

Facility Location Games with Scaling Effects

Yu He

City University of Hong Kong
Hong Kong, Hong Kong SAR
yuhe32-c@my.cityu.edu.hk

Alexander Lam

City University of Hong Kong
Hong Kong, Hong Kong SAR
alexlam@cityu.edu.hk

Minming Li

City University of Hong Kong
Hong Kong, Hong Kong SAR
minming.li@cityu.edu.hk

ABSTRACT

We take the classic facility location problem and consider a variation, in which each agent’s individual cost function is equal to their distance from the facility multiplied by a scaling factor which is determined by the facility placement. In addition to the general class of continuous scaling functions, we also provide results for piecewise linear scaling functions which can effectively approximate or model the scaling of many real world scenarios. We focus on the objectives of total and maximum cost, describing the computation of the optimal solution. We then move to the approximate mechanism design setting, observing that the agents’ preferences may no longer be single-peaked. Consequently, we characterize the conditions on scaling functions which ensure that agents have single-peaked preferences. Under these conditions, we find results on the total and maximum cost approximation ratios achievable by strategyproof and anonymous mechanisms.

KEYWORDS

Facility location; Single-peaked preferences; Scaled costs

ACM Reference Format:

Yu He, Alexander Lam, and Minming Li. 2024. Facility Location Games with Scaling Effects. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

The facility location problem is a widely studied problem, with decades of literature spanning the fields of operations research, economics, and computer science. In the classic variant of the problem, agents are located on a unidimensional interval, and we are tasked with finding an ideal location to place a facility which serves these agents. The agents have single-peaked preferences for the facility location, as they incur a cost equal to their distance from the facility. This problem models many real-world single-peaked social choice and preference aggregation problems. The geographical placement of public facilities such as libraries and parks is an immediate example [15], and the problem can also be applied to choose a political or economic outcome [5, 8], or to decide how to divide a budget [9]. Furthermore, the one-dimensional setting can be generalized to a multi-dimensional setting with an L_1 distance metric [11, 19], or extended to a network setting [22].




Figure 1: Facility location example where the agents (•) are located at $(0, 0, 0.6)$, and the scaling function (represented by the red lines) is $2 - 3y$ for $y \leq 0.5$, and $3y - 1$ for $y > 0.5$. The median and midpoint facility placements are at $y_{med} = 0$ and $y_{mid} = 0.3$ respectively, whilst the optimal facility placements in terms of total and maximum cost are at $y_{TC}^* = y_{MC}^* = 0.5$.

In our setting, we assume that the agents’ locations are private information, and take a strategyproof mechanism design approach to the classic facility location problem. We also add an additional dimension of complexity: each agent’s cost is scaled by an external factor corresponding to the facility’s location on the domain. Specifically, there is a (continuous) *scaling function* which maps the facility location to a positive scaling factor which is multiplied by each agent’s distance to calculate their individual costs. A local minimum of the scaling function implies that the facility is particularly effective when placed at this point. The primary motivation for our model is the placement of a cell-phone tower, which is most effective when it is on top of a hill, and less effective as it is placed lower on the hill. Each signal broadcasted or received by the radio tower has a failure probability depending on its elevation, and this averages out over time to be the facility’s effectiveness in a multiplicative sense. Other examples include the proximity of the facility to a river or public transportation, or of a warehouse to a shipping port. Similarly, placing the facility at a local maximum is relatively ineffective (such as placing the cell phone tower at the bottom of a valley, or a public facility next to a garbage dump), and the scaling function improves up to an extent as the facility moves away from this point.

In our paper, we consider the general class of continuous scaling functions, as well as piecewise linear scaling functions due to their simplicity in modelling real-world problems, and because they can approximate more complex scaling functions. To demonstrate our model, consider the example in Figure 1, where there are two agents at 0 and one agent at 0.6, and the scaling function is $2 - 3y$ for $y \leq 0.5$, and $3y - 1$ for $y > 0.5$. Under the classic model, the median and midpoint mechanisms are known to minimize total and maximum cost, respectively. In our example, they place the facility at 0 and 0.3 respectively, resulting in $0.6 \times 2 = 1.2$ total cost and 0.33 maximum cost. However, placing the facility at the scaling function’s minimum of 0.5 is optimal for both objectives, leading to 0.55 total cost and 0.25 maximum cost.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Table 1: Phantom mechanism approximation ratio results for total and maximum cost. Agents have single-peaked preferences when the scaling function q meets certain conditions, resulting in phantom mechanisms being strategyproof. The term r_q denotes the ratio between the max. and min. values of q , and k denotes the number of line segments in the piecewise linear q . All results are tight in the sense that they are a lower bound for all phantom mechanisms, and that there exists a phantom mechanism with a matching approximation ratio.

	General Pref.	Single-Peaked Pref.	
	Continuous and Piecewise Lin. q	Continuous q	Piecewise Linear q
Total Cost	r_q (Thm. 5.6)	e (Cor. 5.8)	$(1 + \frac{1}{k})^k$ (Cor. 5.9)
Max. Cost	$2r_q$ (Thm. 5.11)	$2e$ (Cor. 5.12)	$2(1 + \frac{1}{k})^k$ (Cor. 5.13)

The goal of our paper is to design strategyproof mechanisms for our scaled facility location problem, that have ideal approximations of the optimal total or maximum cost. These two objectives are standard measures of efficiency and egalitarian fairness in facility location mechanism design [20]. A key consequence of the scaling function is that the agents’ preferences may no longer be single-peaked, and thus we cannot immediately leverage existing characterizations of strategyproof mechanisms under single-peaked preferences (e.g., [5, 16]). We therefore consider the sufficient and necessary conditions on the scaling function such that the agents’ preferences are single-peaked, and design reasonable strategyproof mechanisms for such scaling functions.

Contributions.

- We initiate the study of facility location mechanism design where the facility has a degree of effectiveness depending on where it is placed, and the agents’ costs are scaled by this level of effectiveness. Our results focus on continuous scaling functions, as well as the smaller space of piecewise linear scaling functions.
- We give computational results on the optimal facility placement for the objectives of total cost and maximum cost.
- We show that the agents’ preferences may not be single-peaked in our setting, and characterize the conditions on scaling functions such that agents have single-peaked preferences.
- When scaling functions meet the conditions for agents to have single-peaked preferences, we find the lower bound on the approximation ratios achieved by any strategic and anonymous mechanism, and propose such mechanisms with approximation ratios matching this lower bound. Table 1 summarizes some of our approximation ratio results.

Related Work. The facility location problem can be interpreted as a special case of preference aggregation when agents have symmetric and single-peaked preferences over a continuous range of outcomes, in which there has been decades of research in strategyproof

aggregation rules. For general preferences, the famous Gibbard-Satterthwaite theorem implies that a surjective and strategyproof rule must be dictatorial [10, 21]. This impossibility result vanishes when we consider the restricted domain of single-peaked preferences, in which there have been numerous characterizations of strategyproof rules. Most notably, under this restricted domain, Moulin [16]’s seminal paper characterizes the space of strategyproof and anonymous rules as Phantom mechanisms which output the median of the agent locations and $n+1$ constant ‘Phantom’ values. This characterization was later extended to the domain of symmetric and single-peaked preferences by Border and Jordan [5] and Massó and Moreno De Barreda [14]. Other characterizations of strategyproof rules have been proposed for similar settings by [4, 12, 17, 18].

The paper most similar to our work considers the facility location model with entrance fees [13], in which agents incur, in addition to their distance from the facility, an additive cost which depends on the facility placement. In this model, the authors give strategyproof mechanisms and compute their approximation ratios with respect to the total and maximum cost objectives. The key difference with our model is that our agents experience a multiplicative scaling factor to their cost, whilst their agents experience an additional additive cost. Our work and [13] are inspired by the line of research on approximate mechanism design without money, initially proposed for the facility location problem by Procaccia and Tennenholtz [20]. In their paper, the authors compute the worst-case ratio between the performance of strategyproof mechanisms and welfare-optimal mechanisms. There have since been many papers applying this approach to variations of the facility location problem, such as when there are proportionally fair distance constraints [3], when facilities have capacity limits [1, 2], or when the facility is obnoxious (i.e., agents prefer to be far away from the facility) [7]. A recent survey of related work is given by Chan et al. [6].

2 MODEL

We have a set of agents $N = \{1, \dots, n\}$, where agent i has location x_i on the domain¹ $X := [0, 1]$. Although we consider the space of potentially non-anonymous mechanisms, we assume for simplicity that agent locations are ordered such that $x_1 \leq \dots \leq x_n$. This does not affect the nature of our results. A scaling function $q : X \rightarrow \mathbb{R}_{>0}$ gives the effectiveness of a facility.² We also assume that the scaling function is continuous. As we will show in Proposition 3.1, continuity of q is required for the optimal facility location to be well-defined.

Denoting \mathcal{Q} as the space of all scaling functions, facility location mechanism $f : \mathcal{Q} \times X^n \rightarrow X$ maps the agent location profile $\mathbf{x} = (x_1, \dots, x_n)$ to the location of a facility y . We define the cost incurred by agent i as its distance from the facility multiplied by the scaling factor: $c_i(q, y) := q(y)|y - x_i|$. Finally, we define the total cost of an instance as $TC(q, y, \mathbf{x}) := \sum_i c_i(q, y) = q(y) \sum_i |y - x_i|$, and the maximum cost of an instance as $MC(q, y, \mathbf{x}) := \max_i c_i(q, y) = q(y) \max_i |y - x_i|$. Respectively, we denote the optimal facility location which minimizes the total cost (resp. maximum cost) as y_{TC}^* (resp. y_{MC}^*).

¹Although we consider the unit interval domain, our results can be scaled and shifted to any compact interval on \mathbb{R} . We consider a bounded domain so that a single-peaked linear scaling function does not take negative values.

²If there is a point y where $q(y) = 0$, then it is trivial to place the facility at that point.

It is typically ideal for the mechanism output to be independent of the agents' labelling, so we are concerned with mechanisms that satisfy *anonymity*.

Definition 2.1. A mechanism f is *anonymous* if its output does not change when the agents' labels are permuted.

3 COMPUTING THE OPTIMAL SOLUTION

In this section, we investigate the properties of the optimal facility locations for total cost y_{TC}^* and maximum cost y_{MC}^* . We begin by justifying our assumption that the scaling function q must be continuous, as a discontinuity may result in the optimal solution being not well-defined.

PROPOSITION 3.1. *There exists a scaling function with a discontinuity and an agent location profile such that the optimal facility locations for total cost y_{TC}^* and maximum cost y_{MC}^* are not well-defined.*

PROOF. Consider the location profile $\mathbf{x} = (0, 0.5, 1)$ and the scaling function $q(y) = \begin{cases} 1, & y \in [0, 0.5], \\ 0.5, & y \in (0.5, 1]. \end{cases}$ The optimal facility locations for total cost and maximum cost are the right limit $y^* \rightarrow 0.5^+$, but not at $y^* = 0.5$, so the optimal solutions are not well-defined. \square

3.1 Total Cost

We now give computational results on the optimal facility location for total cost, beginning with the class of linear scaling functions.

PROPOSITION 3.2. *If q is a linear function, the optimal facility location for total cost y_{TC}^* is either on one of the agents' locations, or the minimum value of q .*

PROOF. Since q is a linear function, it is of the form $q(y) = ay + b$, where $a \in \mathbb{R}$ and $b \in \mathbb{R}_{>0}$. If $a = 0$, then $q(y)$ is constant, and thus the optimal facility location is on the median agent.³

Due to symmetry, it suffices to consider the case where the scaling function is increasing, i.e., $a > 0$. From this assumption, we know that the optimal solution cannot lie to the right of x_n .

Since q is increasing with y , its minimum value is on $y = 0$. We denote $x_0 := 0$, and we know the optimal solution y^* is in the interval $[x_0, x_n]$. Suppose the optimal solution y^* is in the interval $[x_{n_1}, x_{n_1+1}]$, where $n_1 \in [0, n-1]$ denotes the number of agents on the left-hand side of y^* . Also, let $n_2 = n - n_1$ denote the number of agents on the right-hand side of y^* .

Recall that the total cost corresponding to y^* , $TC(q, y^*, \mathbf{x})$ is $q(y^*) \sum_i |y^* - x_i|$. Considering the term $\sum_i |y^* - x_i|$, we have

$$\begin{aligned} \sum_{i=1}^n |y^* - x_i| &= \sum_{i=1}^{n_1} (y^* - x_i) + \sum_{i=n_1+1}^n (x_i - y^*) \\ &= (n_1 - n_2)y^* + \sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i. \end{aligned}$$

We remark that $\sum_i |y^* - x_i|$ can also be written as a linear function of y^* : $\sum_i |y^* - x_i| = a'y^* + m$ with $a' := n_1 - n_2$, and $m := \sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i$.

³If there are an even number of agents, we break the tie in favour of the leftmost agent in the optimal interval.

$\sum_{i=1}^{n_1} x_i$. We can therefore express the total cost as the following function of y^* .

$$\begin{aligned} TC(q, y^*, \mathbf{x}) &= q(y^*) \cdot \sum_i |y^* - x_i| \\ &= (ay^* + b)(a'y^* + m) \\ &= aa'(y^*)^2 + (am + a'b)y^* + bm. \end{aligned}$$

Now we know that the total cost function can either be a linear function or a quadratic function of y^* , depending on whether $a' = 0$, $a' < 0$, or $a' > 0$ (recalling that $a > 0$). There are three cases: either the total cost function is a linear function ($a' = 0$), a quadratic function which opens downwards ($a' < 0$), or a quadratic function which opens upwards ($a' > 0$). For the first two cases, it is clear that the minimum value of the total cost function must lie on an endpoint of the interval $[x_{n_1}, x_{n_1+1}]$. It remains to consider the last case ($a' > 0$), where the total cost function $q(y^*) \sum_i |y^* - x_i| = aa'y^* + (am + a'b)y^* + bm$ opens upwards. Consider the total cost function on the larger domain of \mathbb{R} . Since $a' = (n_1 - n_2) > 0$, we know there exists $z_1 \in \mathbb{R}$ such that $a'z_1 + m = 0$. Also, since $a > 0$, there exists $z_2 \in \mathbb{R}$ such that $q(z_2) = 0$. As a result, we see that the total cost function is equal to zero at either one or two points on \mathbb{R} ; we now show that both of these points lie to the left of x_{n_1} .

For the term $a'y^* + m$, we know that $a'z_1 + m = 0$, and thus $z_1 = -\frac{m}{a'} = \frac{\sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i}{n_2 - n_1}$. Consider the numerator $\sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i$. We know there are n_2 agents between $(x_{n_1}, x_n]$, implying $\sum_{i=n_1+1}^n x_i > n_2 \cdot x_{n_1}$. Similarly, we know there are n_1 agents between $[0, x_{n_1}]$, implying $\sum_{i=1}^{n_1} x_i \leq n_1 \cdot x_{n_1}$. Combining these inequalities gives us

$$\sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i > n_2 x_{n_1} - n_1 x_{n_1} > 0.$$

As we have assumed $a' > 0$, we know $n_2 - n_1 < 0$ and thus $z_1 = \frac{\sum_{i=n_1+1}^n x_i - \sum_{i=1}^{n_1} x_i}{n_2 - n_1} < x_{n_1}$.

Now consider z_2 , which satisfies $q(z_2) = 0$. Clearly, $z_2 < 0$ as we must have $q(y) > 0$ for $y \in [0, 1]$ and $a > 0$, implying $z_2 < x_{n_1}$.

Since $z_1, z_2 < x_{n_1}$, the total cost function must be strictly increasing in $[x_{n_1}, x_{n_1+1}]$, so its minimum value must lie on x_{n_1} , which is either an agent location or the minimum value of q . \square

We can generalise this result to continuous, *piecewise* linear scaling functions.

THEOREM 3.3. *Let the scaling function q be a continuous, piecewise linear function. The facility location minimizing total cost y_{TC}^* is either on one of the agents' locations or on a local minimum of q .*

PROOF. First, consider each line segment on q , which is represented by a linear function q_j on the sub-domain $[c_j, d_j]$, where $c_j, d_j \in X$. Let $y_j^* := \arg \min_{y_j \in [c_j, d_j]} TC(q, y_j, \mathbf{x})$ denote the optimal facility placement on $[c_j, d_j]$, and $e_j := \arg \min_{y \in [c_j, d_j]} q(y)$ be the minimum point of q on $[c_j, d_j]$. By Proposition 3.2, we must have $y_j^* \in \{e_j, x_1, \dots, x_n\}$. Therefore, considering all of the q_j functions, the optimal solution y^* must lie at one of the agents' positions or at a local minimum of q_j . \square

COROLLARY 3.4. *When q is a continuous, piecewise linear function with a constant number of local minima, the optimal facility location*

can be found in linear time by simply iterating over the agent locations and the local minima of the scaling function.

As we assume agent locations are private information, there is a concern that agents may misreport their locations to unfairly attain a better facility placement. It is therefore ideal to implement a *strategyproof* mechanism, which does not incentivize agents to lie about their locations.

Definition 3.5 (Strategyproofness). A mechanism $f(\cdot)$ is *strategyproof* if for every agent $i \in N$, we have, for every scaling function q and agent locations x'_i, \mathbf{x}_{-i} and x_i ,

$$c_i(q, f(x_i, \mathbf{x}_{-i})) \leq c_i(q, f(x'_i, \mathbf{x}_{-i})).$$

However, we find that the optimal solution is not strategyproof.

PROPOSITION 3.6. *The optimal facility placement which minimizes total cost is not strategyproof, even if the scaling function is a linear function and there are only 3 agents.*

PROOF. Suppose the scaling function is $q(y) = 1 - \frac{4}{5}y$ and the agent location profile is $x = (0.3, 0.4, 0.7)$. The optimal facility placement is $y = 0.4$, which leads to a cost of 0.204 for agent x_3 . However if agent x_3 misreports $x'_3 = 0.755$, then the facility placement is instead $y' = 1$, which leads to a cost of 0.06. Hence the optimal facility placement is not strategyproof. \square

3.2 Maximum Cost

We now consider the optimal facility location for maximum cost. It is well-known that in the classic setting, placing the facility at the midpoint of the leftmost and rightmost agents is optimal for maximum cost but not strategyproof [20]. This implies that y_{MC}^* is not strategyproof in our setting. However, it is easily computable in constant time when q is a linear function, as there are only three possible solutions.

PROPOSITION 3.7. *If q is a linear function, the optimal facility location for maximum cost satisfies $y_{MC}^* \in \{0, \frac{x_1+x_n}{2}, 1\}$.*

PROOF. Suppose q is of the form $q(y) = ay + b$, and recall that the maximum cost of the instance is $MC(q, y, \mathbf{x}) = \max_i c_i(q, y) = \max_i (ay + b)|y - x_i|$. Due to symmetry it suffices to consider $a > 0$.

Since $a > 0$, we know that $y_{MC}^* \notin (x_n, 1]$, as such a facility placement has a strictly higher cost for all agents than a facility on x_n . We also know that $y_{MC}^* \notin (\frac{x_1+x_n}{2}, x_n]$. This is because for every facility placement on this interval, there is a ‘symmetric’ facility placement $y \in [x_1, \frac{x_1+x_n}{2})$ where the maximum (unscaled) distance is the same, but the scaling function value is lower. We therefore know that $y_{MC}^* \in [0, \frac{x_1+x_n}{2}]$. When the facility satisfies $y \leq \frac{x_1+x_n}{2}$, then we also know that agent n incurs the highest cost, as it has the highest unscaled distance from the facility.

We now know that if $a > 0$, the maximum cost is $MC(q, y, \mathbf{x}) = c_n(q, y) = (ay + b)(x_n - y)$, which is a quadratic function of y that opens downwards. Therefore the maximum cost is minimized when $y = 0$ or $y = \frac{x_1+x_n}{2}$. By symmetry, if $a < 0$, agent 1 incurs the maximum cost of $MC(q, y, \mathbf{x}) = c_1(q, y) = (ay + b)(y - x_1)$, which is minimized when $y = 1$ or $y = \frac{x_1+x_n}{2}$. We therefore have $y_{MC}^* \in \{0, \frac{x_1+x_n}{2}, 1\}$. \square

We can also similarly show that when computing y_{MC}^* for continuous scaling functions, we only need to consider the optimal facility location with respect to the leftmost and rightmost agents.

PROPOSITION 3.8. *For continuous scaling functions, the optimal facility location for maximum cost satisfies*

$$y_{MC}^* \in \left\{ \arg \min_{y \in [0, \frac{x_1+x_n}{2}]} c_n(q, y), \arg \min_{y \in [\frac{x_1+x_n}{2}, 1]} c_1(q, y), \frac{x_1 + x_n}{2} \right\}.$$

PROOF. If $y \in [0, \frac{x_1+x_n}{2}]$, then agent n receives the highest cost in the instance, similar to the proof of Proposition 3.7. A symmetric argument shows that if $y \in [\frac{x_1+x_n}{2}, 1]$, then agent 1 receives the highest cost in the instance. Therefore y_{MC}^* is either on the minimum value of $c_n(q, y)$ on $y \in [0, \frac{x_1+x_n}{2}]$, the minimum value of $c_1(q, y)$ on $y \in [\frac{x_1+x_n}{2}, 1]$, or on $\frac{x_1+x_n}{2}$. \square

We next investigate whether there exist non-trivial strategyproof mechanisms for this setting.

4 ACHIEVING SINGLE-PEAKED PREFERENCES

When agents have single-peaked preferences along a real line domain, Moulin [16] characterized the set of anonymous and strategyproof mechanisms as *phantom mechanisms*, which place the facility at the median of the n agent locations and $n + 1$ constant ‘phantom’ points.

Definition 4.1 (Phantom Mechanism). Given \mathbf{x} and $n + 1$ constant values $0 \leq p_1 \leq p_2 \leq \dots \leq p_{n+1} \leq 1$, a *phantom mechanism* places the facility at $\text{med}\{x_1, \dots, x_n, p_1, \dots, p_{n+1}\}$.

However, we find that in our setting, phantom mechanisms are not necessarily strategyproof. Even the well-studied median mechanism, a phantom mechanism with $\lceil \frac{n+1}{2} \rceil$ phantoms at 0 and $\lfloor \frac{n+1}{2} \rfloor$ phantoms at 1, is not strategyproof.

PROPOSITION 4.2. *The median mechanism is not strategyproof.*

PROOF. Consider the agent location profile $\mathbf{x} = (0.5, 0.7, 0.9)$ and the scaling function $q(y) = \begin{cases} 6 - 10y, & y \leq 0.5, \\ 10y - 4, & y > 0.5. \end{cases}$

The median mechanism places the facility at $y = 0.7$, and it leads to a cost of 0.6 for agent x_3 . However if the agent at x_3 misreports to $x'_3 = 0.5$, then the median facility placement will be 0.5, which leads to a lower cost of 0.4 for the agent at x_3 . Hence the median mechanism is not strategyproof. \square

In fact, if there are at least two unique phantom locations, the phantom mechanism is not strategyproof.

THEOREM 4.3. *Every phantom mechanism with $n + 1$ phantoms on at least two unique phantom locations is not strategyproof.*

PROOF. Suppose the $n+1$ phantoms are located on $k \in \{2, \dots, n+1\}$ unique points (a_1, a_2, \dots, a_k) . We consider an agent profile such that all the agents are in the interval $[a_1, a_2]$. Denoting the number

of phantoms on a_1 as n^* , the facility is placed on the $(n+1-n^*)$ th agent under this profile. Now suppose the scaling function is

$$q(y) = \begin{cases} 3 - \frac{a_1+1}{a_2+1} - \frac{4}{a_1+a_2}y, & y \in [0, \frac{a_1+a_2}{2}], \\ \frac{8}{a_2+1}y - \frac{5a_1+3a_2}{a_2+1}, & y \in [\frac{a_1+a_2}{2}, 1]. \end{cases}$$

Case 1 ($n+1-n^* \in \{2, \dots, n\}$): With all agents located in $[a_1, a_2]$, we also specifically let $x_{n-n^*} = a_1$, $x_{n+1-n^*} = \frac{3a_1+a_2}{4}$, and $x_{n+2-n^*} = \frac{a_1+a_2}{2}$. In such case, the mechanism places the facility at x_{n+1-n^*} , and the agent at x_{n-n^*} incurs cost

$$\begin{aligned} c_{(n-n^*)}(q, f(\mathbf{x})) &= q\left(\frac{3a_1+a_2}{4}\right) \cdot \left|\frac{3a_1+a_2}{4} - a_1\right| \\ &= \left(3 - \frac{a_1+1}{a_2+1} - \frac{4}{a_1+a_2} \cdot \frac{3a_1+a_2}{4}\right) \left(\frac{a_2-a_1}{4}\right) \\ &= \left(1 - \frac{1}{2} \left(\frac{a_1+1}{a_2+1} + \frac{2a_1}{a_1+a_2}\right)\right) \left(\frac{a_2-a_1}{2}\right) \\ &> \left(1 - \frac{a_1+1}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right). \end{aligned}$$

The last inequality is due to $\frac{a_1+a_1}{a_2+a_1} < \frac{a_1+1}{a_2+1}$. If x_{n-n^*} misreports to $x'_{n-n^*} = \frac{a_1+a_2}{2}$, then $f(\mathbf{x}') = \frac{a_1+a_2}{2}$, and their new cost becomes $c_{(n-n^*)}(q, f(\mathbf{x}')) = \left(1 - \frac{a_1+1}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right) < c_{(n-n^*)}(q, f(\mathbf{x}))$.

Case 2 ($n+1-n^* = 1$): Consider the agent location profile \mathbf{x} where $x_1 = \frac{a_1+3a_2}{4}$ and $x_n = a_2$. Here, we have $f(\mathbf{x}) = x_{n+1-n^*} = x_1 = \frac{a_1+3a_2}{4}$. Therefore the cost for agent x_n is

$$\begin{aligned} c_n(q, f(\mathbf{x})) &= q\left(\frac{a_1+3a_2}{4}\right) \cdot \left|\frac{a_1+3a_2}{4} - a_2\right| \\ &= \left(\frac{8}{a_2+1} \cdot \frac{a_1+3a_2}{4} - \frac{5a_1+3a_2}{a_2+1}\right) \left|\frac{a_1+3a_2}{4} - a_2\right| \\ &= \left(\frac{3a_2-3a_1}{a_2+1}\right) \left(\frac{a_2-a_1}{4}\right) \\ &= \frac{3}{2} \cdot \left(1 - \frac{a_1+1}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right) \\ &> \left(1 - \frac{a_1+1}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right). \end{aligned}$$

However if x_n misreports to $x'_n = \frac{a_1+a_2}{2}$, then $f(\mathbf{x}') = \frac{a_1+a_2}{2}$. Agent x_n will then incur cost

$$\begin{aligned} c_n(q, f(\mathbf{x}')) &= q\left(\frac{a_1+a_2}{2}\right) \cdot \left|\frac{a_1+a_2}{2} - a_2\right| \\ &= \left(\frac{8}{a_2+1} \cdot \frac{a_1+a_2}{2} - \frac{5a_1+3a_2}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right) \\ &= \left(1 - \frac{a_1+1}{a_2+1}\right) \left(\frac{a_2-a_1}{2}\right) \\ &< c_n(q, f(\mathbf{x})). \end{aligned}$$

Therefore, a phantom mechanism which does not place all phantoms at the same position is not strategyproof. \square

From this theorem, we arrive at the following corollary. Note that the constant mechanism is equivalent to the phantom mechanism with all $n+1$ phantoms at the same location.

COROLLARY 4.4. *The constant mechanism (which places the facility at some constant $c \in X$) is the only anonymous and strategyproof mechanism in our setting.*

This is because the scaling function may cause the agent's preference to no longer be single-peaked. To show this, we first define each agent's utility as the negative of their cost ($u_i(q, y) := -c_i(q, y)$), so that their preference (with respect to cost) is single-peaked if and only if their utility is single-peaked. We therefore use the terms 'utility functions' and 'preferences' interchangeably throughout the paper.

THEOREM 4.5. *An agent's preference may not be single-peaked, even if the scaling function is a linear function.*

PROOF. Consider the scaling function $q(y) = y + 0.01$ and $x_1 = 0.5$. The derivative of agent 1's utility is $u'_1(q, y) = 2y - 0.49$ when $y < 0.5$, and $u'_1(q, y) = -(2y - 0.49)$ when $y > 0.5$. For $y \in [0, 0.245]$, $u'_1(q, y) \leq 0$. For $y \in [0.245, 0.5]$, $u' \geq 0$. Lastly, for $y \in [0.5, 1]$, $u'_1(q, y) < 0$. Therefore $u_1(q, y)$ is not single-peaked. \square

However, we can characterize the conditions on the scaling functions which ensure that agents have single-peaked preferences. These conditions hold as long as q is continuous, even if it is not differentiable at a countable set of points D .

THEOREM 4.6. *If q is continuous but not differentiable at a countable set of points $D := \{j | q'(j) \text{ does not exist}\}$, the agents' preferences are guaranteed to be single-peaked if and only if $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$.*

PROOF. Recall that the utility function for agent x_i is

$$u_i(q, y) = \begin{cases} q(y)(y - x_i), & y \leq x_i, \\ q(y)(x_i - y), & y \geq x_i. \end{cases}$$

The derivative of u with respect to y is

$$u'_i(q, y) = \begin{cases} q(y) + q'(y)(y - x_i), & y \in [0, x_i] \setminus D, \\ \text{Not differentiable}, & y \in D \cup \{x_i\}, \\ -(q(y) + q'(y)(y - x_i)), & y \in (x_i, 1] \setminus D. \end{cases}$$

(\implies) The utility function is single-peaked, so for each agent i , there exists a value $k_i \in (0, 1)$ such that $u'_i(q, y) \geq 0$ with $y \in [0, k_i]$, and $u'_i(q, y) \leq 0$ with $y \in (k_i, 1]$. Since agents receive zero cost when the facility is placed at their location, we know the utility function's 'peak' lies on their location, and therefore $k_i = x_i$.

As a result, we have $q(y) + q'(y)(y - x_i) \geq 0$ for $y \in [0, x_i] \setminus D$, and $-q(y) + q'(y)(y - x_i) \leq 0$ for $y \in (x_i, 1] \setminus D$. This can also be written as $q(y) + q'(y)(y - x_i) \geq 0$ for all $y \in [0, 1] \setminus (\{x_i\} \cup D)$. Having $x_i \in [0, 1]$ implies that $(y - x_i) \in [-1, 1]$, meaning that our condition is $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$.

(\impliedby) Suppose that a scaling function q satisfies $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$. We know that for every agent x_i , $q(y) + q'(y)(y - x_i) \geq 0$ for all $y \in [0, 1] \setminus (\{x_i\} \cup D)$. The utility function is increasing for $y < x_i$, and decreasing for $y > x_i$. \square

By using this theorem, we can characterize the conditions on linear scaling functions which result in agents having single-peaked preferences.

COROLLARY 4.7. *If q is a linear function of the form $q(y) = ay + b$, the agents' preferences are guaranteed to be single-peaked if and only if $b \geq a$ for $a > 0$, and $b \geq -2a$ for $a < 0$.*

We also observe that a sufficiently large constant c can be added to any continuous scaling function to ensure that every agent has a single-peaked preference. This can be explained by the scaling function becoming ‘closer’ to a constant scaling function.

We can also apply Theorem 4.6 to piecewise linear scaling functions to characterize the conditions for agents to have single-peaked preferences.

PROPOSITION 4.8. *If q is a continuous piecewise linear function where each line segment is of the form $q_j(y) = a_j y + b_j$, the agents’ preferences are guaranteed to be single-peaked if and only if for each line segment, $b_j \geq -2a_j$ when $a_j < 0$, and $b_j \geq a_j$ when $a_j > 0$.*

PROOF. From Theorem 4.6, we know that if the q function is continuous but not differentiable at some set of points D , the utility function is single-peaked if and only if $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$. When $a_j < 0$, the utility function is single-peaked if and only if $a_j y + b_j + a_j \geq 0 \Rightarrow b_j \geq -2a_j y$, and when $a_j > 0$, the utility function is single-peaked if and only if $a_j y + b_j - a_j \geq 0 \Rightarrow b_j \geq a_j(y-1)$. Since $y \in [0, 1]$, these conditions become $b_j \geq -2a_j$ and $b_j \geq a_j$, respectively. \square

5 APPROXIMATION RATIO RESULTS

In this section, we give results on the approximation ratio of certain mechanisms. The approximation ratio of a mechanism quantifies its worst-case performance for a specified objective function.

Definition 5.1 (Total Cost Approximation Ratio). Given a mechanism f , its total cost approximation ratio is

$$\max_{q \in \mathcal{Q}, \mathbf{x} \in X^n} \frac{TC(q, f(\mathbf{x}), \mathbf{x})}{TC(q, y_{TC}^*, \mathbf{x})}.$$

The maximum cost approximation ratio is similarly defined by replacing the total cost expressions $TC(q, \cdot, \mathbf{x})$ and y_{TC}^* with the corresponding maximum cost expressions $MC(q, \cdot, \mathbf{x})$ and y_{MC}^* .

While we have defined the approximation ratio for the general space of scaling functions \mathcal{Q} , we additionally prove approximation ratio bounds for more restricted classes of scaling functions, such as piecewise linear scaling functions, and scaling functions such that agents are guaranteed to have single-peaked preferences.

For simplicity of notation, we let $r_q := \frac{\max_{y \in X} q(y)}{\min_{y \in X} q(y)}$ be the ratio between the highest and lowest values of the scaling function.

We begin with an impossibility result in the general continuous setting, which explains our need to restrict the scaling function parameters to ensure single-peakedness and widen the space of strategyproof and anonymous mechanisms.

THEOREM 5.2. *Every strategyproof and anonymous mechanism has an unbounded approximation ratio for total and maximum cost.*

PROOF. By Moulin [16]’s characterization, a strategyproof and anonymous mechanism in our setting must be a phantom mechanism, as the classic setting is a special case of our setting. By Theorem 4.3, we also know that a phantom mechanism is only strategyproof if all $n + 1$ phantoms are located at the same location. Such a phantom mechanism is equivalent to the constant mechanism, which has an unbounded approximation ratio as it is not unanimous. \square

As we will show, our approximation ratio results are typically a function of r_q . We show that when the scaling function is chosen such that agents have single-peaked preferences, the value of r_q is upper bounded by a constant.

THEOREM 5.3. *When agents’ preferences are single-peaked, $r_q \leq e$.*

PROOF. From Theorem 4.6, we know that $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$, where D denotes the points where $q(y)$ is not differentiable.

First, consider the intervals where $q'(y) > 0$, and consequently q satisfies $q(y) - q'(y) \geq 0$. We define a function $g(y) := e^y q(y)$, where $g'(y) = e^y(q(y) - q'(y))$. We know $g(y)$ is non-decreasing because $q(y) - q'(y) \geq 0$. Consider the points $y_1, y_2 \in [0, 1]$ where $y_1 < y_2$. Since $g(y_1) \leq g(y_2)$, we know that $e^{y_1} q(y_1) \leq e^{y_2} q(y_2)$. Also, since $q(y) > 0$ and $e^y > 0$, we have $\frac{q(y_1)}{q(y_2)} \leq e^{y_2 - y_1}$, which is maximized when $y_1 = 0$ and $y_2 = 1$. Therefore $r_q \leq e$. Furthermore, the function $q(y) = e^y$ satisfies $q(y) - |q'(y)| \geq 0$ and has $r_q = e$.

Lastly, consider the intervals where $q'(y) < 0$, and consequently q satisfies $q(y) + q'(y) \geq 0$. Similarly, we define $g(y) := e^{-y} q(y)$, where $g'(y) = e^{-y}(q'(y) - q(y))$. We know $g(y)$ is non-increasing because $q'(y) - q(y) \leq 0$. Consider the points $y_1, y_2 \in [0, 1]$ where $y_1 < y_2$. Since $g(y_1) \geq g(y_2)$, we know that $e^{-y_1} q(y_1) \geq e^{-y_2} q(y_2)$. Also, since $q(y) > 0$ and $e^y > 0$, we have $\frac{q(y_2)}{q(y_1)} \leq e^{y_2 - y_1}$, which is again maximized when $y_1 = 0$ and $y_2 = 1$. Therefore $r_q \leq e$. \square

When q is a piecewise linear function, the upper bound of r_q is dependent on how many line segment pieces the function has.

THEOREM 5.4. *Suppose the scaling function q is a piecewise linear function with k line segment pieces. When agents’ preferences are single-peaked, $r_q \leq (1 + \frac{1}{k})^k$.*

PROOF. Suppose for $j \in \{1, \dots, k\}$ that each line segment of q is denoted as $q_j(y) = a_j y + b_j$. From Theorem 4.6, for agents to always have single-peaked preferences, we require that $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus \{D\}$, where $\{D\}$ denotes the set of points where $q(y)$ is not differentiable. This implies that for each line segment, if $a_j > 0$, then $a_j y + b_j - a_j > 0 \Rightarrow b_j \geq (1 - y)a_j$. Similarly, if $a_j < 0$, then $b_j \geq -(1 + y)a_j$.

We denote the endpoints of the scaling function’s line segment pieces as the set $S = \{0, s_1, s_2, \dots, s_{k-1}, 1\}$. Consider a line segment piece $q_j = a_j y + b_j$, where $y \in [s_{j-1}, s_j]$. We now find an expression for r_{q_j} , which is the ratio between the maximum and minimum values of the scaling function on $[s_{j-1}, s_j]$.

If $a_j > 0$, then $r_{q_j} = \frac{q_j(s_j)}{q_j(s_{j-1})} = \frac{a_j s_j + b_j}{a_j s_{j-1} + b_j}$. Since we require $b_j \geq (1 - y)a_j$ on $y \in [s_{j-1}, s_j]$, we know that $b_j \geq (1 - s_{j-1})a_j$. Therefore $r_{q_j} = \frac{a_j s_j + b_j}{a_j s_{j-1} + b_j} \leq \frac{a_j s_j + (1 - s_{j-1})a_j}{a_j s_{j-1} + (1 - s_{j-1})a_j} = (1 + s_j - s_{j-1})$.

If $a_j < 0$, then $r_{q_j} = \frac{q_j(s_{j-1})}{q_j(s_j)} = \frac{a_j s_{j-1} + b_j}{a_j s_j + b_j}$. Since we require $b_j \geq -(1 + y)a_j$ on $y \in [s_{j-1}, s_j]$, we know that $b_j \geq -(1 + s_j)a_j$. Therefore $r_{q_j} = \frac{a_j s_{j-1} + b_j}{a_j s_j + b_j} \leq \frac{a_j s_{j-1} - (1 + s_j)a_j}{a_j s_j - (1 + s_j)a_j} = (1 + s_j - s_{j-1})$, and we know that for each q_j , $r_{q_j} \leq (1 + s_j - s_{j-1})$.

Since q is piecewise linear, its global maxima and minima must be elements of S . We take an arbitrary global minimum $s_{min} := \arg \min_{y \in X} q(y)$ and global maximum $s_{max} := \arg \max_{y \in X} q(y)$, where $s_{min}, s_{max} \in S$, and suppose w.l.o.g. that $s_{min} < s_{max}$. It

also suffices to assume that for each line segment q_j between s_{min} and s_{max} , we have $a_j > 0$ (i.e. the function is strictly increasing in $[s_{min}, s_{max}]$). This means that r_q is the product of the r_{q_j} between s_{min} and s_{max} . Finally, since $1 \leq r_{q_j} \leq (1 + s_j - s_{j-1})$, and applying the arithmetic mean-geometric mean inequality, we have $r_q \leq \prod_{i=1}^k (1 + s_i - s_{i-1}) \leq (1 + \frac{1}{k})^k$. \square

When the amount of line segments k approaches infinity, the upper bound of r_q asymptotically meets the upper bound of $r_q = e$ for general scaling functions. Intuitively, this can also be seen by constructing a piecewise linear function which approximates the function $q(y) = e^y$.

We follow up with another result which restricts the space of phantom mechanisms with a bounded approximation ratio.

LEMMA 5.5. *A phantom mechanism which does not place at least one phantom on 0 and 1 has an unbounded approximation ratio for total cost and maximum cost.*

PROOF. Due to symmetry it suffices to consider phantom mechanisms which do not have any phantoms on 1. If all agents are located at 1, the optimal total and maximum cost are 0, but such a phantom mechanism must place the facility at a location strictly left of 1, as all $n + 1$ phantoms are strictly left of 1. Therefore the total cost and maximum cost approximation ratios are unbounded. \square

This lemma shows that when proving the approximation results, it suffices to consider phantom mechanisms where $p_1 = 0$ and $p_{n+1} = 1$. Such a phantom mechanism is equivalent to a phantom mechanism with only $n - 1$ phantoms, and is known to characterize the strategyproof and anonymous mechanisms which place the facility between the leftmost and rightmost agent locations [16].

5.1 Total Cost

We now prove approximation ratio results for the objective of total cost, a standard utilitarian measure of efficiency. We begin by giving a lower bound on the approximation ratio achievable by a phantom mechanism, with a proof which applies to both continuous and piecewise linear scaling functions.

THEOREM 5.6. *For continuous and piecewise linear scaling functions, any phantom mechanism has a total cost approximation ratio of at least r_q .*

PROOF. We consider two cases: either all phantoms are located at 0 or 1, or there is at least one phantom that is not located at 0 or 1. We start with the former case, where the facility is located at some agent. For every such phantom mechanism, we can always construct an instance where the scaling function is highest at the agent that receives the facility, so the corresponding total cost must be at least $\max_y q(y) \sum_i |f_{med}(x) - x_i|$, where f_{med} is the median mechanism. Now suppose that the instance has an even number of agents, with $x_{n/2} \neq x_{n/2+1}$, and that the scaling function is lowest at $(x_{n/2} + x_{n/2+1})/2$. Placing the facility at this point results in the optimal total cost of $\min_y q(y) \sum_i |f_{med}(x) - x_i|$, and thus the total cost approximation ratio must be at least r_q in this case.

Now consider the case where at least one phantom is located in $(0, 1)$. By Lemma 5.5, it suffices to consider phantom mechanisms

with at least one phantom at 0 and 1, so we are interested in where the remaining $n - 1$ phantoms are placed. We show that for every phantom mechanism f where at least one phantom p_j on $(0, 1)$ is not at the lowest point of the scaling function, there exists a phantom mechanism f' that places the $n - 1$ phantoms at the lowest point of the scaling function and has a strictly lower approximation ratio. Specifically, if the scaling function has multiple global minima, the phantoms are placed at the leftmost global minimum on $(0, 1)$. If the scaling function only has global minima on 0 and/or 1, then for sufficiently small $\epsilon > 0$, f' places the phantoms at ϵ or $1 - \epsilon$. For every agent location profile x such that $f(x) = p_j$, we claim there exists another profile x' where, under f' , the total cost ratio with the optimal total cost is strictly lower. We derive x' from x by first moving each agent closer to $f(x) = p_j$ by a sufficiently small constant factor $c > 0$ (so that each agent i 's distance changes from $|x_i - p_j|$ to $c|x_i - p_j|$), and then translating the profile such that the facility is co-located⁴ with the $n - 1$ phantoms of f' and the relative distances between the agents and the facility is the same but multiplied by the sufficiently small constant factor. The distances are still $c|x_i - p_j|$, so the ratio of total unscaled distances is still the same. But now the facility is placed at a point where the scaling function value is strictly lower, so the total cost ratio is lower. Therefore f' has the lowest approximation ratio of all phantom mechanisms which have at least one phantom on $(0, 1)$.

We now show that f' has an approximation ratio of at least r_q . We define a scaling function q which has a unique global minimum at some point $c_1 \in (0, 1)$, a unique global maximum at $c_2 > c_1$, and a point $c_3 > c_2$ where $q(c_3) = q(c_1) + \epsilon$, for sufficiently small $\epsilon > 0$. Consider a location profile where $x_1 = c_2 < x_2 = x_3 = \dots = x_n = c_3$. The mechanism f' places the facility at x_1 , resulting in a total cost of $(n - 1)(x_n - x_1) \max_y q(y)$, whilst the optimal total cost is at most $(x_n - x_1)(\min_y q(y) + \epsilon)$ by placing the facility at c_3 . Therefore f' has an approximation ratio of at least r_q . By exhaustion of cases, we see that any phantom mechanism has a total cost approximation ratio of at least r_q . \square

From this result, we can also easily see that among all of the phantom mechanisms, placing the facility at the median agent gives the best possible approximation ratio for total cost.

THEOREM 5.7. *For continuous and piecewise linear scaling functions, the median mechanism has a total cost approximation ratio of r_q .*

PROOF. The median mechanism minimizes the sum of distances, so its total cost is at most $\max_y q(y) \sum_i |f_{med}(x) - x_i|$, whilst the optimal total cost is at least $\min_y q(y) \sum_i |f_{med}(x) - x_i|$. This shows the approximation ratio is at most r_q , and we have a matching lower bound from Theorem 5.6. \square

Recall that phantom mechanisms characterize all strategyproof and anonymous mechanisms when agents have single-peaked preferences. We can consider scaling functions which ensure that agents have single-peaked preferences, and use Theorems 5.3 and 5.4 to find the specific total cost approximation ratio lower bounds for any strategyproof and anonymous mechanism.

⁴This occurs as long as there are agents on both sides of the phantoms. Since the phantoms are placed at ϵ and $1 - \epsilon$ instead of 0 and 1, there is space on both sides of the phantoms.

COROLLARY 5.8. *For continuous scaling functions, when agents are guaranteed to have single-peaked preferences, any strategyproof and anonymous mechanism has a total cost approximation ratio of at least e .*

COROLLARY 5.9. *For piecewise linear scaling functions with k line segments, when agents are guaranteed to have single-peaked preferences, any strategyproof and anonymous mechanism has a total cost approximation ratio of at least $(1 + \frac{1}{k})^k$.*

Again, the median mechanism would also give the best total cost approximation ratio, and can be suitably implemented for these restricted sets of scaling functions. However, in a scenario where the scaling function does not guarantee single-peaked agent preferences, phantom mechanisms no longer guarantee strategyproofness, and a constant mechanism has an unbounded approximation ratio. The *dictator mechanism*, which selects a specific agent to be the dictator and places the facility at their reported location, is known to be strategyproof for general agent preferences (though it is clearly not anonymous). We show that this mechanism has a bounded approximation ratio for total cost.

PROPOSITION 5.10. *The dictator mechanism has a total cost approximation ratio of $(n - 1)r_q$.*

PROOF. Suppose there are $n - 1$ agents at the same location away from the dictator. Also suppose that the scaling function is highest at the dictator's location, and lowest at the other $n - 1$ agents' location. The total cost approximation ratio is therefore at least $(n - 1)r_q$. To see that the approximation ratio cannot be higher than this, note that the total cost corresponding to the dictator mechanism is at most $d(n - 1) \max_{y \in X} q(y)$ where $d \in (0, 1]$, as the dictator receives 0 cost, and the optimal total cost is at least $d \min_{y \in X} q(y)$ when not all agents are at the same point. \square

5.2 Maximum Cost

We next prove approximation ratio results for the objective of maximum cost, a standard measure of egalitarian fairness. We begin by giving the maximum cost approximation ratio of any phantom mechanism with $p_1 = 0$ and $p_{n+1} = 1$, and again the proof applies to both continuous and piecewise linear scaling functions.

THEOREM 5.11. *For continuous and piecewise linear scaling functions, any phantom mechanism with $p_1 = 0$ and $p_{n+1} = 1$ has a maximum cost approximation ratio of $2r_q$.*

PROOF. We first prove that for every phantom mechanism f with $p_1 = 0$ and $p_{n+1} = 1$, there exists an agent profile where the facility is placed at either x_1 or x_n . Due to symmetry, it suffices to consider the former case, where we require that $p_2 < 1$ and $x_1 = x_2 = \dots = x_{n-1} > p_2$. Consider such an agent location profile where the facility is placed at x_1 and $x_1 \neq x_n$, and also consider a scaling function which has a global maximum at x_1 and a global minimum at $\frac{x_1 + x_n}{2}$. We have $MC(q, f(x), x) = (x_n - x_1) \max_{y \in X} q(y)$ and $MC(q, y_{MC}^*, x) = \frac{x_n - x_1}{2} \min_{y \in X} q(y)$, giving $\frac{MC(q, f(x), x)}{MC(q, y_{MC}^*, x)} = 2r_q$. Thus, for every phantom mechanism with at least one phantom on 0 and 1, we can construct an agent location profile and scaling function such that the cost ratio is $2r_q$, proving the lower bound.

To see that the approximation ratio is at most $2r_q$, note that a phantom mechanism with $p_1 = 0$ and $p_{n+1} = 1$ ensures the facility is placed within $[x_1, x_n]$, so the mechanism's maximum cost is at most $(x_n - x_1) \max_{y \in X} q(y)$, whilst the optimal maximum cost is at least $\frac{x_n - x_1}{2} \min_{y \in X} q(y)$. \square

As in the previous section, we can also derive the maximum cost approximation ratios when the scaling functions are restricted such that agents have single-peaked preferences.

COROLLARY 5.12. *For continuous scaling functions, when agents are guaranteed to have single-peaked preferences, any strategyproof and anonymous mechanism has a maximum cost approximation ratio of at least $2e$.*

COROLLARY 5.13. *For piecewise linear scaling functions with k line segments, when agents are guaranteed to have single-peaked preferences, any strategyproof and anonymous mechanism has a maximum cost approximation ratio of at least $2(1 + \frac{1}{k})^k$.*

Finally, we prove the maximum cost approximation ratio for the dictator mechanism, which ensures strategyproofness for all scaling functions.

PROPOSITION 5.14. *The dictator mechanism has a maximum cost approximation ratio of $2r_q$.*

PROOF. Consider an agent location profile where $x_1 = x_2 = \dots = x_{n-1}$ and $x_n \neq x_1$. When the scaling function has its minimum value at $\frac{x_1 + x_n}{2}$ and its maximum value on the facility/dictator, we have $\frac{MC(q, f(x), x)}{MC(q, y_{MC}^*, x)} = 2r_q$, showing that the approximation ratio is at least $2r_q$. For the matching upper bound, we note that the mechanism's maximum cost is at most $(x_n - x_1) \max_{y \in X} q(y)$ and the optimal maximum cost is at least $\frac{x_n - x_1}{2} \min_{y \in X} q(y)$. \square

6 DISCUSSION

In this paper, we have initiated the study of facility location problems where the agents' costs include a multiplicative scaling factor which depends on where the facility is placed. This problem applies to many real-world scenarios where the facility's effectiveness depends on external factors influencing its location. Our results focus on the objectives of total cost and maximum cost, and on both continuous and piecewise linear scaling functions.

Although we focus on single-peaked agent preferences and have some mechanism design results for general agent preferences, the agent preferences in our setting are a strict subspace of general preferences, taking the form where the metric distance is multiplied by some continuous scaling function. An open question is to characterize the deterministic mechanisms satisfying strategyproofness and surjectivity under these preferences, which we conjecture to be dictatorships.

There are many natural extensions to our paper. We can consider multiple facilities and multiple dimensions, as well as randomized mechanisms. We could also restrict further the space of piecewise linear scaling functions to single-dipped or single-peaked linear scaling functions, and find even lower approximation ratio bounds. Finally, our concept of scaling functions can be applied to existing variations of the facility location problem, such as capacitated facilities [1] and obnoxious facility location [7].

REFERENCES

- [1] H. Aziz, H. Chan, B. E. Lee, B. Li, and T. Walsh. 2020. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1806–1813.
- [2] H. Aziz, H. Chan, B. E. Lee, and D. C. Parkes. 2020. The capacity constrained facility location problem. *Games and Economic Behavior* 124 (2020), 478–490.
- [3] H. Aziz, A. Lam, B. E. Lee, and T. Walsh. 2022. Strategyproof and Proportionally Fair Facility Location. In *Proceedings of the 18th International Conference on Web and Internet Economics*.
- [4] S. Barberà and M. Jackson. 1994. A characterization of strategy-proof social choice functions for economies with pure public goods. *Social Choice and Welfare* 11, 3 (1994), 241–252.
- [5] K. C. Border and J. S. Jordan. 1983. Straightforward Elections, Unanimity and Phantom Voters. *The Review of Economic Studies* 50, 1 (1983), 153–170.
- [6] H. Chan, A. Filos-Ratsikas, B. Li, M. Li, and C. Wang. 2021. Mechanism Design for Facility Location Problems: A Survey. In *30th International Joint Conference on Artificial Intelligence (IJCAI 2021)*. 4356–4365.
- [7] Y. Cheng, W. Yu, and G. Zhang. 2013. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science* 497 (2013), 154–163.
- [8] M. Feldman, A. Fiat, and I. Golomb. 2016. On Voting and Facility Location. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 269–286.
- [9] R. Freeman, D. M. Pennock, D. Peters, and J. W. Vaughan. 2021. Truthful aggregation of budget proposals. *Journal of Economic Theory* 193 (2021), 105234.
- [10] A. Gibbard. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41, 4 (1973), 587–601.
- [11] S. Goel and W. Hann-Caruthers. 2023. Optimality of the coordinate-wise median mechanism for strategyproof facility location in two dimensions. *Social Choice and Welfare* 61, 1 (2023), 11–34.
- [12] A. B. Jennings, R. Laraki, C. Puppe, and E. M. Varloot. 2023. New characterizations of strategy-proofness under single-peakedness. *Mathematical Programming* (2023), 1–32.
- [13] M. Ma, M. Xiao, T. Bai, and B. Khoussainov. 2023. Facility location games with entrance fees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 5797–5804.
- [14] J. Massó and I. Moreno De Barreda. 2011. On strategy-proofness and symmetric single-peakedness. *Games and Economic Behavior* 72, 2 (2011), 467–484.
- [15] E. Miyagawa. 2001. Locating libraries on a street. *Social Choice and Welfare* 18, 3 (2001), 527–541.
- [16] H. Moulin. 1980. On strategy-proofness and single peakedness. *Public Choice* (1980).
- [17] K. Nehring and C. Puppe. 2006. The structure of strategy-proof social choice — Part I: General characterization and possibility results on median spaces. *Journal of Economic Theory* 135, 1 (2006), 269–305.
- [18] K. Nehring and C. Puppe. 2007. Efficient and strategy-proof voting rules: A characterization. *Games and Economic Behavior* 59, 1 (2007), 132–153.
- [19] H. Peters, H. van der Stel, and T. Storcken. 1992. Pareto optimality, anonymity, and strategy-proofness in location problems. *International Journal of Game Theory* 21 (1992), 221–235.
- [20] A. D. Procaccia and M. Tennenholtz. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation* 1, 4, Article 18 (2013).
- [21] M. A. Satterthwaite. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10 (1975), 187–217.
- [22] J. Schummer and R. V. Vohra. 2002. Strategy-proof Location on a Network. *Journal of Economic Theory* 104 (2002), 405–428.