Mechanism Design for Facility Location using Predictions

Author Name
Affiliation
email@example.com

Abstract

We study mechanisms for the facility location problem augmented with predictions of the optimal facility location. We demonstrate that an egalitarian viewpoint which considers *both* the maximum distance of any agent from the facility *and* the minimum utility of any agent provides important new insights compared to a viewpoint that just considers the maximum distance. As in previous studies, we consider performance in terms of consistency (worst case when predictions are accurate) and robustness (worst case irrespective of the accuracy of predictions). By considering how mechanisms with predictions can perform poorly, we design new mechanisms that are more robust. Indeed, by adjusting parameters, we demonstrate how to trade robustness for consistency. We go beyond the single facility problem by designing novel strategy proof mechanisms for locating two facilities with bounded consistency and robustness that use two predictions for where to locate the two facilities.

1 Introduction

In online algorithms, an elegant method to improve (worst-case) performance is to provide predictions about future inputs. Such predictions might come from machine learning methods applied to historical data. For example, a cache scheduler has to decide which pages to evict from the cache without knowing future requests for page access. However, we can use machine learning to predict future cache requests, improving performance of the cache scheduler when these predictions are accurate. Recently, researchers have proposed exploiting predictions in mechanism design, arguing that they will transform the design and analysis of mechanisms in multiagent systems.

Most relevant to this work, Agrawal *et al.* [2022] proposed augmenting mechanisms for facility location with predictions of the optimal location of the facility. Facility location is a classic problem where we decide the location of a facility so as to minimize the distance of agents from the facility. It models a number of collective decision making problems such as deciding the optimal room temperature for a class room, the maintenance budget for an apartment complex, or the best location for a mobile phone tower. Our aim is to use predictions of the optimal facility location to provide better performance guarantees when predictions are accurate (consistency) without sacrificing worst-case performance when they are not (robustness).

We look in more detail at the mechanisms for facility location proposed in [Agrawal *et al.*, 2022] that take account of the predicted optimal location of the facility. We demonstrate the importance of the precise choice of objective. In particular, we show that an egalitarian viewpoint considering *both* the maximum distance of any agent from the facility *and* the minimum utility provides important new insights. We also demonstrate the value of censoring extreme predictions. Insights from our study (such as the value of censoring extreme predictions) could be useful in the design of mechanisms with predictions in other application domains such as fair division, school choice or ad auctions.

2 Facility location

In a facility location problem, we need to decide where to locate a facility to serve a set of agents. We consider n agents located at x_1 to x_n. We assume without loss of generality that $x_1 \leq \ldots \leq x_n$. A mechanism f locates the facility at y. Formally, $f(x_1, \ldots, x_n) = y$. We let d_i be the distance of agent i to the facility: $d_i = |x_i - y|$. As in a number of previous studies, we assume that agents and facilities are on the interval $[0, 1]$, and the utility of agent i is $1 - d_i$. The interval could be $[a, b]$ supposing we normalise by $b - a$.

Having agents and facilities lie on a fixed interval is both practically and theoretically interesting. In practice, agents and facilities can be limited to a fixed interval due to physical constraints. For example, when locating charging stations in a warehouse, robots and charging stations might be limited to the warehouse. As a second example, when setting a thermostat, we might be limited by the boiler. As a third example, when locating a distribution centre, the centre might have to be on the fixed road network. There are thus many settings where locations are limited to an interval. Restricting agents to a fixed interval also limits the extent to which agents can misreport their location to gain advantage. A fixed interval has been used in several recent studies (e.g. [Aziz *et al.*, 2021; Mei *et al.*, 2016]).

Our focus is on egalitarian mechanisms that look to minimize the maximum distance any agent must travel or, equivalently, to maximize the minimum utility of any agent. When consider approximation ratios of the optimal solution, the utility and distance viewpoints offer different insights. Indeed, we will show that the viewpoint of the minimum utility of any agents provides an alternative but useful perspective that is complementary to that provided by the maximum distance.

We consider mechanisms for locating facilities with good normative properties. One such property is unanimity. A mechanism is *unanimous* iff the facility is located where all agents agree. Formally f is unanimous iff for any x, we have $f(x, \ldots, x) = x$. A simple fairness property is anonymity. A mechanism is *anonymous* iff permuting the agents does not change the outcome. Formally f is anonymous iff for any permutation σ, we have $f(x_{\sigma(1)}, \ldots, x_{\sigma(n)}) = f(x_1, \ldots, x_n)$. A mechanism is *Pareto efficient* iff we cannot move the facility location to make one agent better off without hurting other agents. Formally f is Pareto efficient iff for any $x_j, \ldots, x_n$, there does not exist a location z and agent i with $|x_i - z| < |x_i - f(x_1, \ldots, x_n)|$ and $|x_j - z| \leq |x_j - f(x_1, \ldots, x_n)|$ for all $j \in [1, n]$. Another important property is resistance to manipulation. A mechanism is *strategy proof* iff no agent can mis-report their location and reduce their distance to the nearest facility. Formally f is strategy proof iff for any $x_1, \ldots, x_n$, and any agent i, it is not the case that there exists x'_i with $|x_i - f(x_1, \ldots, x'_i, \ldots, x_n)| < |x_i - f(x_1, \ldots, x_i, \ldots, x_n)|$. We will consider how well strategy proof mechanisms approximate an objective like the optimal maximum distance or minimum utility. A mechanism has an approximation ratio ρ for a maximization (minimization) objective iff the objective it returns is at least $1/\rho$ (at most ρ) times the optimal.

We consider a number of strategy proof mechanisms. Many are based on the function $median(z_1, \ldots, z_p)$ which returns z_i where $|\{j | z_j < z_i\}| < \lceil p/2 \rceil$ and $|\{j | z_j > z_i\}| \leq \lfloor p/2 \rfloor$. For example, the GENMEDIAN mechanism locates a facility at $median(x_1, \ldots, x_n, z_1, \ldots, z_{n-1})$ where the $n-1$ parameters z_1 to z_{n-1} are "phantom" agents at fixed locations. Moulin [1980] proved that a mechanism is anonymous, Pareto efficient and strategy-proof iff it is GENMEDIAN. The LEFTMOST mechanism is an instance of GENMEDIAN with $z_i = 0$ for $i \in [1, n)$, locating the facility at the leftmost agent. The RIGHTMOST mechanism is an instance of GENMEDIAN with $z_i = 1$ for $i \in [1, n)$, locating the facility at the rightmost agent. The MEDIAN mechanism is an instance of GENMEDIAN with $z_i = 0$ for $i \leq \lfloor n/2 \rfloor$ and 1 otherwise, locating the facility at the median agent. The MIDORNEAREST mechanism is an instance of GENMEDIAN with $z_i = 1/2$ for $i \in [1, n)$. It locates the facility either at $1/2$ if $x_1 \leq 1/2 \leq x_n$, otherwise at the agent nearest to $1/2$.

In [Agrawal *et al.*, 2022], mechanisms for facility location are augmented with a prediction of the optimal facility location. For example, if x_1 and x_n are the maximum and minimum reported locations of the agents, the mechanism MINMAXP(x_1, x_n, π) returns the predicted solution π as facility location when $x_1 \leq \pi \leq x_n$, otherwise it returns x_1 when $\pi < x_1$, and x_n when $\pi > x_n$. In general, our goal is for predictions to improve the performance when accurate and not to hinder performance when inaccurate. A mechanism with prediction is *α-consistent* iff, when the prediction is correct, the mechanism has an approximation ratio of α or better. A mechanism with prediction is *β-robust* iff, irrespective of the quality of the prediction, the mechanism has an approximation ratio of β or better. The MINMAXP mechanism is strategy proof and, with respect to maximum distance, is 1-consistent and 2-robust (i.e. returns the optimal maximum distance when the prediction is exact, and 2-approximates it otherwise) [Agrawal *et al.*, 2022].

3 Single facility

We begin with the simplest setting where we locate a single facility on the interval $[0, 1]$. You might think that the MINMAXP mechanism was optimal and all that could be usefully said about strategy proof mechanisms that exploit predictions. Clearly no mechanism can do better than 1-consistency. And no deterministic and strategy proof mechanism can be better than 2-robustness (Procaccia and Tennenholts [2013] demonstrate this for the real line. However, the result easily extends to any fixed interval). The MINMAXP mechanism therefore has optimal consistency and robustness with respect to the maximum distance.

An utility viewpoint suggests, however, that there is more to uncover about egalitarian mechanisms exploiting predictions. Consider the subtly different egalitarian objective of the minimum utility, and the approximation ratios that can be achieved of this objective. The MINMAXP mechanism is **far from optimal from this perspective**. In fact, there is no bound on how badly it approximates the minimum utility.

Theorem 1. *The* MINMAXP *mechanism is 1-consistent with respect to the optimal minimum utility, but has no bound on its robustness.*

Proof. If the prediction is accurate then, as the mechanism is strategy proof, $x_1 \leq \pi = \frac{x_1 + x_n}{2} \leq x_n$. The facility is therefore located at this accurate prediction, and the mechanism is 1-consistent. For robustness, suppose $\pi = x_1 = 0$ and $x_n = 1$. Then the facility is located at 0, giving an minimum utility of zero. However, the optimal minimum utility is $1/2$ with the facility at $1/2$. Hence robustness is unbounded. □

Considering the approximation ratio of the optimal maximum distance focuses attention on problem instances where distances are small and all agents are necessarily close to the facility location. It ignores those more challenging problem instances where distances are large and some agents are necessarily far from the facility location. Unfortunately the MINMAXP mechanism may approximate poorly certain instances in which agents must travel large distances.

We compare this lack of robustness with the simple strategy proof MIDORNEAREST mechanism. This has consistency and robustness that is bounded with respect to both maximum distance and minimum utility.

Theorem 2. *The* MIDORNEAREST *mechanism is* $3/2$*-consistent and* $3/2$*-robust with respect to the optimal minimum utility. It is* 2*-consistent and* 2*-robust with respect to the optimal maximum distance.*

Proof. MIDORNEAREST ignores the prediction so consistency is the same as robustness. Theorem 1 in [Walsh, 2024] demonstrates that the mechanism $3/2$-approximates the minimum utility and 2-approximates the maximum distance. $\square$

Note that no deterministic and strategy proof mechanism can do better than $3/2$-robustness with respect to the minimum utility [Walsh, 2024], or 2-robustness with respect to the maximum distance [Procaccia and Tennenholtz, 2013]. The MIDORNEAREST mechanism is actually an instance of the MINMAXP mechanism when the predicted optimal location is $1/2$. It is extreme predictions away from $1/2$ that lead to the lack of robustness of the MINMAXP mechanism.

Of course, the MIDORNEAREST mechanism is not exploiting any information about predicted optimal facility location. Can a mechanism that uses the prediction do better? Yes, mechanisms which are responsive to the predicted optimal location can do better. However, to get good (bounded) robustness with respect to the minimum utility, we must avoid extreme predictions near the interval end points.

We propose here a **new mechanism** guided by **non-extreme predictions** that has bounded robustness. The MINMAXP$_\gamma$ mechanism is a truncated version of the MINMAXP mechanism with a parameter $\gamma \in [0, 1/2]$. It maps the prediction π onto $\max(\gamma, \min(\pi, 1-\gamma))$, and then applies the MINMAXP mechanism to this truncated prediction. This mapping limitxss predictions to the interval $[\gamma, 1-\gamma]$. The MINMAXP$_\gamma$ mechanism is the MINMAXP mechanism when $\gamma = 0$, and the MIDORNEAREST mechanism when $\gamma = 1/2$. For $0 < \gamma < 1/2$, it is a synthesis of the two mechanisms. In fact, it smoothly interpolates between the MINMAXP mechanism (which is optimal with respect to consistency) and the MIDORNEAREST mechanism (which, as we will argue shortly, is optimal with respect to robustness).

Theorem 3. *For $\gamma \in [0, 1/2]$, the MINMAXP$_\gamma$ mechanism is strategy proof, $\frac{(2-\gamma)}{(2-2\gamma)}$-consistent and $\frac{(1+\gamma)}{2\gamma}$-robust with respect to the optimal minimum utility. It is 1-consistent with respect to the optimal maximum distance when $\gamma = 0$, but 2-consistent when $\gamma > 0$. It is always 2-robust with respect to the optimal maximum distance.*

Proof. Strategy proofness is immediate from that of the untruncated mechanism. With respect to the optimal minimum utility, suppose the prediction π is correct. There are five cases. In the first case, $\pi \leq \gamma/2$. Let the minimum utility be $1-b$ with $b \leq \pi$. The mechanism locates the facility at $\pi + b$ giving a minimum utility of $1-2b$. The approximation ratio is thus $\frac{(1-b)}{(1-2b)}$. This has a maximum of $\frac{(2-\gamma)}{(2-2\gamma)}$ when $\pi = b = \gamma/2$. In the second case, $\gamma/2 \leq \pi \leq \gamma$ and the minimum utility is $1-b$ with $b \leq \gamma - \pi$. The mechanism locates the facility at $\pi + b$ giving a minimum utility of $1-2b$. The approximation ratio is thus $\frac{(1-b)}{(1-2b)}$. This again has a maximum of $\frac{(2-\gamma)}{(2-2\gamma)}$ when $\pi = b = \gamma/2$. In the third case, $\gamma/2 \leq \pi \leq \gamma$ and the minimum utility is $1-b$ with $\pi \geq b \geq \gamma - \pi$. The mechanism locates the facility at γ giving a minimum utility of $1-(\gamma-(\pi-b))$. The approximation ratio is thus $\frac{(1-b)}{(1-b-\gamma+\pi)}$. This again has a maximum of $\frac{(2-\gamma)}{(2-2\gamma)}$ when $\pi = b = \gamma/2$. In the fourth case, $\gamma \leq \pi \leq 1-\gamma$. The mechanism locates the facility at π giving the optimal minimum utility and an approximation ratio of 1. In the fifth case, $\pi \geq 1-\gamma$. This is symmetric to the first three cases. Over the five cases, the largest approximation ratio is $\frac{(2-\gamma)}{(2-2\gamma)}$.

Now suppose the prediction is incorrect. There are four cases. In the first case, x_n is less than or equal to the truncated prediction π'. The optimal minimum utility is $1 - \frac{(x_n - x_1)}{2}$. However, the mechanism locates the facility at x_n giving a minimum utility of $1-(x_n - x_1)$. The approximation ratio is therefore $\frac{(2-x_n+x_1)}{2(1-x_n+x_1)}$. This is maximized for $x_1 = 0$ and $x_n = 1-\gamma$ when the ratio is $\frac{(1+\gamma)}{2\gamma}$. In the second case, π' is between x_1 and x_n or equal to x_1, and nearer to x_1 than x_n. The optimal minimum utility is again $1 - \frac{(x_n-x_1)}{2}$. However, the mechanism locates the facility at π'. The minimum utility is $1-(x_n-\pi')$. The approximation ratio is therefore $\frac{(2-x_n+x_1)}{2(1-x_n+\pi')}$. This is maximized for $x_1 = 0$, $x_n = 1-\gamma$ and $\pi' = \gamma$ when the ratio is $\frac{(1+\gamma)}{2\gamma}$. In the third case, π' is between x_1 and x_n or equal to x_n, and not nearer to x_1 then x_n. This is symmetric to the second case. In the fourth case, π' is greater than x_n. This is symmetric to the first case. Over the four cases, the largest approximation ratio is $\frac{(1+\gamma)}{2\gamma}$.

With respect to the optimal maximum distance, suppose the prediction π is correct. For $\gamma = 0$, the MINMAXP$_\gamma$ mechanism is equivalent to MINMAXP mechanism which is 1-consistent and 2-robust. For $\gamma > 0$ there are five cases. In the first case, $\pi \leq \gamma/2$. Let the maximum distance be b with $b \leq \pi$. The mechanism locates the facility at $\pi + b$ giving a maximum distance of $2b$. The approximation ratio is thus 2. In the second case, $\gamma/2 \leq \pi \leq \gamma$ and the maximum distance is b with $b \leq \gamma - \pi$. The mechanism locates the facility at $\pi + b$ giving a maximum distance of $2b$. The approximation ratio is thus again 2. In the third case, $\gamma/2 \leq \pi \leq \gamma$ and the maximum distance is b with $\pi \geq b \geq \gamma - \pi$. The mechanism locates the facility at γ giving a maximum distance of $(\gamma - (\pi - b))$. The approximation ratio is thus $\frac{(b+\gamma-\pi)}{b}$. This has a maximum of 2 when $\pi = b = \gamma/2$. In the fourth case, $\gamma \leq \pi \leq 1-\gamma$. The mechanism locates the facility at π giving the optimal maximum distance and an approximation ratio of 1. In the fifth case, $\pi \geq 1-\gamma$. This is symmetric to the first three cases. Over the five cases, the largest approximation ratio is 2.

Now suppose again that the prediction is incorrect. There are four cases. In the first case, x_n is less than or equal to the truncated prediction π'. The optimal maximum distance is $\frac{(x_n-x_1)}{2}$. However, the mechanism locates the facility at x_n giving a maximum distance of $(x_n - x_1)$. The approximation ratio is therefore 2. In the second case, π' is between x_1 and x_n or equal to x_1, and nearer to x_1 than x_n. The optimal maximum distance is again $\frac{(x_n-x_1)}{2}$. However, the mechanism locates the facility at the truncated prediction π'. The maximum distance is $x_n - \pi'$. The approximation ratio is therefore $\frac{2(x_n-\pi')}{(x_n-x_1)}$. This is maximized for $x_1 = \gamma$, $x_n = 1$ and $\pi' = \gamma$ when the ratio is 2. In the third case, π' is between x_1 and x_n or equal to x_n, and not nearer to x_1 than x_n. This is symmetric to the second case. In the fourth case, π' is greater than x_n. This is symmetric to the first case. Over the four cases, the largest approximation ratio is 2. $\square$

Note that when the prediction is in $[\gamma, 1-\gamma]$, the MINMAXP$_\gamma$ mechanism does even better. In this setting, the mechanism is 1-consistent with respect to minimum utility or maximum distance. It is only with extreme predictions (less than γ or greater than $1-\gamma$) where consistency drops. Note also that by adjusting γ, we can **trade consistency for robustness** (see Figure 1 for a visualization of this). At $\gamma = 0$, the MINMAXP$_\gamma$ mechanism is 1-consistent with respect to minimum utility but has unbounded robustness. Increasing γ decreases robustness but increases consistency. At $\gamma = 1/2$, the mechanism is $3/2$-consistent and $3/2$-robust.

We return now to the reason that we proposed a mechanism that smoothly interpolates between the MINMAXP mechanism (the MINMAXP$_\gamma$ mechanism with $\gamma = 0$) and the MIDORNEAREST mechanism (the MINMAXP$_\gamma$ mechanism with $\gamma = 1/2$). The reason is that the MINMAXP mechanism achieves the optimal consistency, while the MIDORNEAREST mechanism achieves the optimal robustness. To be more precise, the MIDORNEAREST mechanism achieves an optimal 2-approximation of the maximum distance, and an optimal $3/2$-approximation of the minimum utility. Indeed, as we show next, the MIDORNEAREST mechanism is **the unique** anonymous, Pareto efficient and strategy proof mechanism that $3/2$-approximates the minimum utility.

Theorem 4. *No anonymous, Pareto efficient and strategy-proof mechanism besides the* MIDORNEAREST *mechanism has as good an approximation ratio of the minimum utility.*

Proof. Consider any anonymous, Pareto efficient and strategy-proof mechanism. This is a median mechanism with $n-1$ phantoms [Moulin, 1980]. If this is not the MIDORNEAREST mechanism, one of the phantoms will be different to $1/2$. Consider the smallest such phantom a. Suppose $0 \leq a < 1/2$. A dual argument holds for $1/2 < a \leq 1$. Consider one agent at 1 and the remaining agents at a. The facility is located at a, giving a minimum utility of a. The optimal minimum utility is $1/2 + a/2$. Therefore the approximation ratio is $\frac{(1+a)}{2a}$. For $0 \leq a < 1/2$, this is in $(3/2, \infty]$. Hence the approximation ratio is worse than $3/2$. $\square$

4 Randomized mechanisms

We can often achieve better approximation ratios in expectation with randomized mechanisms. A randomized mechanism returns a probability distribution over ex post outcomes. Consider, for example, the randomized and strategy proof LRM mechanism which locates the facility at x_1 with probability $1/4$, at $\frac{(x_1+x_n)}{2}$ with probability $1/2$, and at x_n with probability $1/4$. This achieves an optimal $3/2$-approximation of the maximum distance in expectation (Procaccia and Tennenholtz [2013] show this for the real line but the result easily extends to any interval). The LRM mechanism does not do quite as well at approximating the optimal minimum utility, only 2-approximating it in expectation [Walsh, 2024].

The randomized LRM mechanism can be adapted to take advantage of predictions. Given a parameter $\delta \in [0, 1/2]$, the LRMP mechanism proposed in [Agrawal *et al.*, 2022] uses the LRM mechanism with probability 2δ, and the MINMAXP mechanism with probability $1-2\delta$. This achieves an optimal $1+\delta$-consistency and $2-\delta$-robustness in expectation with respect to the maximum distance (Proposition 1 and Theorem 1 in [Agrawal *et al.*, 2022]).

Theorem 5. *For $\delta \in [0, 1/2]$, the* LRMP *mechanism is $\frac{1}{(1-\delta)}$-consistent and $\frac{1}{\delta}$-robust in expectation with respect to the optimal minimum utility.*

Proof. Suppose the prediction is correct and the optimal minimum utility u. The expected minimum utility is $2\delta\frac{u}{2} + (1-2\delta)u = (1-\delta)u$. Hence it is $\frac{1}{(1-\delta)}$-consistent. Suppose the prediction is incorrect. The expected minimum utility is $2\delta\frac{u}{2} = \delta u$. Hence it is $\frac{1}{\delta}$-robust. $\square$

For a given δ, both the approximation ratios for consistency and robustness are worse with respect to minimum utility compared to the ratios for maximum distance.

As with deterministic mechanisms, censoring extreme facility locations improves performance. Let $y = \max(1/3, \min(x_1, 2/3))$ and $z = \max(1/3, \min(2/3, x_n))$. The LRMT mechanism proposed in [Walsh, 2024] locates the facility at y with probability $1/4$, at $(y+z)/2$ with probability $1/2$ and z otherwise. This mechanism truncates facility locations to $[1/3, 2/3]$. This mechanism is strategy proof and achieves in expectation an optimal $4/3$-approximation of the minimum utility, and a 2-approximation of the maximum distance [Walsh, 2024]. The LRMT mechanism can also be adapted to take advantage of predictions by combining it with the MINMAXP mechanism. Given a parameter $\delta \in [0, 1/2]$, the LRMTP mechanism uses the LRMT mechanism with probability 2δ, and the MINMAXP mechanism with probability $1-2\delta$.

Theorem 6. *For $\delta \in [0, 1/2]$, the* LRMTP *mechanism is strategy proof, $\frac{2}{(2-\delta)}$-consistent and $\frac{2}{3\delta}$-robust in expectation with respect to the optimal minimum utility. It is also $1+2\delta$-consistent and 2-robust in expectation with respect to the optimal maximum distance.*

Proof. Strategy proofness is immediate from the strategy proofness of the constituent mechanisms and the fact that the choice of mechanism is independent of the agents' reports.

Suppose the prediction is correct and the optimal minimum utility is u. The expected minimum utility is $2\delta\frac{3}{4}u + (1-2\delta)u = (1-\frac{\delta}{2})u$. Hence it is $\frac{2}{(2-\delta)}$-consistent. Suppose the prediction is incorrect. The expected minimum utility is $2\delta\frac{3}{4}u = \frac{3\delta}{2}u$. Hence it is $\frac{2}{3\delta}$-robust.

Suppose the prediction is correct and the optimal maximum distance is d. The expected maximum distance is $2\delta 2d + (1-2\delta)d = (1+2\delta)d$. Hence it is $1+2\delta$-consistent. Suppose the prediction is incorrect. With respect to the maximum distance, since it is a probablistic mixture of two 2-robust mechanisms, it is itself 2-robust in expectation. $\square$

In approximating the minimum utility, the LRMTP mechanism outperforms the LRMP mechanism in four ways:

1. For any fixed $\delta > 0$, both the consistency and robustness of LRMTP are better than for LRMP.
2. For any given consistency in the interval $(1, 4/3]$, LRMTP achieves an expected robustness than is three times smaller than for LRMP.
3. For any given robustness greater than or equal to $4/3$, LRMTP achieves a smaller expected consistency than LRMP.
4. LRMTP achieves a consistency in $[1, 3/4]$, while LRMP achieves a consistency in $[1, 2]$.

On the other hand, in approximating the maximum distance, the LRMP mechanism outperforms the LRMTP mechanism again in four ways:

1. For any fixed $\delta > 0$, both the consistency and robustness of LRMP are better than for LRMTP.
2. For any given consistency $c \in (1, 3/2]$, LRMP achieves an expected robustness of $3 - c$ which is strictly smaller than the fixed 2-robustness of LRMTP.
3. LRMP achieves a robustness in $[3/2, 2]$, while LRMTP is only ever 2-robust.
4. LRMP achieves a consistency in $[1, 3/2]$, while LRMTP achieves a consistency in $[1, 2]$.

By adjusting δ, both mechanisms again **trade consistency for robustness**. At $\delta = 0$, both LRMP and LRMTP are 1-consistent with respect to minimum utility but have unbounded robustness. Increasing δ decreases robustness but increases consistency. At $\delta = 1/2$, LRMP is 2-consistent and 2-robust with respect to minimum utility, while LRMTP is $4/3$-consistent and $4/3$-robust. See Figure 1 for a visualization.

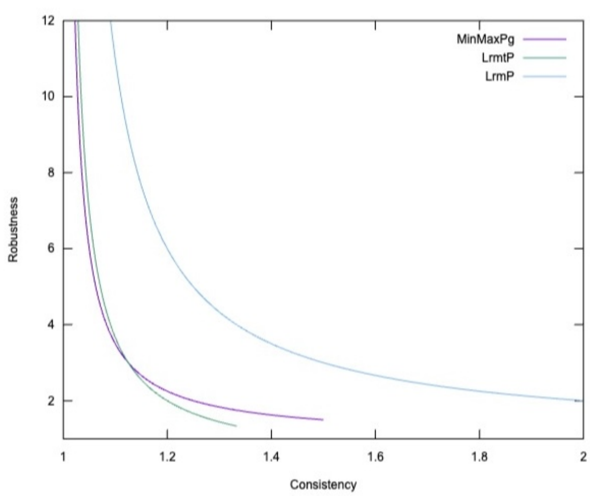

Figure 1: Trade-off between consistency (x-axis) and robustness (y-axis) with respect to the minimum utility for the MINMAXP$_\gamma$ mechanism when varying $\gamma \in [0, 1/2]$, and for the LRMP and LRMTP mechanisms when varying $\delta \in [0, 1/2]$.

5 Two facilities

We now design several new mechanisms for locating two facilities in which the mechanism is provided with **two predictions**, one for the optimal location of the leftmost facility, and another for the optimal location of the rightmost facility. Xu and Lu [2022] propose a deterministic mechanism with predictions for the two facility problem that is $(1 + n/2)$-consistent and $(2n - 1)$-robust with respect to the maximum distance. They observe:

> *"Whether there is a mechanism with $o(n)$-consistent and a bounded robustness is a very interesting open question."*

We answer this open question positively in two ways. First, we design a novel deterministic mechanism for approximating the minimum utility with bounded consistency and robustness. Second, we design a novel randomized mechanism for approximating the maximum distance with bounded consistency and robustness.

The MINMAX2P mechanism locates the two facilities by applying the MINMAXP mechanism to each of the predictions in turn. That is, if x_1 and x_n are the maximum and minimum reported locations of the agents, and π_1 and π_2 are the two predicted locations of the facilities, then MINMAX2P(x_1, x_n, π_1, π_2) locates one facility at MINMAXP(x_1, x_n, π_1) and the other at MINMAXP(x_1, x_n, π_2). The next theorem demonstrates that the MINMAX2P mechanism is 1-consistent with respect to the maximum distance or minimum utility, $3/2$-robust with respect to the minimum utility, but has unbounded robustness with respect to the maximum distance.

As with one facility, we also adapt the mechanism to censor extreme predictions. This again lets us **trade consistency for robustness**. Given $\lambda \in [0, 1/4]$, MINMAX2P$_\lambda$ (x_1, x_n, π_1, π_2) maps the leftmost prediction π_1 onto $\pi_1' = \max(\lambda, \min(\pi, 1 - 3\lambda))$, the rightmost prediction π_2 onto $\pi_2' = \max(3\lambda, \min(\pi, 1 - \lambda))$, and then applies MINMAX2P $(x_1, x_n, \pi_1', \pi_2')$. For $\lambda = 0$, predictions are not censored. For $\lambda = 1/4$, the leftmost prediction is mapped onto $1/4$ while the rightmost prediction is mapped onto $3/4$. More generally, the leftmost prediction is mapped into $[\lambda, 1 - 3\lambda]$, and the rightmost prediction into $[3\lambda, 1 - \lambda]$.

Theorem 7. *For $\lambda \in [0, 1/4]$, the* MINMAX2P$_\lambda$ *mechanism is strategy proof, $\frac{(2-\lambda)}{(2-2\lambda)}$-consistent and $\frac{(3+2\lambda)}{2(1+2\lambda)}$-robust with respect to the optimal minimum utility. At $\lambda = 0$, it is 1-consistent and $3/2$-robust while at $\lambda = 1/4$, it is $7/6$-consistent and $7/6$-robust.*

With respect to the optimal maximum distance, MINMAX2P$_\lambda$ *is 1-consistent and has unbounded robustness at $\lambda = 0$, and has unbounded consistency and robustness for $\lambda > 0$.*

Proof. Strategy proofness is immediate from that of the untruncated mechanism. With respect to the optimal minimum utility, suppose the two predictions are correct. Using a similar case analysis to Theorem 1, the worst case is when the optimal facility location is halfway between 0 and λ, and agents served by this facility are in $[0, \lambda]$ including at the endpoints of the interval. Suppose in this case that the optimal minimum utility is $1 - b$. Then $\pi_1 = b = \lambda/2$, and $x_1 = 0$. The minimum utility of the solution returned by the MINMAX2P$_\lambda$ mechanism is just $1-2b$. This gives an approximation ratio of

$\frac{(1-b)}{(1-2b)} = \frac{(2-\lambda)}{(2-2\lambda)}$. The MINMAX2P$_\lambda$ mechanism is therefore $\frac{(2-\lambda)}{(2-2\lambda)}$-consistent.

Now we consider the possibility that the predictions are incorrect. The leftmost truncated prediction is in $[\lambda, 1-3\lambda]$, and the rightmost prediction in $[3\lambda, 1-\lambda]$. Again using a similar case analysis to Theorem 1, the worst case is when the facilities are located by the mechanism at their most extreme point (i.e. $x_1 = \lambda, x_n = 1-\lambda$). The optimal minimum utility in this setting is $\frac{(3+2\lambda)}{4}$, while the minimum utility of the solution returned by the MINMAX2P$_\lambda$ mechanism is just $\frac{(1+2\lambda)}{2}$. This gives an approximation ratio of $\frac{(3+2\lambda)}{2(1+2\lambda)}$. The MINMAXP$_\gamma$ mechanism is therefore $\frac{(3+2\lambda)}{2(1+2\lambda)}$-robust with respect to the minimum utility.

With respect to the optimal maximum distance, suppose the predictions are correct. For $\lambda = 0$, the MINMAX2P$_\lambda$ mechanism is equivalent to MINMAX2P mechanism. It is easy to see that this is 1-consistent. For $\lambda > 0$, consider agents at 0 and 1. The optimal maximum distance is zero, but the MINMAX2P$_\lambda$ mechanism locates facilities at λ and $1-\lambda$, giving a maximum distance of λ. The consistency is therefore unbounded. Similarly the robustness is unbounded with respect to the maximum distance irrespective of λ. □

As with locating a single facility, we can achieve better consistency and robustness in expectation with randomized mechanisms. The randomized RANDENDS mechanism (called Mechanism 2 by Procaccia and Tennenholtz [2013]) locates two facilities in three ways: (1) at x_1 and x_n with probability $1/2$; (2) at x_1+2d and x_n-2d with probability $1/6$ where $2d$ is the optimal minimum distance; (3) and at x_1+d and x_n-d with the remaining probability $1/3$. The mechanism is strategy proof and $5/3$-approximates the maximum distance in expectation [Procaccia and Tennenholtz, 2013]. It achieves an even better approximation ratio with respect to the minimum utility.

Theorem 8. *The* RANDENDS *mechanism* $9/7$*-approximates the optimal minimum utility in expectation.*

Proof. With probability $1/3$, the mechanism has a minimum utility $1-d$, and with probability $2/3$, it has a minimum utility of $1-2d$. The expected minimum utility is thus $\frac{(1-d+2-4d)}{3} = \frac{(3-5d)}{3}$. This compares to an optimal minimum utility of $1-d$. The approximation ratio is thus $\frac{3(1-d)}{(3-5d)}$. This is maximized for $d = 1/4$ when it is $9/7$. Hence, the mechanism $9/7$-approximates the minimum utility. □

The randomized RANDENDS mechanism can be augmented to take advantage of two predictions for the optimal locations of the two facilities. Given a parameter $\theta \in [0, 1/2]$, the RANDENDS2P uses RANDENDS with probability 2θ, and MINMAX2P with probability $1-2\theta$.

Theorem 9. *The* RANDENDS2P *mechanism is strategy proof. With respect to the optimal minimum utility, it is* $\frac{9}{(9-4\theta)}$*-consistent and* $\frac{9}{2(3+\theta)}$*-robust for* $\theta \in [0, 1/2]$. *With respect to the optimal maximum distance, it is* $\frac{(3+4\theta)}{3}$*-consistent for* $\theta \in [0, 1/2]$, $5/3$*-robust for* $\theta = 1/2$. *and has unbounded robustness for* $\theta < 1/2$.

Proof. Strategy proofness is immediate from the strategy proofness of the constituent mechanisms and the fact that the choice of mechanism is independent of the agents' reports.

Suppose the prediction is correct and the optimal minimum utility u. The expected minimum utility is $2\theta\frac{7}{9}u + (1-2\theta)u = (1-\frac{4\theta}{9})u$. Hence it is $\frac{9}{(9-4\theta)}$-consistent. Suppose the prediction is incorrect. The expected minimum utility is $2\theta\frac{7}{9}u + (1-2\theta)\frac{2}{3}u = \frac{2(3+\theta)}{9}u$. Hence it is $\frac{9}{2(3+\theta)}$-robust.

Suppose the prediction is correct and the optimal maximum distance is d. The expected maximum distance is $2\theta\frac{5}{3}d + (1-2\theta)d = \frac{(3+4\theta)}{3}d$. Hence it is $\frac{(3+4\theta)}{3}$-consistent. □

This mechanism again lets us **trade consistency for robustness**. At $\theta = 0$, the RANDENDS2P mechanism is 1-consistent and $3/2$-robust with respect to minimum utility. Increasing θ decreases robusness but increases consistency. At $\theta = 1/2$, it is $9/7$-consistent and $9/7$-robust.

6 Characterization of consistent mechanisms

You might wonder why we have considered so far mechanisms that are mostly based on (truncated versions of) MINMAXP or mixing together combinations of mechanisms including MINMAXP. We now give a result that characterizes strategy proof mechanisms with predictions which achieve good levels of consistency. This characterization result demonstrates the central role played by the MINMAXP mechanism. When predictions are extreme (i.e. the optimal facility location is predicted to be at 0 or 1), there are multiple mechanisms that are 1-consistent (e.g. MIDORNEAREST, MEDIAN and MINMAXP). But when predictions are not extreme, the only mechanism that is better than 2-consistent with respect to the maximum distance is MINMAXP.

Theorem 10. *For non-extreme predictions, the only deterministic, strategy proof, anonymous and Pareto efficient mechanism for locating a single facility that is better than* 2*-consistent with respect to the maximum distance is the* MINMAXP *mechanism which is* 1*-consistent. It is also the only such mechanism that is* 1*-consistent with respect to the minimum utility.*

Proof. Any such mechanism is a generalized median mechanism with $n-1$ phantoms [Moulin, 1980]. If the phantoms are all at the predicted facility location, then we have the MINMAXP mechanism. Suppose instead that one or more of those phantoms is not at the predicted and correct facility location π. And suppose $\pi \leq 1/2$. There is a dual argument for $\pi \geq 1/2$. By the non-extreme assumption $\pi > 0$. Let ρ be the largest such phantom different to π. There are three cases. In the first case $\rho < \pi$. Consider $n-1$ agents at ρ and one at $2\pi-\rho$. As $0 < \pi \leq 1/2$, $2\pi-\rho$ is in $(\pi, 1]$. The optimal facility location is π as required. However, the mechanism locates the facility at ρ which is twice the optimal maximum distance from the agent at $2\pi-\rho$. In the second case $\pi < \rho \leq 2\pi$. Consider one agent at $2\pi-\rho$, and the other $n-1$ agents at ρ.

Note that $2\pi - \rho$ is in $[0, \pi)$. The optimal facility location is π as required. However, the mechanism locates the facility at ρ which is twice the optimal maximum distance from the agent at $2\pi - \rho$. In the third case $2\pi < \rho$. Consider one agent at 0, and the other $n-1$ agents at 2π. Note that 2π is in $(0, 1]$. The facility is located at 2π at twice the optimal maximum distance from the agent at 0. In each case, the mechanism is 2-consistent with respect to the maximum distance. By a similar argument, MINMAXP is the only such mechanism that is 1-consistent with respect to the minimum utility. □

Note that predictions can only be extreme when locating a facility on an interval. On the (unbounded) real line, predictions are never extreme and MINMAXP is the unique deterministic mechanism achieving better than 2-consistency.

7 Related work

There is a considerable literature on augmenting algorithms with predictions to improve worst-case performance (see [Mitzenmacher and Vassilvitskii, 2020] for a survey). Several recent surveys also summarize the considerable literature on mechanism design for facility location [Cheng and Zhou, 2015; Chan *et al.*, 2021].

Starting with Procaccia and Tennenholtz [2009], the analysis of strategy proof mechanisms for facility location has focused on approximating the total and maximum distance that agents must travel to the nearest facility. Many subsequent works asked similar questions about the design of approximate and strategy proof mechanisms for facility location (e.g. [Fotakis and Tzamos, 2010; Escoffier *et al.*, 2011; Lu *et al.*, 2010; Fotakis and Tzamos, 2013; Zhang and Li, 2014; Tang *et al.*, 2020; Aziz *et al.*, 2020; Goel and Hann-Caruthers, 2023]). Indeed, one of those recent surveys [Chan *et al.*, 2021] describes the design of strategy proof mechanisms which approximate well the total or maximum distance as the "classic setting" for approximate mechanism design.

Some recent work on approximate mechanism design for facility location has, however, considered other objectives such as the utility of agents as this can uncover fresh insight. For example, Walsh [2021; 2024] has looked at strategy proof mechanisms for facility location optimizing both the maximum distance and the minimum utility. As a second example, Han *et al.* [2023] look to optimize several objectives from the l-centrum family of metrics (which includes total and maximum distance) simultaneously. As a third example, Aziz *et al.* [2022] identified srategy proof mechanism that satisfy proportional fairness, a normative condition on the utility of agents. As a fourth example, Mei *et al.* [2019] considered strategy proof mechanisms maximizing the "happiness" of agents where happiness is a normalized utility.

For the online version of the facility location problem, Jiang *et al.* [2022] study online algorithms guided by predictions. Such online algorithms must irrevocably assign each agent to an open facility upon its arrival or must decide to open a new facility (at cost) to which to assign it. They provide a near-optimal online algorithm that offers a smooth tradeoff between the prediction error and the competitive ratio. Here, by comparison, we do not have to make online decisions but suppose the mechanism has access to location data for all agents simultaneously. Almanza *et al.* [2021] and Fotakis *et al.* [2021] also look at the online version of the facility location problem with predictions, and propose algorithms with good competitive ratios.

For obnoxious facility location, where the goal of agents is to be as far away from the facility as possible, Istrate and Bonchis [2022] study mechanism design with predictions. They present strategy proof mechanisms that explore the tradeoff between robustness and consistency on various metrics such as intervals, squares, circles, trees and hypercubes.

8 Conclusions

	max distance		**min utility**	
1 facility, determin	**consis**	**robust**	**consis**	**robust**
lower bound	1	2	1	$3/2$
MINMAXP	1	2	**1**	$\boldsymbol{\infty}$
MIDORNEAREST	**2**	**2**	$\mathbf{3/2}$	$\mathbf{3/2}$
MINMAXP$_\gamma$, $\gamma > 0$	**2**	**2**	$\boldsymbol{\frac{(2-\gamma)}{(2-2\gamma)}}$	$\boldsymbol{\frac{(1+\gamma)}{2\gamma}}$
MINMAXP$_\gamma$, $\gamma = 1/2$	**2**	**2**	$\mathbf{3/2}$	$\mathbf{3/2}$
1 facility, random				
lower bound	1	$3/2$	1	$4/3$
LRMP	$1+\delta$	$2-\delta$	$\boldsymbol{\frac{1}{(1-\delta)}}$	$\boldsymbol{\frac{1}{\delta}}$
LRMP, $\delta = 1/2$	$3/2$	$3/2$	**2**	**2**
LRMTP	$\mathbf{1+2\delta}$	**2**	$\boldsymbol{\frac{2}{(2-\delta)}}$	$\boldsymbol{\frac{2}{3\delta}}$
LRMTP, $\delta = 1/2$	**2**	**2**	$\mathbf{4/3}$	$\mathbf{4/3}$
2 facilities, determin				
lower bound	1	$n-2$	1	$10/9$
MINMAX2P	**1**	∞	**1**	$\mathbf{3/2}$
MINMAX2P$_\lambda$, $\lambda > 0$	∞	∞	$\boldsymbol{\frac{(2-\lambda)}{(2-2\lambda)}}$	$\boldsymbol{\frac{(3+2\lambda)}{2(1+2\lambda)}}$
MINMAX2P$_\lambda$, $\lambda = 1/4$	∞	∞	$\mathbf{7/6}$	$\mathbf{7/6}$
2 facilities, random				
lower bound	1	$3/2$	1	$10/9$
RANDENDS2P, $\theta = 0$	**1**	∞	**1**	$\mathbf{3/2}$
RANDENDS2P, $\theta < 1/2$	$\boldsymbol{\frac{(3+4\theta)}{3}}$	∞	$\boldsymbol{\frac{9}{(9-4\theta)}}$	$\boldsymbol{\frac{9}{2(3+\theta)}}$
RANDENDS2P, $\theta = 1/2$	$\mathbf{5/3}$	$\mathbf{5/3}$	$\mathbf{9/7}$	$\mathbf{9/7}$

Table 1: Summary of **consis**tency and **robust**ness results with respect to the optimal **max**imum **distance** or **min**imum **utility** of agents for **determin**istic or **random**ized strategy proof mechanisms. **Bold** font for those results proved here.

Our examination of mechanisms for facility location augmented with predictions of the optimal location demonstrates that an egalitarian viewpoint considering *both* the maximum distance any agent travels *and* the least utility of any agent provides a more complex picture of performance than one considering just maximum distance alone. Our results are summarized in Table 1. By considering how mechanisms can perform poorly, we proposed new deterministic and randomized mechanisms for locating a single facility that achieve bounded robustness with respect to both maximum distance and minimum utility. For locating two facilities, we also designed novel mechanisms with predictions with bounded robustness and consistency. These new mechanisms let us smoothly trade consistency for robustness. A repeated idea to obtain good performance was to censor extreme predictions.

References

[Agrawal *et al.*, 2022] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. In David M. Pennock, Ilya Segal, and Sven Seuken, editors, *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, pages 497–528. ACM, 2022.

[Almanza *et al.*, 2021] Matteo Almanza, Flavio Chierichetti, Silvio Lattanzi, Alessandro Panconesi, and Giuseppe Re. Online facility location with multiple advice. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 4661–4673, 2021.

[Aziz *et al.*, 2020] H. Aziz, H. Chan, B.E. Lee, B. Li, and T. Walsh. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In Vincent Conitzer and Fei Sha, editors, *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*. AAAI Press, 2020.

[Aziz *et al.*, 2021] Haris Aziz, Alexander Lam, Barton E. Lee, and Toby Walsh. Strategyproof and proportionally fair facility location. *CoRR*, abs/2111.01566, 2021.

[Aziz *et al.*, 2022] Haris Aziz, Alexander Lam, Barton E. Lee, and Toby Walsh. Strategyproof and proportionally fair facility location. In Kristoffer Arnsfelt Hansen, Tracy Xiao Liu, and Azarakhsh Malekian, editors, *Web and Internet Economics - 18th International Conference, WINE 2022, Troy, NY, USA, December 12-15, 2022, Proceedings*, volume 13778 of *Lecture Notes in Computer Science*, page 351. Springer, 2022.

[Chan *et al.*, 2021] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problems: A survey. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4356–4365. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Survey Track.

[Cheng and Zhou, 2015] Yukun Cheng and Sanming Zhou. A survey on approximation mechanism design without money for facility games. In David Gao, Ning Ruan, and Wenxun Xing, editors, *Advances in Global Optimization*, pages 117–128, Cham, 2015. Springer International Publishing.

[Escoffier *et al.*, 2011] B. Escoffier, L. Gourvès, N. Kim Thang, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In R.I. Brafman, F.S Roberts, and A. Tsoukiàs, editors, *Algorithmic Decision Theory*, pages 67–81, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[Fotakis and Tzamos, 2010] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In A. Saberi, editor, *Internet and Network Economics*, pages 234–245, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[Fotakis and Tzamos, 2013] D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. In F.V. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, editors, *Automata, Languages, and Programming*, pages 449–460, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[Fotakis *et al.*, 2021] Dimitris Fotakis, Evangelia Gergatsouli, Themis Gouleakis, and Nikolas Patris. Learning augmented online facility location. *CoRR*, abs/2107.08277, 2021.

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

[Han *et al.*, 2023] Yue Han, Christopher Jerrett, and Elliot Anshelevich. Optimizing multiple simultaneous objectives for voting and facility location. In Brian Williams, Yiling Chen, and Jennifer Neville, editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 5665–5672. AAAI Press, 2023.

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

[Jiang *et al.*, 2022] Shaofeng H.-C. Jiang, Erzhi Liu, You Lyu, Zhihao Gavin Tang, and Yubo Zhang. Online facility location with predictions. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.

[Lu *et al.*, 2010] P. Lu, X. Sun, Y. Wang, and Z.A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 315–324, New York, NY, USA, 2010. ACM.

[Mei *et al.*, 2016] Lili Mei, Minming Li, Deshi Ye, and Guochuan Zhang. Strategy-proof mechanism design for facility location games: Revisited (extended abstract). In Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls, editors, *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 1463–1464. ACM, 2016.

[Mei *et al.*, 2019] Lili Mei, Minming Li, Deshi Ye, and Guochuan Zhang. Facility location games with distinct desires. *Discrete Applied Mathematics*, 264:148–160, 2019.

[Mitzenmacher and Vassilvitskii, 2020] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.

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

[Procaccia and Tennenholtz, 2009] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In John Chuang, Lance Fortnow, and Pearl Pu, editors, *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6–10, 2009*, pages 177–186. ACM, 2009.

[Procaccia and Tennenholtz, 2013] A.D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. *ACM Trans. Econ. Comput.*, 1(4):18:1–18:26, December 2013.

[Tang *et al.*, 2020] Zhongzheng Tang, Chenhao Wang, Mengqi Zhang, and Yingchao Zhao. Mechanism design for facility location games with candidate locations. Technical report, CoRR archive within arXiv.org, Cornell University Library, 2020.

[Walsh, 2021] Toby Walsh. Strategy proof mechanisms for facility location at limited locations. In Duc Nghia Pham, Thanaruk Theeramunkong, Guido Governatori, and Fenrong Liu, editors, *PRICAI 2021: Trends in Artificial Intelligence - 18th Pacific Rim International Conference on Artificial Intelligence*, volume 13031 of *Lecture Notes in Computer Science*, pages 113–124. Springer, 2021.

[Walsh, 2024] Toby Walsh. Approximate mechanism design for facility location with multiple objectives. In *Proceedings of ECAI 2024*, Frontiers in Artificial Intelligence and Applications. IOS Press, 2024.

[Xu and Lu, 2022] Chenyang Xu and Pinyan Lu. Mechanism design with predictions. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 571–577. International Joint Conferences on Artificial Intelligence Organization, 7 2022.

[Zhang and Li, 2014] Q. Zhang and M. Li. Strategyproof mechanism design for facility location games with weighted agents on a line. *Journal of Combinatorial Optimization*, 28(4):756–773, Nov 2014.