP mechanism that achieves 1-consistency and 2-robustness. That is, the mechanism outputs a facility location that optimally minimizes the maximum cost objective when predictions are correct and approximately minimizes it within a factor of 2 when predictions are incorrect.

While learning-augmented mechanism design for facility location under the social cost and maximum cost objectives is well-understood in the standard setting, it is not immediately clear whether predictions can overcome the impossibility results of the more general family of L_p-norm social cost objectives for any $p \geq 1$ that provides a continuous interpolation between the social cost objective (when $p = 1$) and maximum cost objective (when $p = \infty$). Concretely, as p increases, the objective becomes increasingly sensitive to large individual agent costs, thereby emphasizing the need to decrease their costs. Without prediction, in the standard setting, Feigenbaum et al. [17] showed that the median mechanism is $2^{1-1/p}$-approximate for minimizing the L_p-norm social cost objective for any $p \geq 1$ and that no deterministic SP mechanism can provide a better approximation ratio. They also provided a lower bound that holds for a broad class of randomized SP mechanisms. Feldman and Wilf [18] studied the case for $p = 2$ and established tight bounds for both deterministic and randomized mechanisms. Therefore, in this work, we consider the study of learning-augmented mechanism design under the L_p-norm social cost objective for $p \in (1, \infty)$ and investigate the extent to which predictions can overcome existing impossibility results. We aim to systematically analyze the consistency-robustness trade-off across a family of objectives ranging from social cost to maximum cost.

1.1 Our Contributions

We consider the study of learning-augmented mechanism design for facility location under the L_p-norm social cost objective in the standard setting of locating a single facility on the line. We examine the problem of designing SP mechanisms that, given a possibly imperfect prediction π of the optimal facility location, elicit location preferences from n agents truthfully and determine facility locations

to approximately minimize the L_p-norm social cost objective for $p \in (1, \infty)$ (when $p = 1$, the median mechanism is already SP and optimal [25]; when $p = +\infty$, the maximum cost objective in learning-augmented settings has been thoroughly studied [1, 4, 8]). The performance of a mechanism is measured by its *consistency* (approximation ratio when the prediction is perfect) and *robustness* (approximation ratio for an arbitrary inaccurate prediction). Our main technical contributions are summarized as follows.

Deterministic SP Mechanisms: Tight Trade-off. We propose a family of deterministic SP generalized median mechanisms parameterized by a fraction $t \in (0, 1)$ of phantom points placed at the prediction π. For the upper bound, we prove that this family achieves consistency of $\left(1 + (\frac{1-t}{1+t})^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$ and robustness of $\left(1 + (\frac{1+t}{1-t})^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$ for the L_p-norm social cost objective (Theorem 1). Our proof introduces a novel scaling argument to construct worst-case instances for both consistency and robustness. Also it is the first analysis of generalized median mechanisms under the L_p norm social cost objective, which contributes technically even without predictions. We denote these two expressions by $A_p(t)$ and $B_p(t)$, respectively. Furthermore, we establish a matching lower bound (Theorem 2), showing that no deterministic SP mechanism with the same consistency can achieve better robustness. Hence, our mechanism attains the optimal trade-off between these two competing measures. Compared with the prediction-free upper and lower bound results of $2^{1-1/p}$ from [17], predictions can help to overcome such impossibility results and any improvement in consistency is necessarily accompanied by a loss in robustness. Beyond the two endpoint guarantees, Section 3.3 shows that the mechanism degrades smoothly with the normalized prediction error η: its pointwise approximation ratio is at most $\min\{A_p(t) + \eta, B_p(t)\}$. Appendix C gives the full proof and corresponding expected guarantees for predictions that are approximately accurate with high probability. Figure 1 illustrates this trade-off of consistency/robustness for $p = 2$ and compares our results with the prediction-free results of [17].

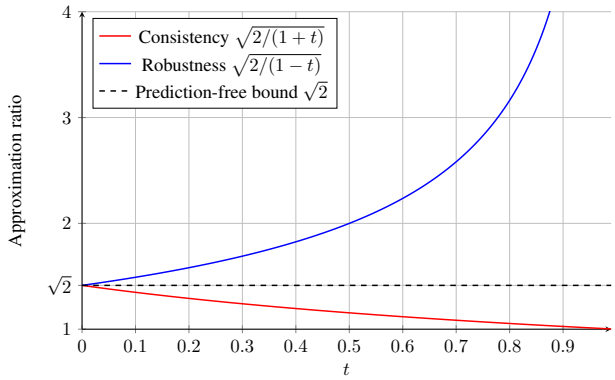


Figure 1: Trade-off between consistency (red) and robustness (blue) for $p = 2$.

Randomized SP Mechanisms: Upper and Lower Bounds. For upper bounds, we use the centroid random dictatorship (CRD) guarantee of Chan et al. [11] and mix CRD with our prediction-based deterministic mechanism. This gives an explicit finite-n randomized SP trade-off for every $p \in (1, +\infty)$ and places the formerly separate $p = 2$ and $n = 2$ mechanisms in one framework. The CRD component strictly improves the deterministic prediction-free ratio for every finite $p \geq 2$; for $1 < p < 2$, the improvement depends on n and, in the large-n limit, begins at approximately $p = 1.683$. For lower bounds, we establish a general bound $\beta(u) > 1$ on the robustness of any $(1+u)$-consistent mechanism (Theorem 4). While Balkanski et al. [4] gave a tight result for maximum cost ($p = +\infty$), their argument relies on the objective not being strictly convex and does not extend to our intermediate L_p setting. Our lower bound applies throughout $p \in (1, +\infty)$, although its limiting $p = +\infty$ specialization is weaker than their tight bound.

1.2 Additional Related Work

L_p-Norm Social Costs. We provide additional related work specific to the standard mechanism design setting for locating a single facility. For deterministic SP mechanisms, in addition to the results

of Feigenbaum et al. [17] on the line, existing studies have examined the L_p-norm social cost objective in higher-dimensional spaces. In d-dimensional Euclidean spaces, the coordinate-wise median (CM) mechanism, which locates the facility at the median of each coordinate, is SP and achieves a $\sqrt{d}$-approximation for the total-cost objective [23]. In the two-dimensional Euclidean space, Goel and Hann-Caruthers [19] showed that the CM mechanism achieves the best possible approximation ratio among all deterministic SP mechanisms. Recently, [11] studied the tight approximation ratios of CM and explored two randomized mechanisms. Our work sits at the intersection of these threads, being the first to study learning-augmented facility location under the L_p-norm social cost.

L_p Spaces. Separately, a distinct line of research considers d-dimensional L_p spaces where agent costs themselves are measured by L_p-norm distances. In this setting, Tang et al. [26] characterized the group strategyproof mechanisms, Lin [22] characterized two-agent strategyproof mechanisms, Chan et al. [8] established that CM is 2-approximate for the maximum cost in any two-dimensional L_p space, and Gravin and Jia [20] recently showed that CM achieves at most a 3-approximation for the total cost in any L_p space, improving upon the earlier $\sqrt{d}$ bound [23].

For randomized SP mechanisms, Procaccia and Tennenholtz [25] designed the LRM mechanism, which returns the leftmost agent, the rightmost agent, and their midpoint with probabilities $\frac{1}{4}$, $\frac{1}{4}$, and $\frac{1}{2}$, respectively; this mechanism is SP and achieves a 1.5-approximation for the maximum-cost objective. For $p = 2$, Feldman and Wilf [18] studied the mechanism that returns each agent's location with probability $\frac{1}{2n}$ and their average with probability $\frac{1}{2}$. More recently, Chan et al. [11] analyzed this centroid random dictatorship mechanism for every $p \geq 1$ and every n in arbitrary-dimensional Euclidean spaces. Their result applies, in particular, to the one-dimensional setting considered here.

Learning-Augmented Mechanism Design for Facility Location. Subsequent studies following the work of [1, 27] have focused on several directions: handling outlier predictions [6], studying obnoxious facility settings where agents prefer to be far from the facility [21, 15], examining the price of anarchy in metric spaces [12], and considering alternative success metrics such as the quality of recommendation [14]. Balkanski et al. [4] further investigated how the nature of the prediction affects the design of randomized mechanisms for the maximum cost objective. [3] analyzed the envy ratio objective, which is different from previous work. Last, [5, 9] studied the high dimensional spaces.

2 Preliminaries

For every finite $n \geq 1$, consider a set of agents $N = \{1, 2, \ldots, n\}$ located on the real line $\mathbb{R}$. The agents' locations (or preferences) are represented by a profile $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$. Without loss of generality, we assume $x_1 \leq x_2 \leq \ldots \leq x_n$. The distance between two points $x, y \in \mathbb{R}$ is denoted by $d(x, y)$. Given a deterministic facility location $y \in \mathbb{R}$, the cost incurred by agent $i \in N$ is the distance to the facility, i.e.,

$$\text{cost}(x_i, y) = d(x_i, y).$$

Our objective is to minimize the L_p-norm social cost for $p \geq 1$, defined for a deterministic facility location $y \in \mathbb{R}$ as

$$\text{SC}_p(\mathbf{x}, y) = \left(\sum_{i \in N} \text{cost}(x_i, y)^p \right)^{\frac{1}{p}} = \left(\sum_{i \in N} d(x_i, y)^p \right)^{\frac{1}{p}}.$$

This objective generalizes two fundamental criteria. When $p = 1$, it coincides with the total cost, capturing efficiency. As $p \to \infty$, it converges to the maximum cost, reflecting fairness or equity considerations [7]. For intermediate values $1 < p < \infty$, the L_p-norm provides a continuous spectrum of objectives that interpolate between efficiency and fairness. As p increases, the objective becomes increasingly sensitive to larger individual costs, implicitly assigning greater weight to agents who experience higher disutility, while still accounting for all agents. This makes the L_p framework a natural model for fairness-aware optimization.

In the learning-augmented setting, the mechanism additionally has access to a prediction $\pi \in \mathbb{R}$ of the optimal facility location. A mechanism is required to work uniformly over every finite input size and is formally a family $f = (f_n)_{n \geq 1}$. For each n, a deterministic mechanism $f_n : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ maps

a location profile $\mathbf{x}$ and a prediction π to a facility location $y = f_n(\mathbf{x}, \pi)$, whereas a randomized mechanism $f_n : \mathbb{R}^n \times \mathbb{R} \to \Delta(\mathbb{R})$ outputs a probability distribution over locations. Here, $\Delta(\mathbb{R})$ denotes the set of all probability distributions over $\mathbb{R}$. We suppress the subscript n when the input size is clear.

For a randomized mechanism, the expected cost of agent i is

$$\text{cost}(x_i, f(\mathbf{x}, \pi)) = \mathbb{E}_{y \sim f(\mathbf{x}, \pi)}[d(x_i, y)].$$

The social objective of a randomized mechanism is the expected realized L_p-norm social cost:

$$\text{SC}_p(\mathbf{x}, f(\mathbf{x}, \pi)) = \mathbb{E}_{y \sim f(\mathbf{x}, \pi)} \left[\text{SC}_p(\mathbf{x}, y)\right] = \mathbb{E}_{y \sim f(\mathbf{x}, \pi)} \left[\left(\sum_{i \in N} d(x_i, y)^p \right)^{\frac{1}{p}} \right].$$

That is, for each realized location in the lottery, we first compute its L_p social cost, and then take the expectation over the lottery. This is distinct from taking the L_p-norm of the agents' expected costs.

A mechanism is *unanimous* if it outputs the common location whenever all agents report the same location; that is, $f(\mathbf{x}, \pi) = c$ whenever $x_i = c$ for all $i \in N$. Note that unanimity rules out trivial mechanisms that completely ignore agent reports and always output the prediction.

A mechanism is *strategyproof* (SP) if no agent can reduce their cost by misreporting their location, regardless of the reports of other agents. Formally, a mechanism f is SP if, for every $n \geq 1$, profile $\mathbf{x} \in \mathbb{R}^n$, prediction π, agent $i \in N$, and alternative report $x_i' \in \mathbb{R}$,

$$\text{cost}(x_i, f(x_i', \mathbf{x}_{-i}, \pi)) \geq \text{cost}(x_i, f(\mathbf{x}, \pi)).$$

A mechanism is *group strategyproof* (GSP) if no coalition of agents can jointly misreport in a way that strictly benefits all of its members. Specifically, for every $n \geq 1$, every prediction $\pi \in \mathbb{R}$, all $\mathbf{x} \in \mathbb{R}^n$, all coalitions $S \subseteq N$, and all misreports $\mathbf{x}_S' \in \mathbb{R}^{|S|}$, there exists some agent $i \in S$ such that

$$\text{cost}(x_i, f(\mathbf{x}_S', \mathbf{x}_{-S}, \pi)) \geq \text{cost}(x_i, f(\mathbf{x}, \pi)).$$

Clearly, every GSP mechanism is also SP.

The performance of a learning-augmented mechanism is evaluated in terms of its *consistency* and *robustness*. Let y^* denote an optimal facility location that minimizes $\text{SC}_p(\mathbf{x}, \cdot)$ for a given profile $\mathbf{x}$. A mechanism f is α-*consistent* if it achieves an α-approximation uniformly over every finite n when the prediction is perfect, i.e., when $\pi = y^*$:

$$\forall n \geq 1,\ \forall \mathbf{x} \in \mathbb{R}^n : \quad \text{SC}_p(\mathbf{x}, f_n(\mathbf{x}, y^*)) \leq \alpha \cdot \text{SC}_p(\mathbf{x}, y^*).$$

The mechanism is β-*robust* if it guarantees a β-approximation uniformly over every finite n and arbitrary predictions:

$$\forall n \geq 1,\ \forall \mathbf{x} \in \mathbb{R}^n,\ \forall \pi \in \mathbb{R} : \quad \text{SC}_p(\mathbf{x}, f_n(\mathbf{x}, \pi)) \leq \beta \cdot \text{SC}_p(\mathbf{x}, y^*).$$

Equivalently, both approximation ratios are worst-case suprema over profiles of all finite sizes. For prediction-free mechanisms, the approximation ratio coincides with the robustness guarantee in the learning-augmented framework.

Any mechanism with bounded robustness must be unanimous. Indeed, if all agents are located at the same point a, then the optimal social cost is zero, and a bounded approximation ratio forces the mechanism to output a. Consequently, we restrict attention to unanimous mechanisms throughout this paper.

When $p = 1$, [25] shows that the median mechanism is GSP and optimal for minimizing social cost. When $p = \infty$, [1, 4, 8] provide tight bounds and complete characterizations for the maximum-cost objective. Accordingly, our analysis focuses on the intermediate regime $1 < p < \infty$.

3 Deterministic Mechanisms

In this section, we provide tight upper and lower bounds for deterministic mechanisms in Sections 3.1 and 3.2, respectively.

3.1 Upper Bound

Feigenbaum et al. [17] show that the median mechanism achieves a $2^{1-1/p}$-approximation for the L_p-norm social cost, and that this bound is tight among deterministic prediction-free SP mechanisms. Motivated by generalized-median characterizations for single-peaked domains [24], we design a mechanism in this class. A generalized median mechanism returns the median of the n reported locations together with at most $n-1$ fixed points. For convenience, when there are two median points, the mechanism always selects the left one.

Mechanism 1. *Given a profile $\mathbf{x}$, run the generalized median mechanism with $t \cdot n$ fixed points ($t \in [0, 1)$), all located at the prediction π. We choose t such that $t \cdot n$ is an integer.*

We discuss some special cases here. When $t \cdot n = n - 1$, this mechanism coincides with the MINMAXP mechanism [1], which outputs $y = \max(x_1, \min(x_n, \pi))$. When $t = 0$, it reduces to the standard median mechanism. The following theorem establishes the relationship between the parameter t and the resulting consistency and robustness guarantees.

Theorem 1. *Mechanism 1 is GSP, $\left(1 + \left(\frac{1-t}{1+t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$-consistent, and $\left(1 + \left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$-robust.*

Group strategyproofness follows directly from the properties of generalized median mechanisms. Therefore, we focus on analyzing the approximation ratios. Rather than presenting a single long proof, we introduce several auxiliary lemmas that facilitate the analysis. Specifically, we establish separate lemmas that will be used to bound consistency and robustness.

For convenience, given a profile $\mathbf{x}$ and the mechanism output y, we partition the agents into three subsets: $N = N_1 \cup N_2 \cup N_3$. An agent $i \in N_1$ satisfies $x_i < y$, an agent $i \in N_2$ satisfies $x_i = y$, and an agent $i \in N_3$ satisfies $x_i > y$. Let $n_1 = |N_1|$, $n_2 = |N_2|$, and $n_3 = |N_3|$. We will use the following observation.

Observation 1. *For an agent with $x_i \le y$ (respectively, $x_i \ge y$), moving it to any location $x_i' \le y$ (respectively, $x_i' \ge y$) will not change the mechanism's output y.*

Robustness Analysis

We first construct a worse-case instance for the robustness. Note that the maximum L_p-norm social cost for any $y \in [x_1, x_n]$ is attained at either $y = x_1$ or $y = x_n$, because $\mathrm{SC}_p^p(\mathbf{x}, y) = \sum_{i=1}^n |x_i - y|^p$ is a convex function of y when $p > 1$ (as each term $|x_i - y|^p$ is convex). Furthermore, the convexity guarantees that any stationary point satisfying $\frac{\partial \mathrm{SC}_p^p(\mathbf{x}, y)}{\partial y} = 0$ is the unique global minimizer. We begin with providing a basic lemma that simplifying the profile without improving the ratios.

Lemma 1. *If $n_1 > 1$ (respectively, $n_3 > 1$), then for any profile $\mathbf{x}$, there exists an agent $i \in N_1$ (respectively, $i \in N_3$) such that moving all agents in N_1 (respectively, N_3) to x_i, yielding a new profile $\mathbf{x}^{(i)}$, does not improve the approximation ratio.*

Next we show that we can further move all agents on one side to the mechanism output itself to construct the "worst" case.

Lemma 2. *Let $n_1, n_3 > 0$ and $n_2 \ge 0$. For any profile $\mathbf{x}$ with mechanism output $y = b$, suppose all agents in N_1 are located at $a < b$, all the agents in N_2 are located at b, and all agents in N_3 are located at $c > b$. Then there exists $k \in \{1, 3\}$ such that moving all agents in N_k to b does not improve the approximation ratio.*

The following characterizes the optimal solution and the optimal L_p social cost for the worst-case instance in Lemma 2.

Lemma 3. *For a profile with n_1 agents located at a and $n - n_1$ agents located at b, where $a < b$ and $0 < n_1 < n$, the optimal facility location and the optimal L_p social cost are given by*

$$y^* = \frac{a + rb}{1 + r}, \qquad \mathrm{SC}_p(\mathbf{x}, y^*) = (b - a) \cdot \left(\frac{n - n_1}{(1 + r)^{p-1}}\right)^{\frac{1}{p}}, \qquad \textit{where} \quad r = \left(\frac{n - n_1}{n_1}\right)^{\frac{1}{p-1}}.$$

Now we are ready to prove the robustness ratio.

Proposition 1. *Mechanism 1 is* $\left(1+\left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}$*-robust.*

Proof. We analyze two possible cases for the mechanism's output y.

Case 1: $y \in \{x_1, x_n\}$**.** By symmetry, assume $y = x_1$. Let $a = x_1$, and suppose n_1 agents are located at a. Since the mechanism is a generalized median mechanism with $t \cdot n$ fixed points at π, the number of agents at a satisfies $n_1 \geq \frac{n+t\cdot n}{2} - t \cdot n = \frac{(1-t)\cdot n}{2}$, regardless of the location of the prediction π.

By Lemma 1, the worst-case ratio occurs when all agents to the right of a are moved to a common location b, and by Observation 1 the mechanism's output remains unchanged. For this profile, Lemma 3 yields

$$\mathrm{ALG}^p = (n-n_1)\cdot(b-a)^p, \qquad \mathrm{OPT}^p = \frac{n-n_1}{(1+r)^{p-1}}\cdot(b-a)^p,$$

where $r = \left(\frac{n-n_1}{n_1}\right)^{\frac{1}{p-1}} = \left(\frac{n}{n_1}-1\right)^{\frac{1}{p-1}}$. Consequently,

$$\frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} = \frac{(n-n_1)\cdot(b-a)^p}{\frac{n-n_1}{(1+r)^{p-1}}\cdot(b-a)^p} = (1+r)^{p-1} = \left(1+\left(\frac{n}{n_1}-1\right)^{\frac{1}{p-1}}\right)^{p-1}.$$

This ratio is increasing as n_1 decreases. Substituting the lower bound $n_1 \geq \frac{(1-t)\cdot n}{2}$ gives

$$\frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} \leq \left(1+\left(\frac{2}{1-t}-1\right)^{\frac{1}{p-1}}\right)^{p-1} = \left(1+\left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{p-1}.$$

Taking the p-th root yields the desired bound on the L_p social cost:

$$\frac{\mathrm{ALG}}{\mathrm{OPT}} \leq \left(1+\left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}.$$

Case 2: $x_1 < y < x_n$**.** Let $b = y$, and denote by n_1, n_2, n_3 the numbers of agents to the left of b, at b, and to the right of b, respectively. By Lemma 1, there exist points $a < b$ and $c > b$ such that moving all agents in N_1 to a and all agents in N_3 to c does not improve the approximation ratio; Observation 1 ensures the output remains unchanged. Then, by Lemma 2, either moving all agents at a to b, or moving all agents at c to b, would not improve the approximation ratio, and again the output is unchanged. This reduces the profile to Case 1, and the same bound applies. □

Consistency Analysis

Next, we continue with the consistency guarantee. Again we first construct a worst-case profile for consistency. Given profile $\mathbf{x}$, let y be the mechanism output and y^* be the optimal solution.

Lemma 4. *If the mechanism outputs $y < y^*$, then moving any agent $i \in N_1$ (i.e., $x_i < y$) to location $x_i' = y$ does not improve the approximation ratio. Symmetrically, if $y > y^*$, then moving any agent $i \in N_3$ (i.e., $x_i > y$) to $x_i' = y$ does not improve the approximation ratio.*

Proposition 2. *Mechanism 1 is* $(1+(\frac{1-t}{1+t})^{\frac{1}{p-1}})^{\frac{p-1}{p}}$*-consistent.*

Proof. Recall that a mechanism is α-consistent if it achieves an α-approximation when the prediction is perfect, i.e., $\pi = y^*$. If the mechanism output y coincides with y^*, the approximation ratio is 1. By symmetry, we focus on the case where $y < y^*$.

Our goal is to construct the worst-case profile $\mathbf{x}$ that maximizes the ratio $\frac{\mathrm{SC}_p(\mathbf{x},y)}{\mathrm{SC}_p(\mathbf{x},y^*)}$. We proceed by simplifying the profile through several steps that do not improve the ratio:

1. *Left-side Concentration:* By Lemma 4, for any agent i with $x_i < y$, moving x_i to the location of y does not improve the approximation ratio. After repeatedly applying this, all agents in N_1 have been moved to y, namely, no agent is located to the left of y. Assume $y = 0$ without loss of generality.

2. *Right-side Concentration:* By Lemma 1, there exists a location $L > y$ such that moving all agents in N_3 to L does not improve the ratio.

After these steps, the profile is reduced to a two-point distribution: n_1 agents are located at 0 and $n_3 = n - n_1$ agents are located at L. Moreover, the output remains to be $y = 0$ due to Observation 1, which indicates that the number of agents at $y = 0$ is at least the median threshold, i.e., $n_1 \geq \frac{n+tn}{2} \geq \frac{1+t}{2}n$. For this two-point profile, the mechanism's cost satisfies $\mathrm{ALG}^p = n_3 \cdot L^p$. According to Lemma 3, the optimal social cost satisfies

$$\mathrm{OPT}^p = \frac{n_3}{(1+r)^{p-1}} L^p, \quad \text{where } r = \left(\frac{n_3}{n_1}\right)^{\frac{1}{p-1}}.$$

The approximation ratio is therefore

$$\frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} = \frac{n_3 L^p}{\frac{n_3}{(1+r)^{p-1}} L^p} = (1+r)^{p-1} = \left(1 + \left(\frac{n_3}{n_1}\right)^{\frac{1}{p-1}}\right)^{p-1}. \tag{1}$$

This ratio is monotonically increasing with respect to n_3/n_1. To maximize the ratio, we need to maximize n_3 (or equivalently, minimize n_1) subject to the constraint $n_1 \geq \frac{1+t}{2}n$. Setting n_1 to its lower bound $\frac{1+t}{2}n$, we have:

$$\frac{n_3}{n_1} \leq \frac{\frac{(1-t)n}{2}}{\frac{(1+t)n}{2}} = \frac{1-t}{1+t}.$$

Substituting this back into (1) gives

$$\frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} \leq \left(1 + \left(\frac{1-t}{1+t}\right)^{\frac{1}{p-1}}\right)^{p-1}.$$

Taking the p-th root, we obtain the consistency guarantee claimed. □

3.2 Lower Bound

We now turn to the lower bound. The following theorem establishes a tight lower bound.

Theorem 2. *Let $M = (M_n)_{n\geq 1}$ be a deterministic SP mechanism family. For the L_p-norm social cost with $p \in (1, \infty)$, if M is uniformly $(1 + (\frac{1-t}{1+t})^{\frac{1}{p-1}})^{\frac{p-1}{p}}$-consistent over all finite n for some $t \in [0, 1)$, then its uniform robustness is at least $(1 + (\frac{1+t}{1-t})^{\frac{1}{p-1}})^{\frac{p-1}{p}}$.*

Therefore, among all deterministic strategyproof mechanisms, Mechanism 1 achieves the optimal trade-off between consistency and robustness.

Proof sketch. Assume bounded robustness, which implies unanimity. Place a_n agents at 0 and $b_n = n - a_n$ at $L > 0$, choosing $a_n/n < (1+t)/2$ with $a_n/n \to (1+t)/2$, and provide the correct prediction. Strategyproofness and unanimity force the output y into $[0, L]$, while consistency excludes $y = 0$. Keeping the prediction fixed, move the agents at L to y one at a time; strategyproofness preserves the output. On the resulting two-cluster profile, Lemma 3 gives ratio $(1+(a_n/b_n)^{1/(p-1)})^{(p-1)/p}$, tending to $B_p(t)$. Since uniform robustness bounds every finite instance, it also bounds this limit. The full proof appears in Appendix A.

3.3 Intermediate Prediction Quality

The endpoint guarantees also yield a bound for imperfect but accurate predictions. Fix a profile with optimum y^* and positive optimal cost $\mathrm{OPT}_p(\mathbf{x}) = \mathrm{SC}_p(\mathbf{x}, y^*)$. Define the normalized error $\eta(\mathbf{x}, \pi) = n^{1/p}|\pi - y^*|/\mathrm{OPT}_p(\mathbf{x})$. Writing $M_t(\mathbf{x}, \pi)$ for the output of Mechanism 1, we obtain

$$\frac{\mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, \pi)\big)}{\mathrm{OPT}_p(\mathbf{x})} \leq \min\{A_p(t) + \eta(\mathbf{x}, \pi),\ B_p(t)\}.$$

Indeed, the mechanism is a fixed order statistic and hence 1-Lipschitz in the prediction, while Minkowski's inequality bounds the change in social cost by $n^{1/p}$ times the facility displacement. Comparing with the output under the correct prediction gives the first bound; robustness gives the second. Zero optimal cost is handled by unanimity. Appendix C provides the full proof and extensions to probabilistic prediction guarantees.

For a fixed profile, suppose an exogenous random prediction Π satisfies $\Pr[\eta(\mathbf{x}, \Pi) \le \varepsilon] \ge 1 - \delta$, where $\varepsilon \ge 0$ and $\delta \in [0, 1]$. Applying the error-dependent bound on this event and robustness otherwise gives

$$\frac{\mathbb{E}_\Pi[\mathrm{SC}_p(\mathbf{x}, M_t(\mathbf{x}, \Pi))]}{\mathrm{OPT}_p(\mathbf{x})} \le (1 - \delta) \min\{A_p(t) + \varepsilon, B_p(t)\} + \delta B_p(t).$$

In particular, $\varepsilon = 0$ covers predictions that are exactly correct with probability at least $1 - \delta$; the expectation here is over prediction uncertainty, not internal mechanism randomization.

Alternatively, if only the expected normalized error is known, $\mathbb{E}_\Pi[\eta(\mathbf{x}, \Pi)] \le \bar{\eta}$, the same pointwise bound yields $\mathbb{E}_\Pi[\mathrm{SC}_p(\mathbf{x}, M_t(\mathbf{x}, \Pi))]/\mathrm{OPT}_p(\mathbf{x}) \le \min\{A_p(t) + \bar{\eta}, B_p(t)\}$, without any further assumption on the prediction error distribution.

4 Randomized Mechanisms

In this section, we analyze randomized mechanisms. We first obtain a finite-n upper bound by mixing our deterministic prediction-based mechanism with centroid random dictatorship (CRD), and then establish a general lower bound on the trade-off between consistency and robustness for randomized SP mechanisms.

4.1 Upper Bound

We first record that randomized SP mechanisms can be mixed while preserving their guarantees.

Observation 2. *Suppose Mechanism A is strategyproof, α_1-consistent, and β_1-robust for the L_p social cost, Mechanism B is strategyproof, α_2-consistent, and β_2-robust, and Mechanism C runs A with probability λ and B with probability $1 - \lambda$. Then Mechanism C is also strategyproof, $(\lambda\alpha_1 + (1 - \lambda)\alpha_2)$-consistent, and $(\lambda\beta_1 + (1 - \lambda)\beta_2)$-robust.*

Mechanism 2. *(*CRD*) Let $\bar{x} = \frac{1}{n}\sum_{i \in N} x_i$. Return $y = \bar{x}$ with probability $\frac{1}{2}$, and return each $y = x_i$ with probability $\frac{1}{2n}$.*

The following prediction-free guarantee was recently established for arbitrary-dimensional Euclidean spaces and therefore applies to the line as a special case [11]. Let $q_n = 1 - 1/n$.

Lemma 5 ([11])**.** *For every $n \ge 2$, Mechanism 2 is SP and has approximation ratio at most*

$$\rho_{p,n} = \begin{cases} \dfrac{1}{2} q_n^{\frac{2}{p} - 1} \left(2^{\frac{2}{p} - 1} + 2^{\frac{1}{p}}\right), & 1 \le p \le 2, \\ \dfrac{1}{2} \left(2^{1 - \frac{2}{p}} q_n^{1 - \frac{2}{p}} + 2^{1 - \frac{1}{p}}\right), & 2 \le p \le +\infty. \end{cases}$$

Recall that the guarantees of Mechanism 1 are

$$A_p(t) = \left(1 + \left(\frac{1 - t}{1 + t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}, \qquad B_p(t) = \left(1 + \left(\frac{1 + t}{1 - t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}.$$

Mechanism 3. *For a fixed $n \ge 2$, run Mechanism 1 with probability $\lambda \in [0, 1]$ and Mechanism 2 with probability $1 - \lambda$.*

To retain the finite-size improvement in Lemma 5, we state the guarantee for each n; "size-n" below means that the worst case is restricted to n-agent profiles.

Theorem 3. *Fix $n \ge 2$, $p \in (1, +\infty)$, and an admissible $t \in [0, 1)$. Mechanism 3 is SP and, on n-agent instances, has size-n consistency and robustness guarantees given by*

$$\alpha_{p,n}(t, \lambda) = \lambda A_p(t) + (1 - \lambda)\rho_{p,n}, \qquad \beta_{p,n}(t, \lambda) = \lambda B_p(t) + (1 - \lambda)\rho_{p,n}.$$

Proof. Mechanism 1 has guarantees $A_p(t)$ and $B_p(t)$ by Theorem 1, while Mechanism 2 ignores the prediction and is $\rho_{p,n}$-approximate by Lemma 5. The claim follows from Observation 2. □

This theorem unifies the earlier special cases. When $p = 2$, CRD is the mechanism studied by Feldman and Wilf [18], and the cited L_2 analysis gives $\rho_{2,n} = (1 + \sqrt{2})/2$ [17, 11]. Thus the upper-bound pair becomes

$$\alpha_{2,n}(t, \lambda) = \lambda\sqrt{\frac{2}{1+t}} + (1-\lambda)\frac{1+\sqrt{2}}{2}, \qquad \beta_{2,n}(t, \lambda) = \lambda\sqrt{\frac{2}{1-t}} + (1-\lambda)\frac{1+\sqrt{2}}{2}.$$

Together with the $p = 2$ specialization of Theorem 4 below, this gives explicit upper and lower trade-off bounds, although a gap remains.

When $n = 2$, CRD selects the midpoint with probability $1/2$ and each report with probability $1/4$, so it is exactly the LRM mechanism [25], with $\rho_{p,2} = (1 + 2^{1-1/p})/2$. In particular, mixing it with the two-agent MinMaxP mechanism recovers the upper-bound line $\alpha + \beta = 1 + 2^{1-1/p}$.

Finally, compared with the deterministic prediction-free ratio $D_p = 2^{1-1/p}$, CRD satisfies $\rho_{p,n} < D_p$ for every finite $p \geq 2$. For $1 < p < 2$, the comparison depends on n; in the large-n limit, the crossover is the solution of $\frac{1}{2}(2^{2/p-1} + 2^{1/p}) = 2^{1-1/p}$, namely $p \approx 1.683$. We do not expect the resulting randomized upper bound to match Theorem 4 in general.

4.2 Lower Bound

We next establish a lower bound on the trade-off between consistency and robustness for randomized SP mechanisms.

Theorem 4. *Let u be a parameter so that $0 \leq u \leq \left(\frac{3^p+1}{2^{p+1}}\right)^{\frac{1}{p}} - 1$. Any randomized SP mechanism that achieves $(1+u)$-consistency must have robustness at least*

$$\beta \geq 2^{1-\frac{1}{p}}\left((1-T)^p + T^p\right)^{\frac{1}{p}},$$

where $T \in [0, 1]$ is the unique solution to $(1+T)^p + (1-T)^p = 2(1+u)^p$.

Two familiar objectives provide useful comparisons. For $p = +\infty$, taking the limit gives $T = u$, so our lower bound is $2 - 2u$ for $u \in [0, 1/2]$. This is weaker than the tight lower bound $2 - u$ of Balkanski et al. [4], as expected from our argument's greater generality. For $p = 2$, we have $T = \sqrt{u^2 + 2u}$ and hence

$$\beta \geq \sqrt{2}\sqrt{1 + 4u + 2u^2 - 2\sqrt{u^2+2u}}, \qquad 0 \leq u \leq \frac{\sqrt{5}}{2} - 1.$$

The randomized upper and lower bounds do not coincide. Our results are limited to a single facility on the line; closing this gap and extending the analysis to richer metric spaces remain open.

5 Conclusions

In this paper, we studied learning-augmented facility location mechanisms under the L_p-norm social cost on the line. For deterministic strategyproof mechanisms, we proposed a family of generalized median rules parameterized by the weight assigned to the prediction, and established a tight trade-off between consistency and robustness. For randomized strategyproof mechanisms, we combined our deterministic mechanism with CRD to obtain a finite-n upper bound for every $p \in (1, +\infty)$, unifying the squared-cost and two-agent cases, and derived a general lower bound on the robustness of any $(1+u)$-consistent rule. A gap remains between the randomized upper and lower bounds. Closing this gap, and extending our results beyond the line to tree metrics, higher-dimensional Euclidean spaces, or arbitrary graphs, are interesting directions for future work. In particular, it would be valuable to determine whether randomization improves the asymptotic order of the trade-off as consistency approaches one, or only its constants. Another direction is to select the trust parameter using externally calibrated prediction-quality estimates while preserving strategyproofness.

Acknowledgments and Disclosure of Funding

Hau Chan is supported by the National Science Foundation under grants IIS:RI #2302999 and IIS:RI #2414554. The work is supported in part by the Guangdong Provincial/Zhuhai Key Laboratory of IRADS (2022B1212010006), by AIRH, BNBU (2020KSYS007), and by ISC of BNU-Zhuhai. Chenhao Wang is supported by the Guangdong Provincial General Grant of 2026A1515010196, and by BNBU under grant R0400004-24B, R0800018-25. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies. There is no competing-interests disclosure.

References

[1] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. Learning-augmented mechanism design: leveraging predictions for facility location. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 497–528, 2022.

[2] Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.

[3] Haris Aziz, Yuhang Guo, Alexander Lam, and Houyu Zhou. Learning-augmented facility location mechanisms for envy ratio. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2025.

[4] Eric Balkanski, Vasilis Gkatzelis, and Golnoosh Shahkarami. Randomized strategic facility location with predictions. *Advances in Neural Information Processing Systems*, 37:35639–35664, 2024.

[5] Zohar Barak. Facility location mechanism design: Breaking the deterministic barrier. *arXiv preprint arXiv:2605.24750*, 2026.

[6] Zohar Barak, Anupam Gupta, and Inbal Talgam-Cohen. MAC advice for facility location mechanism design. In *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[7] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problems: a survey. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4356–4365, 2021.

[8] Hau Chan, Jianan Lin, and Chenhao Wang. Strategyproof mechanisms for facility location with prediction under the maximum cost objective. In *ECAI 2025 - 28th European Conference on Artificial Intelligence*, pages 1318–1325, 2025.

[9] Hau Chan, Jianan Lin, and Chenhao Wang. Learning-augmented strategyproof facility location in $\mathbb{R}^d$ with ℓ_p distances. *arXiv preprint arXiv:2609.12829*, 2026.

[10] Hau Chan, Jianan Lin, and Chenhao Wang. Obnoxious facility location problems: Strategyproof mechanisms optimizing l_p-aggregated utilities and costs. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems*, pages 496–504, 2026.

[11] Hau Chan, Jianan Lin, and Chenhao Wang. Strategyproof mechanisms for euclidean facility location problems under L_p-norm social cost. *arXiv preprint arXiv:2606.08621*, 2026.

[12] Qingyun Chen, Nick Gravin, and Sungjin Im. Strategic facility location via predictions. In *Proceedings of the 20th International Conference of Web and Internet Economics (WINE)*, 2024.

[13] Yukun Cheng, Wei Yu, and Guochuan Zhang. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science*, 497:154–163, 2013.

[14] George Christodoulou, Alkmini Sgouritsa, and Ioannis Vlachos. Mechanism design augmented with output advice. In *Advances in Neural Information Processing Systems: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[15] Jiazhu Fang, Qizhi Fang, Wenjing Liu, and Qingqin Nong. Mechanism design with predictions for facility location games with candidate locations. In *Theory and Applications of Models of Computation - 18th Annual Conference, (TAMC)*, volume 14637, pages 38–49, 2024.

[16] Itai Feigenbaum and Jay Sethuraman. Strategyproof mechanisms for one-dimensional hybrid and obnoxious facility location models. In *AAAI workshop: Incentive and trust in E-communities*, volume WS-15-08, pages 8–13, 2015.

[17] Itai Feigenbaum, Jay Sethuraman, and Chun Ye. Approximately optimal mechanisms for strategyproof facility location: minimizing L_p norm of costs. *Mathematics of Operations Research*, 42(2):434–447, 2017.

[18] Michal Feldman and Yoav Wilf. Strategyproof facility location and the least squares objective. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 873–890, 2013.

[19] Sumit Goel and Wade Hann-Caruthers. Optimality of the coordinate-wise median mechanism for strategyproof facility location in two dimensions. *Social Choice and Welfare*, 61(1):11–34, 2023.

[20] Nikolai Gravin and Jianhao Jia. Approximation guarantees of median mechanism in $\mathbb{R}^d$. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 495–506, 2025.

[21] Gabriel Istrate and Cosmin Bonchis. Mechanism design with predictions for obnoxious facility location. *CoRR*, abs/2212.09521, 2022.

[22] Jianan Lin. Nearly complete characterization of 2-agent deterministic strategyproof mechanisms for single facility location in l_p space. In *International Conference on Combinatorial Optimization and Applications*, pages 411–425. Springer, 2020.

[23] Reshef Meir. Strategyproof facility location for three agents on a circle. In *Algorithmic Game Theory - 12th International Symposium (SAGT)*, volume 11801 of *Lecture Notes in Computer Science*, pages 18–33. Springer, 2019.

[24] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[25] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)*, 1(4):18:1–18:26, 2013.

[26] Pingzhong Tang, Dingli Yu, and Shengyu Zhao. Characterization of group-strategyproof mechanisms for facility location in strictly convex space. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 133–157, 2020.

[27] Chenyang Xu and Pinyan Lu. Mechanism design with predictions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 571–577, 2022.

[28] Deshi Ye, Lili Mei, and Yong Zhang. Strategy-proof mechanism for obnoxious facility location on a line. In *International computing and combinatorics conference*, pages 45–56. Springer, 2015.

[29] Houyu Zhou, Minming Li, and Hau Chan. Strategyproof mechanisms for group-fair facility location problems. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 613–619, 07 2022.

A Omitted Proofs for Section 3

Proof of Theorem 1

Proof of Lemma 1. By symmetry, it suffices to consider the case $n_1 > 1$. We prove the claim by contradiction. Suppose that for a given profile $\mathbf{x}$, no such agent $i \in N_1$ exists. For each agent $j \in N_1$, let $\mathbf{x}^{(j)}$ denote the profile obtained by moving all agents in N_1 to location x_j. Let ALG_j and OPT_j denote the mechanism's social cost and the optimal social cost for profile $\mathbf{x}^{(j)}$, respectively. For the original profile $\mathbf{x}$, we use ALG and OPT. We continue to denote the mechanism's output by y.

By the assumption of contradiction, for every $j \in N_1$, the approximation ratio for $\mathbf{x}^{(j)}$ is strictly better than that for $\mathbf{x}$. Formally,

$$\forall j \in N_1, \quad \frac{\mathrm{ALG}_j^p}{\mathrm{OPT}_j^p} < \frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} \quad \Longleftrightarrow \quad \mathrm{OPT}_j^p > \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \mathrm{ALG}_j^p.$$

Summing over all $j \in N_1$, we obtain

$$\begin{aligned}
\sum_{j \in N_1} \mathrm{OPT}_j^p &> \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \sum_{j \in N_1} \mathrm{ALG}_j^p \\
&= \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \sum_{j \in N_1} \left(n_1 \cdot (y - x_j)^p + \sum_{i \in N_3} (x_i - y)^p \right) \\
&= \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \left(n_1 \cdot \sum_{j \in N_1} (y - x_j)^p + n_1 \cdot \sum_{i \in N_3} (x_i - y)^p \right) \\
&= \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot n_1 \cdot \mathrm{ALG}^p = n_1 \cdot \mathrm{OPT}^p.
\end{aligned}$$

By the definition of OPT_j, for profile $\mathbf{x}^{(j)}$, any facility location $y' \in \mathbb{R}$ yields a social cost at least OPT_j. That is,

$$\forall j \in N_1, \ \forall y' \in \mathbb{R}, \quad \mathrm{OPT}_j^p \leq n_1 \cdot |y' - x_j|^p + \sum_{i \in N_2 \cup N_3} |y' - x_i|^p.$$

Summing over all $j \in N_1$ gives

$$\begin{aligned}
\forall y' \in \mathbb{R}, \quad \sum_{j \in N_1} \mathrm{OPT}_j^p &\leq \sum_{j \in N_1} \left(n_1 \cdot |y' - x_j|^p + \sum_{i \in N_2 \cup N_3} |y' - x_i|^p \right) \\
&= n_1 \cdot \sum_{j \in N_1} |y' - x_j|^p + n_1 \cdot \sum_{i \in N_2 \cup N_3} |y' - x_i|^p \\
&= n_1 \cdot \sum_{i \in N} |y' - x_i|^p.
\end{aligned}$$

Taking the minimum over $y' \in \mathbb{R}$, we obtain

$$\sum_{j \in N_1} \mathrm{OPT}_j^p \leq n_1 \cdot \mathrm{OPT}^p,$$

which contradicts the earlier inequality. This completes the proof. □

Proof of Lemma 2. We again proceed by contradiction. Assume, for the sake of contradiction, that moving the agents in N_1 to $y = b$ and moving the agents in N_3 to $y = b$ both strictly improve the approximation ratio. Denote the resulting profiles by $\mathbf{x}^{(a)}$ and $\mathbf{x}^{(c)}$, respectively. Let ALG_a, ALG_c and OPT_a, OPT_c denote the mechanism's social cost and the optimal social cost for profiles $\mathbf{x}^{(a)}$ and $\mathbf{x}^{(c)}$, respectively.

By the assumption,

$$\forall j \in \{a, c\}, \quad \frac{\mathrm{ALG}_j^p}{\mathrm{OPT}_j^p} < \frac{\mathrm{ALG}^p}{\mathrm{OPT}^p} \quad \Longleftrightarrow \quad \mathrm{OPT}_j^p > \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \mathrm{ALG}_j^p.$$

Summing over $j \in \{a, c\}$ yields

$$\begin{aligned}
\mathrm{OPT}_a^p + \mathrm{OPT}_c^p &> \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot (\mathrm{ALG}_a^p + \mathrm{ALG}_c^p) \\
&= \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \left(\sum_{i \in N_1} (b-a)^p + \sum_{i \in N_3} (c-b)^p \right) \\
&= \frac{\mathrm{OPT}^p}{\mathrm{ALG}^p} \cdot \mathrm{ALG}^p = \mathrm{OPT}^p.
\end{aligned}$$

It is clear that the optimal facility location y^* for the original profile $\mathbf{x}$ lies in the interval $[a, c]$. We consider three cases.

Case 1: $y^* = b$**.** In this case,

$$\mathrm{OPT}^p = n_1 \cdot (b-a)^p + n_3 \cdot (c-b)^p.$$

Observe that

$$\mathrm{OPT}_a^p \leq \mathrm{SC}_p^p \left(\mathbf{x}^{(a)}, b \right) = n_3 \cdot (c-b)^p$$

and

$$\mathrm{OPT}_c^p \leq \mathrm{SC}_p^p \left(\mathbf{x}^{(c)}, b \right) = n_1 \cdot (b-a)^p.$$

Consequently,

$$\mathrm{OPT}_a^p + \mathrm{OPT}_c^p \leq \mathrm{OPT}^p,$$

which is a contradiction.

Case 2: $y^* < b$**.** Here,

$$\mathrm{OPT}^p = n_1 \cdot (y^* - a)^p + n_2 \cdot (b - y^*)^p + n_3 \cdot (c - y^*)^p.$$

We have the upper bounds

$$\mathrm{OPT}_a^p \leq \mathrm{SC}_p^p \left(\mathbf{x}^{(a)}, b \right) = n_3 \cdot (c-b)^p$$

and

$$\mathrm{OPT}_c^p \leq \mathrm{SC}_p^p \left(\mathbf{x}^{(c)}, y^* \right) = n_1 \cdot (y^* - a)^p + (n_2 + n_3) \cdot (b - y^*)^p.$$

Summing these inequalities gives

$$\begin{aligned}
\mathrm{OPT}_a^p + \mathrm{OPT}_c^p &\leq n_1 \cdot (y^* - a)^p + (n_2 + n_3) \cdot (b - y^*)^p + n_3 \cdot (c-b)^p \\
&= n_1 \cdot (y^* - a)^p + n_2 \cdot (b - y^*)^p + n_3 \cdot ((b - y^*)^p + (c-b)^p) \\
&\leq n_1 \cdot (y^* - a)^p + n_2 \cdot (b - y^*)^p + n_3 \cdot (c - y^*)^p \\
&= \mathrm{OPT}^p.
\end{aligned}$$

This again contradicts the earlier inequality.

Case 3: $y^* > b$**.** This case is symmetric to Case 2.

Since all possible cases lead to a contradiction, the lemma holds. □

Proof of Lemma 3. It is clear that the optimal solution y^* must lie in the interval $[a, b]$. Thus, it suffices to minimize SC_p^p. For a given location $y \in [a, b]$, we have

$$\mathrm{SC}_p^p(\mathbf{x}, y) = n_1 \cdot (y-a)^p + (n - n_1) \cdot (b-y)^p.$$

Taking the derivative with respect to y, we obtain

$$\frac{\partial \mathrm{SC}_p^p(\mathbf{x}, y)}{\partial y} = p \left(n_1 \cdot (y-a)^{p-1} - (n - n_1) \cdot (b-y)^{p-1} \right).$$

Since SC_p^p is convex with respect to y, its minimum is achieved when the derivative equals zero. Therefore, the optimal solution y^* satisfies

$$n_1 \cdot (y^* - a)^{p-1} = (n - n_1) \cdot (b - y^*)^{p-1}.$$

Rearranging terms yields

$$\frac{y^* - a}{b - y^*} = \left(\frac{n - n_1}{n_1}\right)^{\frac{1}{p-1}} = r,$$

which implies

$$y^* = \frac{a + rb}{1 + r}.$$

We now compute the optimal social cost. Using the expression for y^*, we have

$$y^* - a = \frac{r}{1+r}(b-a), \qquad b - y^* = \frac{1}{1+r}(b-a).$$

Substituting into SC_p^p, we obtain

$$\begin{aligned}\mathrm{SC}_p^p(\mathbf{x}, y^*) &= n_1 \left(\frac{r}{1+r}(b-a)\right)^p + (n - n_1)\left(\frac{1}{1+r}(b-a)\right)^p \\ &= \frac{(b-a)^p}{(1+r)^p}\left(n_1 r^p + (n - n_1)\right).\end{aligned}$$

Since $r^{p-1} = \frac{n-n_1}{n_1}$, we have $n_1 r^p = (n - n_1)r$, and thus

$$n_1 r^p + (n - n_1) = (n - n_1)(r + 1).$$

Therefore,

$$\mathrm{SC}_p^p(\mathbf{x}, y^*) = \frac{(b-a)^p}{(1+r)^{p-1}} \cdot (n - n_1),$$

which implies

$$\mathrm{SC}_p(\mathbf{x}, y^*) = (b-a) \cdot \left(\frac{n - n_1}{(1+r)^{p-1}}\right)^{\frac{1}{p}}.$$

This completes the proof. □

Proof of Lemma 4. By symmetry, it suffices to prove the case $y < y^*$; the case $y^* < y$ follows analogously by considering an agent in N_3 moving leftward to y.

Fix an agent $i \in N_1$, so $x_i < y$. For $s \in [x_i, y]$, let $\mathbf{x}(s)$ denote the profile obtained from $\mathbf{x}$ by replacing x_i with s. Define

$$F_s(z) := \sum_{j \neq i} |x_j - z|^p + |s - z|^p,$$

and let

$$z^*(s) := \arg\min_{z \in \mathbb{R}} F_s(z).$$

Since $1 < p < \infty$, the function $F_s(z)$ is strictly convex in z, and hence $z^*(s)$ is uniquely defined for every $s \in [x_i, y]$.

We first show that $z^*(s) > y$ for all $s \in [x_i, y]$. Since $z^*(x_i) = y^* > y$, it is enough to show that $F_s'(y) < 0$ throughout the motion. For any $s < y$,

$$F_s'(y) = p \sum_{\substack{j \neq i \\ x_j < y}} (y - x_j)^{p-1} + p(y - s)^{p-1} - p \sum_{x_j > y} (x_j - y)^{p-1}.$$

As s increases, the term $p(y-s)^{p-1}$ weakly decreases, and therefore

$$F_s'(y) \leq F_{x_i}'(y).$$

Because the unique minimizer of F_{x_i} is $y^* > y$, the function F_{x_i} is still decreasing at y, so $F_{x_i}'(y) < 0$. Hence $F_s'(y) < 0$ for every $s \in [x_i, y]$, which implies $z^*(s) > y$.

Therefore, throughout the motion, the prediction $\pi = z^*(s)$ remains to the right of y. Since the mechanism is the generalized median of the n reported locations together with tn fixed points at π, moving one report from $x_i < y$ to any $s \in [x_i, y]$ does not change the relevant order statistic. Thus, the mechanism output remains equal to y for all $s \in [x_i, y]$.

Now define

$$A(s) := \mathrm{SC}_p^p(\mathbf{x}(s), y) = F_s(y), \qquad B(s) := \mathrm{SC}_p^p(\mathbf{x}(s), z^*(s)) = F_s(z^*(s)).$$

We analyze the quantity

$$R(s) := \frac{A(s)}{B(s)},$$

which is the p-th power of the approximation ratio; monotonicity of $R(s)$ is equivalent to monotonicity of the approximation ratio itself.

Since y is fixed and $s < y$,

$$A'(s) = \frac{d}{ds} F_s(y) = -p(y - s)^{p-1}.$$

For $B(s)$, the envelope theorem yields

$$B'(s) = \frac{\partial}{\partial s} F_s(z^*(s)) = -p(z^*(s) - s)^{p-1},$$

because $\frac{\partial}{\partial z} F_s(z^*(s)) = 0$ at the minimizer and $z^*(s) > y \geq s$.

Hence

$$R'(s) = \frac{A'(s)B(s) - A(s)B'(s)}{B(s)^2},$$

and the numerator is

$$\begin{aligned} A'(s)B(s) - A(s)B'(s) &= -p(y - s)^{p-1} B(s) + p(z^*(s) - s)^{p-1} A(s) \\ &= p\Big(A(s)(z^*(s) - s)^{p-1} - B(s)(y - s)^{p-1}\Big). \end{aligned}$$

Since y is not optimal, we have $A(s) > B(s)$. Moreover, because $z^*(s) > y \geq s$,

$$(z^*(s) - s)^{p-1} > (y - s)^{p-1} > 0.$$

It follows that

$$A'(s)B(s) - A(s)B'(s) > 0,$$

and therefore $R'(s) > 0$ for all $s \in [x_i, y)$.

Thus $R(s)$ is monotonically increasing as the agent moves rightward from x_i to y. In particular, replacing x_i by y cannot decrease the approximation ratio. Equivalently, moving any agent in N_1 to y does not improve the approximation ratio. □

Proof of Theorem 2

Proof. Let

$$\alpha(t) = \left(1 + \left(\frac{1 - t}{1 + t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}.$$

If M has unbounded robustness, the claim is immediate. Hence, assume its robustness is bounded; as noted in Section 2, every M_n is then unanimous.

For each integer n, let

$$a_n = \left\lceil \frac{1 + t}{2} n \right\rceil - 1 \qquad \text{and} \qquad b_n = n - a_n.$$

Since $t < 1$, there exists n_0 such that $1 \leq a_n \leq n - 1$ for every $n \geq n_0$. Fix any such $n \geq n_0$. Consider the profile $\mathbf{x}$ with a_n agents at 0 and b_n agents at $L > 0$, and let the prediction be accurate, i.e., $\pi = y^*$. Denote the output of M_n by y.

We first claim that $y \in [0, L]$. If $y > L$, iteratively moving every agent at 0 to L does not change the output. Indeed, in each step, applying SP once with true location 0 and once with true location L forces the two outcomes to coincide. The resulting unanimous profile has all agents at L but still outputs $y > L$, a contradiction. The symmetric argument applies if $y < 0$: move every agent at L to 0 to obtain a unanimous profile that still outputs $y < 0$. Thus, $y \in [0, L]$.

Moreover, $y \neq 0$. By the definition of a_n, we have

$$a_n < \frac{1+t}{2}n, \qquad \text{and hence} \qquad \frac{b_n}{a_n} > \frac{1-t}{1+t}.$$

If $y = 0$, Lemma 3 gives

$$\begin{aligned}\left(\frac{\mathrm{SC}_p(\mathbf{x}, 0)}{\mathrm{SC}_p(\mathbf{x}, y^*)}\right)^p &= \left(1 + \left(\frac{b_n}{a_n}\right)^{\frac{1}{p-1}}\right)^{p-1} \\ &> \left(1 + \left(\frac{1-t}{1+t}\right)^{\frac{1}{p-1}}\right)^{p-1} = \alpha(t)^p,\end{aligned}$$

contradicting uniform $\alpha(t)$-consistency. Therefore, $y > 0$.

Construct a second profile $\mathbf{x}'$ by moving, one at a time, all b_n agents at L to y, while keeping the prediction fixed. The same pair of SP inequalities shows that the output remains y after every move. Hence, on $\mathbf{x}'$, there are a_n agents at 0 and b_n agents at y, and M_n outputs y. Let y'^* be the optimal location for $\mathbf{x}'$. Lemma 3 therefore yields the finite-n witness

$$\frac{\mathrm{SC}_p(\mathbf{x}', y)}{\mathrm{SC}_p(\mathbf{x}', y'^*)} = \left(1 + \left(\frac{a_n}{b_n}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}.$$

Since $a_n/n \to (1+t)/2$ and $b_n/n \to (1-t)/2$, the uniform robustness β of M, which is a supremum over all finite n, satisfies

$$\begin{aligned}\beta &\geq \sup_{n \geq n_0} \left(1 + \left(\frac{a_n}{b_n}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}} \\ &\geq \lim_{n\to\infty} \left(1 + \left(\frac{a_n}{b_n}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}} = \left(1 + \left(\frac{1+t}{1-t}\right)^{\frac{1}{p-1}}\right)^{\frac{p-1}{p}}.\end{aligned}$$

This is an exact lower bound on the uniform worst-case ratio; the sequence of finite instances need not attain its supremum at any fixed n. □

B Omitted Proofs for Section 4

Proof of Theorem 4

The proof proceeds through some auxiliary lemmas below on two-agent instances.

Lemma 6. *For $p \geq 1$,*

$$2^p \geq 2 \cdot \left(\frac{2^{1-\frac{1}{p}} + 1}{2}\right)^p.$$

Proof. Both sides of the inequality are positive, so taking the p-th root (a strictly increasing operation) yields the equivalent inequality

$$2 \geq 1 + 2^{\frac{1}{p}-1} = 2^{\frac{1}{p}} \cdot \frac{2^{1-\frac{1}{p}} + 1}{2}.$$

Equality occurs precisely when $\frac{1}{p} - 1 = 0$, that is, when $p = 1$. □

The following lemma uses the consistency guarantee to bound the *expected absolute deviation* of the output from the perfect prediction in a symmetric two-agent instance.

Lemma 7. *Let $u \in \left[0, \frac{1}{2}\left(2^{1-1/p} - 1\right)\right)$ and let f be a randomized strategyproof mechanism with $(1+u)$-consistency and bounded robustness. For any two-agent profile (x_1, x_2) with $x_1 \le x_2$, let the prediction be the optimal location $\pi = y^* = \frac{x_1+x_2}{2}$. Then the expected absolute deviation satisfies*

$$d := \mathbb{E}\big[|y - \pi|\big] \le \frac{\Delta}{2} \cdot T,$$

where $\Delta = x_2 - x_1$ and $T \in \left[0, \frac{1}{2}\right]$ is the unique solution to $(1+T)^p + (1-T)^p = 2(1+u)^p$.

Proof. Let $\Delta = x_2 - x_1$. If $x_1 = x_2$, then the unanimity forces $y = \pi$ with probability 1, so $d = 0$ and the inequality holds trivially. Assume now $x_1 < x_2$.

For any facility location $y \in \mathbb{R}$, define $h(y) = \big(|y - x_1|^p + |y - x_2|^p\big)^{1/p}$, which is the L_p social cost of the two-agent profile. The function h is convex and, on the interval $[x_1, x_2]$, it is symmetric around π and strictly increasing in $[\pi, y]$. Moreover, for any $y \in [x_1, x_2]$ we can write

$$h(y) = \psi\big(|y - \pi|\big),$$

where $\psi : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ is defined by

$$\psi(z) = \left(\left|\frac{\Delta}{2} + z\right|^p + \left|\frac{\Delta}{2} - z\right|^p\right)^{\frac{1}{p}}.$$

Since h is convex and symmetric, ψ is convex and strictly increasing on $\mathbb{R}_{\ge 0}$.

Because the mechanism is $(1+u)$-consistent and $\pi = y^*$, we have

$$\mathbb{E}\big[h(y)\big] \le (1+u) \cdot h(\pi) = (1+u) \cdot 2^{\frac{1-p}{p}} \Delta.$$

We may apply Jensen's inequality to the convex function ψ:

$$\mathbb{E}\big[h(y)\big] = \mathbb{E}\big[\psi(|y - \pi|)\big] \ge \psi\big(\mathbb{E}[|y - \pi|]\big) = \psi(d).$$

Combining the two inequalities yields

$$\psi(d) \le (1+u) \cdot 2^{\frac{1-p}{p}} \Delta.$$

Dividing both sides by $\frac{\Delta}{2}$ and using the explicit form of ψ, we obtain

$$\left(\left|1 + \frac{2d}{\Delta}\right|^p + \left|1 - \frac{2d}{\Delta}\right|^p\right)^{\frac{1}{p}} \le 2^{\frac{1}{p}}(1+u).$$

Raising to the p-th power and dividing by 2 gives

$$\frac{1}{2}\left(\left|1 + \frac{2d}{\Delta}\right|^p + \left|1 - \frac{2d}{\Delta}\right|^p\right) \le (1+u)^p.$$

Define $t := \frac{2d}{\Delta} \in \mathbb{R}_{\ge 0}$. Then the left-hand side is $\frac{1}{2}\big(|1+t|^p + |1-t|^p\big)$. Notice that the function $g(t) = |1+t|^p + |1-t|^p$ is strictly increasing on $\mathbb{R}_{\ge 0}$. According to Lemma 6, $g(1) = 2^p \ge 2(1+u)^p$ always holds. Therefore, t must be in $[0, 1]$ and we have

$$\frac{1}{2}\big((1+t)^p + (1-t)^p\big) \le (1+u)^p.$$

Let $T \in [0, 1]$ be the unique solution to $(1+T)^p + (1-T)^p = 2(1+u)^p$. Then the inequality forces $t \le T$, which is equivalent to $d \le \frac{\Delta}{2} \cdot T$. □

We are now ready to prove the main lower bound theorem.

Proof of Theorem 4. Let f be any randomized strategyproof mechanism that achieves $(1+u)$-consistency, where $0 \le u \le \left(\frac{3^p+1}{2^{p+1}}\right)^{\frac{1}{p}} - 1$. We first show that

$$\left(\frac{3^p+1}{2^{p+1}}\right)^{\frac{1}{p}} - 1 \le \frac{2^{1-\frac{1}{p}}-1}{2}.$$

Note that it is equivalent to prove $(3^p+1)^{1/p} \le 2 + 2^{1/p}$. For $p \ge 1$, Minkowski's inequality in $\mathbb{R}^2$ states that for any vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^2$, $\|\mathbf{x}+\mathbf{y}\|_p \le \|\mathbf{x}\|_p + \|\mathbf{y}\|_p$. Choose $\mathbf{x} = (2,0)$ and $\mathbf{y} = (1,1)$. Then $\mathbf{x}+\mathbf{y} = (3,1)$ and

$$\left(3^p+1\right)^{\frac{1}{p}} = \|\mathbf{x}+\mathbf{y}\|_p \le \|\mathbf{x}\|_p + \|\mathbf{y}\|_p = \left(2^p+0\right)^{\frac{1}{p}} + \left(1^p+1^p\right)^{\frac{1}{p}} = 2 + 2^{\frac{1}{p}},$$

as desired. Hence, the parameter u satisfies the condition in Lemma 7.

Fix $\Delta > 0$ and consider the two-agent instance $\mathbf{x} = (0, \Delta)$ with prediction $\pi = \frac{\Delta}{2}$. Note that π coincides with the optimal location y^* for this profile. By Lemma 7, the expected absolute deviation of the mechanism's output y from π satisfies

$$d = \mathbb{E}\big[|y-\pi|\big] \le \frac{\Delta}{2} \cdot T,$$

where $T \in [0,1]$ is the unique solution to $(1+T)^p + (1-T)^p = 2(1+u)^p$. It is easy to see that $T \in \left[0, \frac{1}{2}\right]$, because when u attains its maximum permissible value $\left(\frac{3^p+1}{2^{p+1}}\right)^{1/p} - 1$ the solution is $T = \frac{1}{2}$, and T increases with u.

Now consider a different profile $\mathbf{x}' = \left(0, \frac{\Delta}{2}\right)$ under the *same* prediction $\pi = \frac{\Delta}{2}$. Let y' denote the random output of the mechanism on $\mathbf{x}'$ with this prediction. We examine the incentive of agent 2, whose true location in $\mathbf{x}'$ is $\frac{\Delta}{2}$. If agent 2 were to misreport its location as Δ (while agent 1 reports truthfully), the reported profile would become $\mathbf{x} = (0, \Delta)$. Strategyproofness requires that agent 2 does not strictly benefit from this deviation, i.e.,

$$\mathbb{E}\big[\big|y' - \tfrac{\Delta}{2}\big|\big] \le \mathbb{E}\big[\big|y - \tfrac{\Delta}{2}\big|\big] = \mathbb{E}\big[|y-\pi|\big] \le \frac{\Delta}{2} \cdot T. \tag{2}$$

We now analyze the expected L_p social cost of the mechanism on profile $\mathbf{x}'$. Define a convex function $f(z) = \left(|z|^p + |z - \Delta/2|^p\right)^{1/p}$ for $z \in \mathbb{R}$. By Jensen's inequality,

$$\mathbb{E}\big[\mathrm{SC}_p(\mathbf{x}', y')\big] = \mathbb{E}\big[f(y')\big] \ge f\big(\mathbb{E}[y']\big). \tag{3}$$

Set $z = \mathbb{E}[y']$ and we obtain

$$\frac{\Delta}{4} \le \frac{\Delta}{2}(1-T) \le z,$$

where the first inequality is due to $T \le \frac{1}{2}$, and the second follows from (2) and $|\mathbb{E}[y'] - \frac{\Delta}{2}| \le \mathbb{E}\big[\big|y' - \frac{\Delta}{2}\big|\big]$. On this interval, f is increasing, so its minimum is achieved at the left endpoint $z = \frac{\Delta}{2}(1-T)$. Consequently,

$$f\big(\mathbb{E}[y']\big) \ge \left(\left(\frac{\Delta}{2}(1-T)\right)^p + \left(\frac{\Delta}{2}T\right)^p\right)^{\frac{1}{p}} = \frac{\Delta}{2} \cdot \left((1-T)^p + T^p\right)^{\frac{1}{p}}. \tag{4}$$

The optimal social cost for $\mathbf{x}'$ is attained at $y^* = \frac{\Delta}{4}$, giving

$$\mathrm{OPT}' = 2^{\frac{1}{p}} \cdot \frac{\Delta}{4}.$$

Combining (3) and (4), the robustness ratio satisfies

$$\beta \ge \frac{\mathbb{E}[\mathrm{SC}_p(\mathbf{x}', y')]}{\mathrm{OPT}'} \ge \frac{\frac{\Delta}{2} \cdot \left((1-T)^p + T^p\right)^{\frac{1}{p}}}{2^{\frac{1}{p}} \cdot \frac{\Delta}{4}} = 2^{1-\frac{1}{p}} \cdot \left((1-T)^p + T^p\right)^{\frac{1}{p}}.$$

This completes the proof. □

C Intermediate Prediction Quality

The consistency and robustness guarantees capture the two extremes in which the prediction is respectively perfect or arbitrary. We now give an error-dependent guarantee that interpolates between them. Fix an agent profile $\mathbf{x}$ and an admissible t, and write $M_t(\mathbf{x}, \pi)$ for the output of Mechanism 1 under prediction π.

Lemma 8. *For any two predictions $\pi, \pi' \in \mathbb{R}$,*

$$|M_t(\mathbf{x}, \pi) - M_t(\mathbf{x}, \pi')| \le |\pi - \pi'|.$$

Proof. Mechanism 1 returns a fixed order statistic of the n reports and the tn copies of the prediction. Replacing π by π' changes every entry of this augmented profile by at most $|\pi - \pi'|$. Since a fixed order statistic is 1-Lipschitz with respect to the ℓ_∞ distance between its input vectors, its output changes by at most the same amount. □

We will also use the following consequence of the reverse triangle inequality and Minkowski's inequality: for any two facility locations $y, z \in \mathbb{R}$,

$$\begin{aligned} |\mathrm{SC}_p(\mathbf{x}, y) - \mathrm{SC}_p(\mathbf{x}, z)| &\le \left(\sum_{i \in N} \big| |x_i - y| - |x_i - z| \big|^p \right)^{1/p} \\ &\le n^{1/p} |y - z|. \end{aligned} \tag{5}$$

Let y^* be the optimal facility location and write $\mathrm{OPT}_p(\mathbf{x}) = \mathrm{SC}_p(\mathbf{x}, y^*)$. For $\mathrm{OPT}_p(\mathbf{x}) > 0$, define the normalized prediction error

$$\eta(\mathbf{x}, \pi) = \frac{n^{1/p} |\pi - y^*|}{\mathrm{OPT}_p(\mathbf{x})}.$$

When $\mathrm{OPT}_p(\mathbf{x}) = 0$, all agents are colocated and Mechanism 1 outputs their common location, so the guarantee below is trivially satisfied.

Corollary 1. *For every profile with $\mathrm{OPT}_p(\mathbf{x}) > 0$ and every prediction π, Mechanism 1 satisfies*

$$\frac{\mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, \pi)\big)}{\mathrm{OPT}_p(\mathbf{x})} \le \min\left\{ A_p(t) + \eta(\mathbf{x}, \pi),\ B_p(t) \right\}.$$

Proof. Applying (5), Lemma 8, and the consistency guarantee of Theorem 1 gives

$$\begin{aligned} \mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, \pi)\big) &\le \mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, y^*)\big) + n^{1/p} \, |M_t(\mathbf{x}, \pi) - M_t(\mathbf{x}, y^*)| \\ &\le A_p(t) \, \mathrm{OPT}_p(\mathbf{x}) + n^{1/p} |\pi - y^*| \\ &= \big(A_p(t) + \eta(\mathbf{x}, \pi)\big) \mathrm{OPT}_p(\mathbf{x}). \end{aligned}$$

The robustness guarantee of Theorem 1 independently bounds the same ratio by $B_p(t)$. Taking the smaller of the two bounds proves the claim. □

This pointwise result also covers probabilistic prediction models. Let Π be an exogenous random prediction for a fixed profile, and let $\eta(\mathbf{x}, \Pi)$ be its normalized error. If

$$\Pr\big[\eta(\mathbf{x}, \Pi) \le \varepsilon\big] \ge 1 - \delta,$$

then Corollary 1 and robustness imply

$$\mathbb{E}_\Pi \left[\frac{\mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, \Pi)\big)}{\mathrm{OPT}_p(\mathbf{x})} \right] \le (1 - \delta) \min\{A_p(t) + \varepsilon, B_p(t)\} + \delta B_p(t).$$

In particular, setting $\varepsilon = 0$ gives the guarantee when the prediction is exactly correct with probability at least $1 - \delta$. Alternatively, if only the expected normalized error is known and $\mathbb{E}[\eta(\mathbf{x}, \Pi)] \le \bar{\eta}$, then

$$\mathbb{E}_\Pi \left[\frac{\mathrm{SC}_p\big(\mathbf{x}, M_t(\mathbf{x}, \Pi)\big)}{\mathrm{OPT}_p(\mathbf{x})} \right] \le \min\{A_p(t) + \bar{\eta}, B_p(t)\}.$$

Thus, the mechanism's performance interpolates continuously between consistency and robustness rather than treating every imperfect prediction identically.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We have ensured that all the claims the abstract and introduction accurately reflect the theoretical contributions and scope of this work, and they are formally developed and proven in the next following sections.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: This paper discussed all the gaps between the upper bounds and the lower bounds, and illustrated the limitations of the results.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The full assumptions are provided in the section of preliminaries, and the results with proof sketch are provided in the next sections. The omitted proofs can be found in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: This paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [N/A]

Justification: This paper does not include experiments.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: This paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: This paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: This paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: We confirm that this paper adheres to all its requirements since it is purely theoretical research.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We do not have any negative societal impacts. Also this paper brings foundational and theoretical work; a concise discussion of its potential long-term societal implications is provided in the introduction.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This paper does not include experiments.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: No existing code, data, or model assets are used.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: No new assets are released.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The research does not involve crowdsourcing or human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The research does not involve human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs were not used as an important, original, or non-standard component of the core methodology.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.