Designing Optimal Mechanisms to Locate Facilities with Insufficient Capacity for Bayesian Agents

Gennaro Auricchio[1], **Jie Zhang**[2]

[1]University of Padova, Department of Mathematics
[2]University of Bath, Department of Computer Science
gennaro.auricchio@unipd.it, jz2558@bath.ac.uk

Abstract

In this paper, we study the Facility Location Problem with Scarce Resources (FLPSR) under the assumption that agents' type follows a probability distribution on $[0, 1]$. In the FLPSR, the goal is to identify the optimal locations for one or more capacitated facilities to maximize Social Welfare (SW), defined as the sum of the utilities of all agents. Since the total capacity of the facilities is insufficient to serve all agents, they compete in a First-Come-First-Served game to get accommodated. The main contribution of the paper ties Optimal Transport theory to the problem of selecting a truthful mechanism tailored to the agents' distributions. For the case of a single facility, we show that an optimal mechanism always exists. We examine three classes of probability distributions and characterize the optimal mechanism analytically or provide a routine to numerically compute it. We extend our results to the case in which we have two capacitated facilities to place. Initially, we assume that agents are independent and identically distributed, but our techniques generalize to scenarios where agents are not identically distributed. Finally, we validate our findings through several numerical experiments, including: (i) deriving optimal mechanisms for the class of beta distributions, (ii) assessing the Bayesian approximation ratio of these mechanisms for small numbers of agents, and (iii) assessing how quickly the expected mechanism SW converges to its limit.

Code — https://github.com/gauricchio93-boop/aaai26_FLP_with_Scarce_resources
Extended version — https://arxiv.org/abs/2412.00563

Introduction

The m-Capacitated Facility Location Problem (m-CFLP) extends the m-Facility Location Problem (m-FLP) by incorporating a constraint that limits the number of agents a facility can accommodate (Brimberg et al. 2001; Pal, Tardos, and Wexler 2001; Aardal et al. 2015). Both m-FLP and m-CFLP are fundamental subproblems in various applications within social choice theory, including disaster relief (Balcik and Beamon 2008), supply chain management (Melo, Nickel, and da Gama 2009), healthcare systems (Ahmadi-Javid, Seyedi, and Syam 2017), and public facility accessibility (Barda, Dupuis, and Lencioni 1990). At its core, the m-CFLP involves determining the locations of m facilities based on the positions of n agents, with each facility subject to a capacity constraint that limits the number of agents it can serve. While the algorithmic aspects of this problem have been extensively studied in the literature (Brimberg et al. 2001), the mechanism design aspects have only recently begun to attract attention of computer scientists (Aziz et al. 2020b). In mechanism design, facility location problems are examined under the assumption that each agent gets a utility that depends on the location of the facility. When agents self-report their positions, optimizing a communal objective is susceptible to manipulation driven by the agents' selfishness. Thus, a property that a mechanism must possess is *truthfulness*, which ensures that no agent can benefit by misrepresenting their private information. Truthfulness, however, forces the mechanism to select suboptimal locations, causing an efficiency loss described by the *approximation ratio*—the worst-case ratio between the objective achieved by the mechanism and the optimal objective attainable (Nisan and Ronen 1999).

In this paper, we study the Facility Location Problem with Scarce Resources (FLPSR), a variant of the m-CFLP where the facilities cannot accommodate all the agents (Aziz et al. 2020b). This framework differs from the one proposed in (Aziz et al. 2020a) for two reasons: (i) the total capacity of the facilities is lower than the total number of agents, leaving some agents unaccommodated, and (ii) the mechanism designer does not enforce an agent-to-facility assignment. Consequently, after the facility positions are determined, agents compete in a First-Come-First-Served (FCFS) game to access the facilities. While the classic worst-case analysis provides valuable insights for scenarios where the mechanism designer lacks prior information on the agents' positions, it does not offer flexibility for tuning the mechanism. In fact, if the agents belong to a population whose density is known, the best mechanism according to worst-case analysis may not be the optimal choice, as the next example shows.

Example 1 (Motivating Example.) *Let us consider the case in which n agents living on a street, modeled as the segment $[0, 1]$, participate in an eliciting routine to place a single facility able to accommodate 20% of the population. The agents' density on the road μ is known and is represented by the density function $f_\mu(x) = 2(1-x)$, with $x \in [0, 1]$. Once the facility position y is determined, the set of agents accom-*

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

modated by y is the 20% of the population that is closer to y, i.e. the set of agents accommodated by y is $B_R(y)$ where the radius $R > 0$ is such that $\mu(B_R(y)) = 0.2$. Finally, an agent located at x that gets accommodated by y gets an utility equal to $1 - |x - y|$. From (Aziz et al. 2020b), it is known that the truthful mechanism achieving the lowest approximation ratio is the median mechanism, namely Med, which places the facility at the position of the median agent. We give a graphical representation of the problem in Figure 1. For n large enough, the median mechanism locates the facility at the median of μ (De Haan and Taconis-Haantjes 1979), that is $y = 1 - \frac{1}{\sqrt{2}} \approx 0.29$. The median mechanism thus induces an expected Social Welfare (SW)–defined as the expected total utility accrued by the agents–equal to

$$\mathbb{E}_\mu[SW_{Med}(\vec{X})] \approx \int_{0.29-R}^{0.29+R} (1 - |x - 0.29|) f_\mu(x) dx \approx 0.19,$$

where $R \sim 0.07$ is the radius of the ball centered in 0.29 that encompasses 20% of the population. Let us now consider the decile mechanism (DM), which places the facility at the location of the $\lfloor \frac{n}{10} \rfloor$ leftmost agent. For n large enough, the decile mechanism places the facility at $y' \sim 0.05$. Therefore the expected SW of the decile mechanism is

$$\mathbb{E}_\mu[SW_{DM}(\vec{X})] \approx \int_{0.05-R}^{0.05+R} (1 - |x - 0.05|) f_\mu(x) dx \approx 0.20,$$

where $R \sim 0.06$ is the radius of the ball centred in 0.05 that encompasses 20% of the population. Thus, contrary to the classic analysis (Aziz et al. 2020b), DM induces a higher expected SW than the median mechanism.

Motivated by this example, we study the problem of identifying the optimal mechanism for the Facility Location Problem with Scarce Resources (FLPSR) when the agents' positions follow a probability distribution. Under these conditions, we connect the FLPSR to Optimal Transport theory and show that the expected SW converges when the number of agents goes to infinity and characterize its limit.

Our Contribution. First, we consider the one facility case. Given a probability measure μ and the facility capacity q, we introduce the radius function $R_{\mu,q}$ which, given y returns the unique value of r that satisfies $\mu([y-r, y+r]) = q$. The radius function identifies the set of agents accommodated by the facility located at y when the number of agents goes to infinity. Through $R_{\mu,q}$, we connect the FLPSR to Optimal Transport, allowing us to (i) characterize the asymptotic expected SW induced by a facility position y and (ii) show that for every probability distribution there exists a mechanism whose asymptotic expected SW is equal to the optimal expected SW attainable. We then extend our results to the case where agents are not identically distributed.

Second, we extend our study to the two facilities case. We are interested in Equilibrium Stable (ES) percentile mechanisms, *i.e.* mechanisms whose outcome induce a unique Social Welfare (Auricchio, Clough, and Zhang 2024). In this case, defining an ES mechanism whose asymptotic SW coincides with the optimal SW attainable is often impossible. For example, no optimal ES mechanism can be defined when the sum of the capacities is larger than $\frac{2}{3}$. We then characterize a set of necessary and sufficient conditions under which there exists an optimal ES mechanism and use it to retrieve the optimal mechanism for specific cases. When an optimal mechanism does not exist, we consider the best mechanisms, *i.e.* the ES mechanisms that attain the maximum SW. In particular, we characterize the best mechanism when μ and $\vec{q} = (q_1, q_2)$ satisfy some mild assumptions. We conclude by proposing a search routine that finds the best ES mechanism for the cases not covered by previous results.

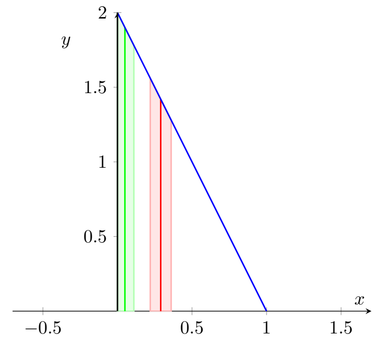

Figure 1: The agents' distribution f_μ along with the facility identified by the Median Mechanism (in red) and the facility identified by the Decile Mechanism (in green). The red and green area describes the set of agents accommodated by the red and green facility, respectively.

Lastly, we run several numerical experiments to validate our findings. In particular, we empirically assess the Bayesian approximation ratio of the percentile mechanisms found via our routines for small number of agents. Our results confirm that our mechanisms achieve a small Bayesian approximation ratio (< 1.02) even when the number of agents is less than 100. We also evaluate the speed at which the expected SW attained by the mechanism converges to its theoretical limit. All our experimental results confirm our theoretical results and proves that expected SW of the mechanisms obtained by our routine is optimal or quasi-optimal.

Due to space limits, proofs are deferred to the Appendix.

Related Work. The mechanism design aspects of the m-FLP were firstly considered in (Procaccia and Tennenholtz 2013), where the authors studied the problem of eliciting the position of a facility from the reports of n self-interested agents. Following this seminal work, various mechanisms with constant approximation ratios for placing one or two facilities on lines (Filos-Ratsikas et al. 2017), trees (Feldman and Wilf 2013), and generic metric spaces (Meir 2019) were introduced. In this work, we consider the Facility Location Problem with Scarce Resources introduced in (Aziz et al. 2020b), where the authors studied the problem of locating a facility with insufficient capacity in $[0, 1]$. This framework was later extended to the case in which there is more than one facility (Auricchio, Clough, and Zhang 2024).

Bayesian mechanism design (Hartline and Lucier 2010; Chawla and Sivan 2014) is an alternative framework for mechanism design where agents' private information follow a probability distribution. This framework has been explored for routing games (Gairing, Monien, and Tiemann 2005), combinatorial mechanisms (Lucier and Borodin 2010), and auction mechanism design (Hartline and Roughgarden 2009). To our knowledge, only two other papers have studied a variant of the FLP within a Bayesian mechanism design framework: (Auricchio and Zhang 2026, 2024) which studies the classic m-FLP problem where facilities do not have a capacity limit and (Auricchio, Zhang, and Zhang 2024) which examines the m-CFLP under the assumption that the mechanism designer can force agents to use a specific facility. Notice that previous works on Bayesian mechanism design for FLP variants assume that the total facility capacity exceeds the number of agents. In our case in study this assumption does not hold; making prior techniques inapplicable as they overlook the FCFS game dynamics and fail to ensure the Equilibrium Stability of the mechanisms.

Lately, Optimal Transport-based methods have found their application within the broad landscape of Theoretical Computer Science. Notable examples include Computer Vision (Rubner, Tomasi, and Guibas 1998; Pele and Werman 2009), Computational Statistics (Levina and Bickel 2001), and Machine Learning in general (Scagliotti 2023; Frogner et al. 2015; Cuturi and Doucet 2014). However, there have been limited advancements in applying OT theory to mechanism design. To the best of our knowledge, only auction design (Daskalakis, Deckelbaum, and Tzamos 2013), the m-FLP (Auricchio and Zhang 2026), and the m-CFLP (Auricchio, Zhang, and Zhang 2024) leverage OT theory to design mechanisms tuned to the distribution of agents' preferences.

Preliminaries

First we recall the basic notions on FLP with Scarce Resources, Bayesian Mechanisms Design and OT.

The FLP with Scarce Resources. Let $\vec{x} \in [0,1]^n$ be the position of n agents in $[0,1]$ and let $\vec{q} = (q_1, \dots, q_m)$ be the m-dimensional vector containing the percentage capacities of the m facilities, so that the j-th facility accommodate up to $\lfloor q_j n \rfloor$ agents. The total capacity of the facilities is less than the number of agents, hence $\sum_{j\in[m]} q_j < 1$. We denote with $\vec{y} = (y_1, \dots, y_m)$ the positions of the facilities, so that the j-th facility has position y_j and capacity q_j.

Once $\vec{y}$ is fixed, agents compete in a First-Come-First-Served (FCFS) game to access the facilities. In a FCFS game each agent selects one of the facilities, so that the set of strategies of each agent is $[m] := \{1, 2, \dots, m\}$. Then, the $\lfloor q_j n \rfloor$ agents that are closer to y_j and that have selected strategy j get accommodated by y_j and receive a utility equal to $1 - |x_i - y_j|$, otherwise the agent gets zero utility. Given a facility position $\vec{y}$, the FCFS game has at least one pure Nash Equilibrium (Auricchio, Clough, and Zhang 2024). The SW of the facility location $\vec{y}$ according to a NE, namely γ, is the sum of all the agents' utilities when agents play according to γ, *i.e.* $SW_\gamma(\vec{x}, \vec{y}) = \sum_{i\in[n]} u_i(\vec{x}, \vec{y}; \gamma)$. When $m = 1$, the NE of the FCFS game is unique, however, when $m \geq 2$, the NE is not unique and the SW depends on the NE (Auricchio, Clough, and Zhang 2024). The Facility Location Problem with Scarce Resources (FLPSR) consists in finding the locations $\vec{y}$ that maximizes the SW across all the NEs.

Bayesian Mechanism Design. Given m capacities $\vec{q}$, a mechanism for the FLPSR is a function $M : [0,1]^n \to \mathbb{R}^m$ that maps a vector containing the agents' reports to a facility location $\vec{y}$. In what follows, we consider the family of percentile mechanisms (Sui, Boutilier, and Sandholm 2013). Given a percentile vector $\vec{p} \in [0,1]^m$, the percentile mechanism induced by $\vec{p}$, namely $\mathcal{PM}_{\vec{p}}$, determines the positions of the facilities by placing the j-th facility at the position of the $\lfloor p_j(n-1) \rfloor + 1$ leftmost agent. Every percentile mechanism is uniquely determined by the percentile vector $\vec{p}$.

A mechanism M is said to be truthful if no agents can misreport their true position to increase their utility, regardless of the others agents' reports and the other agents' strategies adopted in the FCFS game. More formally,

$$\max_{s_i\in[m]} u_i(\vec{x}, M(\vec{x}); s_i, \vec{s}_{-i}) \geq \max_{s_i'\in[m]} u_i(\vec{x}, M(\vec{x}'); s_i', \vec{s}_{-i}),$$

where (i) $\vec{x}' = (x_i', \vec{x}_{-i})$ for every $x_i' \in [0,1]$, (ii) s_i represents the i-th agent strategy, and (iii) $\vec{x}_{-i}$ and $\vec{s}_{-i}$, are the vectors containing the positions and strategies of the other $n-1$ agents. A mechanism is Equilibrium Stable (ES) if, for every possible input $\vec{x}$, all the NEs induced by the mechanism's output attain the same SW. If $m = 1$ each truthful mechanism is ES, however, when $m \geq 2$ a percentile mechanism induced by $\vec{v}$ is ES if and only if

$$\lfloor v_{j+1}(n-1) \rfloor - \lfloor v_j(n-1) \rfloor \geq \lfloor (q_{j+1} + q_j)(n-1) \rfloor, \quad (1)$$

for every $j = 1, \dots, m-1$. Following the Bayesian mechanism design framework, we assume that agents' type is represented by a random variable X_i with an associated probability distribution μ_i. Given a percentile mechanism $\mathcal{PM}_{\vec{p}}$, the Bayesian approximation ratio of $\mathcal{PM}_{\vec{p}}$ is the ratio between the expected SW of $\mathcal{PM}_{\vec{p}}$ and the expected optimal SW, that is $B_{ar}^{(n)}(f) := \frac{\mathbb{E}[SW_{opt}(\vec{X}_n)]}{\mathbb{E}[SW_{\vec{p}}(\vec{X}_n)]}$, where $SW_{opt}(\vec{X}_n)$ is the maximum SW attainable across all the NEs and $\mathbb{E}$ denotes the expected value over the joint distribution of $\vec{X}_n$.

Optimal Transport. Let $\mathcal{P}(\mathbb{R})$ be the set of probability measures over $\mathbb{R}$. Given $\mu \in \mathcal{P}(\mathbb{R})$ and a positive measure ν over $\mathbb{R}$, such that $\nu(\mathbb{R}) \leq 1$, the Wasserstein distance between μ and ν is defined as

$$W_1(\mu, \nu) = \min_{\pi\in\Pi(\mu,\nu)} \int_{\mathbb{R}\times\mathbb{R}} |x - y| d\pi, \quad (2)$$

where $\Pi(\mu, \nu)$ is the set of transportation plans between μ and ν, *i.e.* the probability measures over $\mathbb{R} \times \mathbb{R}$ whose first marginal is μ and the second marginal stochastically dominates ν (Pele and Werman 2009). For a complete introduction to the Wasserstein distances, we refer to (Villani 2009).

Basic Assumptions. Finally, we outline the set of assumptions that we tacitly assume to be true throughout the rest of the paper. First, since agents are located over $[0,1]$, we assume that the agents' distribution μ is supported over $[0,1]$.

Second, we assume that μ is absolutely continuous, *i.e.* it is induced by a probability density f_μ. Without loss of generality, we assume f_μ to be continuous and positive up to a finite number of points, as any probability measure can be approximated by such distribution (Villani 2009).

Third, we restrict our study only percentile mechanisms since, to the best of our knowldge, these are the only mechanisms know to be ES (Auricchio, Clough, and Zhang 2024).

The One Facility Case

Given a n dimensional vector containing the positions of n agents $\vec{x} \in [0,1]^n$, let y be the position at which the facility is located. We denote with μ_i the probability distribution associated to the i-th agent and denote with q the percentage capacity of the facility, so that the facility opened at y is able to accommodate $\lfloor qn \rfloor$ agents. Since we have only one facility located at y, the set of agents accommodated according to the FCFS game is, up to ties, the set of the $k = \lfloor qn \rfloor$ closest agents to y, for every $\vec{x}$. For the sake of simplicity, we first consider the case in which each agent is independent and identically distributed, we then extend our finding to cases in which agents are not identically distributed.

Characterizing the Optimal Mechanism

In this section, we connect the FLPSR with Optimal transport theory and show that there exists a percentile mechanism whose expected SW is optimal when the number of agents increases, regardless of the agents' distribution μ and the capacity q. Let $R_{\mu,q} : [0,1] \to [0,1]$ be the radius function, implicitly defined by the following equation

$$\mu([y - R_{\mu,q}(y), y + R_{\mu,q}(y)]) = q. \tag{3}$$

Lemma 1 *Given μ and $q \in [0,1]$, for every $y \in [0,1]$ the identity $\mu([y-h, y+h]) = q$ is satisfied by a unique value of h. In particular, for any y, the value $R_{\mu,q}(y)$ that satisfies (3) is unique. Moreover, $R_{\mu,q}$ is continuous over $[0,1]$, is differentiable over $(0,1)$, and the following identity holds*

$$R'_{\mu,q}(y) = \begin{cases} 1 & \text{if } y \in [F_\mu^{[-1]}(1-\frac{q}{2}), 1] \\ -1 & \text{if } y \in [0, F_\mu^{[-1]}(\frac{q}{2})] \\ \frac{f_\mu(y-R_{\mu,q}(y)) - f_\mu(y+R_{\mu,q}(y))}{f_\mu(y-R_{\mu,q}(y)) + f_\mu(y+R_{\mu,q}(y))} & \text{otherwise} \end{cases}$$

Having established the properties of the radius function $R_{\mu,q}$, we connect the FLPSR to Optimal Transport.

Lemma 2 *Given $\vec{x} \in \mathbb{R}^n$ and $y \in \mathbb{R}$, we have that*

$$SW(\vec{x}, y) = q - W_1(\mu_{\vec{x}}, q\delta_y), \tag{4}$$

where δ_y is the Dirac's delta centred in y and $\mu_{\vec{x}} = \frac{1}{n}\delta_{x_i}$.

Lemma 2 points out that the set of agents accommodated by a facility in y according to the NE of the FCFS game characterizes the optimal transportation plan between $\mu_{\vec{x}}$ and $q\delta_y$. Owing to (4), we compute the limit for $n \to \infty$ of the expected SW for any $y \in [0,1]$ and any $\mu \in \mathcal{P}([0,1])$.

Theorem 1 *Let $X \sim \mu$ be the random variable describing the agents' type, then we have*

$$\lim_{n\to\infty} \mathbb{E}[SW(\vec{X}; y)] = q - W_1(\mu, q\delta_y), \tag{5}$$

where $y \in \mathbb{R}$ is the facility position and $\mathbb{E}$ is the expected value with respect to μ.

Since q is a constant, the position y that maximizes the SW of the problem is the solution to the following problem

$$\min_{y\in[0,1]} \mathcal{W}(y) := \min_{y\in[0,1]} \int_{y-R_{\mu,q}(y)}^{y+R_{\mu,q}(y)} |x-y| d\mu. \tag{6}$$

Owing to the properties of μ, $\mathcal{W}$ is continuous and thus problem (6) admits a solution. Moreover, any solution to the problem identifies an optimal percentile mechanism.

Theorem 2 *Problem (6) admits a solution. Given a solution $\bar{y}$, $\vec{p} = (F_\mu(\bar{y}))$ induces an optimal percentile mechanism.*

As an example, we implement Theorem 2 to retrieve the optimal percentile mechanism for the uniform distribution.

Example 2 *Let $q \in [0,1]$ be the capacity of a facility and let μ be the uniform distribution over $[0,1]$. First, given $y \in [0,1]$, the value $R_{\mu,q}(y)$ is determined by the equation*

$$\min\{y + R_{\mu,q}(y), 1\} - \max\{y - R_{\mu,q}(y), 0\} = q.$$

It is then easy to see that

$$R_{\mu,q}(y) = \begin{cases} q - y & \text{if } y < \frac{q}{2} \\ \frac{q}{2} & \text{if } \frac{q}{2} \le y \le 1 - \frac{q}{2} \\ q - (1-y) & \text{otherwise.} \end{cases} \tag{7}$$

Plugging (7) into the right-hand side of (6), we obtain

$$\mathcal{W}(y) = \begin{cases} \frac{1}{2}y^2 + \frac{1}{2}(q-y)^2 & \text{if } y < \frac{q}{2} \\ \frac{q^2}{4} & \text{if } \frac{q}{2} \le y \le 1 - \frac{q}{2} \\ \frac{1}{2}(q-y)^2 + \frac{1}{2}(1-y)^2 & \text{otherwise.} \end{cases}$$

In particular $R_{\mu,q}$ is minimized whenever $y \in [\frac{q}{2}, 1 - \frac{q}{2}]$, hence any $\vec{p} = (F_\mu(y))$ such that $y \in [\frac{q}{2}, 1 - \frac{q}{2}]$ is optimal.

We then compute the derivative of $\mathcal{W}$, this result enables us to use root-finding methods to compute the optimal percentile mechanism given $\mu \in \mathcal{P}([0,1])$ and $q \in [0,1]$.

Theorem 3 *Given μ and q, the following identity holds $\mathcal{W}'(y) = 2R_{\mu,q}(y)R'_{\mu,q}(y) - \Delta_\mu(y)$, where $\Delta_\mu(y) = F_\mu(y + R_{\mu,q}(y)) + F_\mu(y - R_{\mu,q}(y)) - 2F_\mu(y)$.*

Computing the Optimal Mechanism

We now apply Theorems 2 and 3 to compute the optimal percentile mechanisms for three classes of probability distributions: monotone distributions, Single-Peaked (SP) distributions, and Single-Dipped (SD) distributions.

First, we restrict the set in which the optimal position y is.

Lemma 3 *Let q be a capacity and μ be a probability distribution. Then any solution to problem (6), namely $\bar{y}$, is such that $\bar{y} \in [\frac{q}{2}, 1 - \frac{q}{2}]$.*

Notice that if $q = 1$, Lemma 3 implies that the only optimal position is the median of μ, which is consistent with previous results (Aziz et al. 2020b).

Monotone distributions. First, we consider all the probability measures whose density is either non-increasing or non-decreasing. This class of probability distributions includes Triangular probability distributions as well as the exponential and the Chi-square distributions restricted to $[0, 1]$. Without loss of generality, we consider the case in which the density is non-increasing. For this class of problems, the optimal percentile mechanism depends only on the value of q.

Theorem 4 *Let f_μ be monotone non-increasing. Given $q \in [0, 1]$, the optimal percentile vector is $\vec{p} = (\frac{q}{2})$.*

Single-Peaked (SP) Distributions. We now consider the class of SP distributions, *i.e.* distributions whose density has a unique maximum, including the Beta (with positive parameters), Bates, and Truncated Gaussian distributions. If μ is an SP and symmetric distribution, the best percentile mechanism is the median mechanism regardless of q.

Theorem 5 *The best percentile mechanism for a symmetric Single-Peaked distribution is the median mechanism.*

Indeed, if μ is symmetric and SP, the median of μ is the unique point at which $R_{\mu,q}$ is minimized and $\Delta_\mu = 0$. For asymmetric SP distributions, this argument does not hold, thus to retrieve the optimal percentile mechanism we search for the zeros of $\mathcal{W}'$ (Pasupathy and Kim 2011).

Single-Dipped (SD) measures. Lastly, we consider the class of SD distributions, *i.e.* measures whose density has a unique minimum. This class of probability measures includes quadratic distributions, arcsine distributions, and Beta distributions with negative parameters. When the SD measure is symmetric the optimal percentile vector depends only on the capacity q. For asymmetric SD distributions, there are only two possible optimal percentile mechanisms.

Theorem 6 *Let $\mu \in \mathcal{P}(\mathbb{R})$ be SD and q be a capacity. Denoted with $\vec{p}$ the optimal percentile vector, we have $\vec{p} = (\frac{q}{2})$ if μ is symmetric. Otherwise $\vec{p} = (\frac{q}{2})$ or $\vec{p} = (1 - \frac{q}{2})$.*

Dropping the I.D. Assumption

To conclude, we extend our results to the case in which agents are not identically distributed. Denoted with μ_i the probability distribution of the i-th agent, we have

$$\mathbb{E}[SW(\vec{X}; y)] = \sum_{i=1}^{n} \frac{1}{n} \int_0^1 |x - y| d\mu_i = \int_0^1 |x - y| d\tilde{\mu},$$

where $\tilde{\mu} = \frac{1}{n} \sum_{i=1}^{n} \mu_i$ and $y \in [0, 1]$. Let us now assume that each μ_i can be expressed as a conditional law $\mu(\circ|\theta)$ where $\theta \in \Theta$ is a parameter that describes the different agents' distributions. Under this assumption, we have $\mu_i = \mu(\circ|\theta_i)$, hence $\frac{1}{n} \sum_{i=1}^{n} \mu_i \to \mu(\circ|\theta)\eta(\theta)$ as $n \to \infty$, where $\eta \in \mathcal{P}(\Theta)$ is the probability distribution that describes how likely is that an agent distribution is $\mu(\circ|\theta)$.

Theorem 7 *Let (X, θ) be a random vector whose associated law is $\mu(x, \theta) = \mu(x|\theta)\eta(\theta)$. Given $y \in [0, 1]$, denoted by $\bar{\mu} = \int_\Theta \mu(\circ|\theta) d\eta$ and $\mathbb{E}$ the expected value with respect to $\bar{\mu}$, we have*

$$\lim_{n \to \infty} \mathbb{E}[SW(\vec{X}; y)] = q(1 - W_1(\bar{\mu}, q\delta_y)).$$

Owing to Theorem 7, we can compute the optimal mechanism by setting $\mu := \bar{\mu} = \int_\Theta \mu(\circ|\theta) d\eta$.

The Two Facilities Case

We now consider the case in which we have to locate two facilities whose percentage capacities are q_1 and q_2, respectively. Without loss of generality, we assume that $q_1 \geq q_2$. Moreover, we assume that agents are i.i.d., as the case in which agents are not i.i.d. can be obtained via Theorem 7.

When $m = 2$, the agents accommodated by a facility located at y is no longer the set of $\lfloor qn \rfloor$ closest agents to y (see the Appendix). For this reason, we define a new radius function $\vec{R}_{\mu,\vec{q}} : [0, 1]^2 \to [0, 1]^2$ as it follows: when $y_1 \leq y_2$, we set $\vec{R}_{\mu,\vec{q}}(\vec{y}) = (R_1(y_1, y_2), R_2(y_1, y_2))$, where $R_i := R_i(y_1, y_2)$ are the unique values satisfying

$$\mu\bigg(\Big[y_1 - R_1, \min\Big\{\max\Big\{\frac{z}{2}, y_2 - R_1\Big\}, y_1 + R_2\Big\}\Big]\bigg) = q_1$$

$$\mu\bigg(\Big[\max\Big\{\min\Big\{y_2 - R_2, \frac{z}{2}\Big\}y_2 - R_2\Big\}, y_2 + R_2\Big]\bigg) = q_2.$$

where $z = y_1 + y_2$. When $y_2 < y_1$ the definition is obtained by swapping y_1 with y_2 and q_1 with q_2. Given the capacities $\{q_i\}_{i=1,2}$, we search for $\vec{y} = (y_1, y_2)$ that minimizes

$$\mathcal{W}(y_1, y_2) = \int_{y_1 - R_1}^{z_1} |x - y_i| d\mu + \int_{z_2}^{y_2 - R_2} |x - y_i| d\mu,$$

where $z_1 = \min\Big\{\max\Big\{\frac{y_1+y_2}{2}, y_2 - R_1\Big\}, y_1 + R_2\Big\}$, $z_2 = \max\Big\{\min\Big\{y_2 - R_2, \frac{y_1+y_2}{2}\Big\}y_2 - R_2\Big\}$, and $\{q_i\}_{i=1,2}$. Unfortunately, it is not possible to extend Theorem 2 to the case in which $m = 2$. Indeed, the percentile mechanism induced by the couple (y_1, y_2) that minimizes $\mathcal{W}$ is, in general, not ES. Given the set containing all the minimizers of $\mathcal{W}$, we then outline a condition to check whether there exists an ES percentile mechanism whose limit expected SW is optimal.

Theorem 8 *Given μ and $\vec{q}$, let $\mathcal{Y}$ be the set of minimizers of $\mathcal{W}$. Then, $\mathcal{Y}$ is non empty and there exists an optimal ES percentile mechanism if and only if exists $\vec{y} \in \mathcal{Y}$ such that*

$$F_\mu(y_2) - F_\mu(y_1) \geq q_1 + q_2. \tag{8}$$

From Theorem 8 we infer two sets of conditions under which no ES mechanism is optimal.

Corollary 1 *Given μ and $\vec{q}$, if $q_1 + q_2 \geq \frac{2}{3}$ then no ES percentile mechanism is optimal. Moreover, if μ is monotone or Single Peaked, no percentile mechanism is ES and optimal.*

Lastly, we show that only SD distributions admit an optimal ES percentile mechanism, when $q_1 + q_2 \leq \frac{2}{3}$.

Corollary 2 *Let $\vec{q}$ be a capacity vector such that $q_1 + q_2 \leq \frac{2}{3}$. If μ is a symmetric SD distributions, the percentile mechanism induced by $\vec{v} = (\frac{q_1}{2}, 1 - \frac{q_2}{2})$ is ES and optimal.*

Searching the Best Percentile Mechanism

In this section, we retrieve the percentile mechanism that attains the lowest Bayesian approximation ratio by minimizing $\mathcal{W}$ under the additional ES constraint (8).

Theorem 9 *Let y_1 and y_2 be such that $F_\mu(y_2) - F_\mu(y_1) \geq q_1 + q_2$, then $\vec{R}_{\mu,\vec{q}}(\vec{y}) = (R_{\mu,q_1}(y_1), R_{\mu,q_2}(y_2))$. Therefore, if $\mathcal{PM}_{\vec{p}}$ is an ES percentile mechanism, then*

$$\lim_{n\to\infty} \mathbb{E}[SW_{\vec{p}}(\vec{X})] = Q - \sum_{i=1}^{m} \int_{B_{R_{\mu,q}(y_i)}} |x - y_i| d\mu, \quad (9)$$

where $y_i = F_\mu^{[-1]}(p_i)$ and $R_{\mu,q}$ is the radius function in (3).

We now propose a search algorithm to find the best ES percentile mechanism given μ, $\vec{q}$, and a tolerance $\delta > 0$. Our search method hinges upon the following result, that reduces the space in which the minimizers of $\mathcal{W}$ are located.

Lemma 4 *Let μ be a probability measure and let $\vec{q} = (q_1, q_2)$ be a capacity vector. Let $Q = q_1 + q_2$ be the total capacity of the facilities and let $(\bar{y}_1, \bar{y}_2) \in \mathcal{Y}$ be a minimizer of $\mathcal{W}$ (without loss of generality, we assume that $\bar{y}_1 \leq \bar{y}_2$, as the other case is symmetric).*

Then, we have that $\bar{y}_1 \in [0, F_\mu^{[-1]}(1-Q)]$. Moreover, denoted with $M = F_\mu^{[-1]}(F_\mu(\bar{y}_1) + Q)$, $\bar{y}_2$ is such that (i) $\bar{y}_2 \in [M, F_\mu(1-\frac{q_2}{2})]$ if $M \leq F_\mu(1-\frac{q_2}{2})$; or (ii) $\bar{y}_2 = M$.

In Algorithm 1, we present a routine that, given μ, $\vec{q}$ and a tolerance parameter $\delta > 0$, finds the best ES percentile mechanism by discretizing the set of feasible solutions into intervals of length δ. It then searches for the best solution among the feasible one determined by the previous discretization. Since this procedure generates at most $\lceil \delta^{-1} \rceil$ intervals, the time-complexity of Algorithm 1 is $O(\delta^{-2})$ as it needs to compare every possible couple. The asymptotic SW induced by the percentile mechanism associated with the vector $\vec{y}_\delta$ returned by Algorithm 1 is at most δ less than the SW induced by the best percentile mechanism.

Theorem 10 *Let μ be a probability measure and let $\vec{q}$ be the capacity vector. Given $\delta > 0$, let $\vec{y}_\delta$ be the vector containing the positions returned by Algorithm 1. Then, $\left|\mathcal{W}(\vec{y}_\delta) - \min_{\vec{y}\ s.t.\ |F_\mu(y_2)-F_\mu(y_1)|\geq q_1+q_2} \mathcal{W}(\vec{y})\right| \leq \delta$.*

Numerical Experiments

In this section, we run numerical experiments to validate our theoretical findings. The aim of our tests is threefold: (i) we assess the feasibility of Algorithm 1 in computing the optimal percentile mechanism for locating one or two facilities tailored to a given distribution. (ii) We evaluate the performance of the computed mechanisms when we have a small number of agents. (iii) We measure the speed at which the expected SW attained by the mechanism converges. Due to space limits, part of the results are deferred to the Appendix.

Throughout our experiments, we consider both the cases in which we have one or two facilities to locate, hence $m = 1, 2$. When $m = 1$, we consider facilities whose capacity q ranges in $\{0.2, \ldots, 0.9\}$. When $m = 2$, we consider $\vec{q} = (q_1, q_2)$ with $q_i \in \{0.2, 0.3, 0.4\}$ for a total of 6 different capacity vectors, up to symmetries. When the agents' positions are i.i.d., we consider Beta distributions $\mathcal{B}(\alpha, \beta)$ with $\alpha, \beta \in \{2, \ldots, 6\}$. We focus on this class of distributions as it can be adjusted to fit various types of data, allowing it to model symmetric and asymmetric probability distributions.

Algorithm 1: Search Routine to minimize $\mathcal{W}$

```
Initialize m ← 100, val ← 0, y_1, y_2 ← 0
Set T = {t_0 = 0, ..., t_{N_1} = F_μ^{[-1]}(1 − Q)} s.t. t_i < t_{i+1} and |t_i − t_{i+1}| ≤ δ
for t ∈ T do
    M = F_μ^{[-1]}(F_μ(t) + Q)
    if M < F_μ(1 − q_2/2) then
        Set S = {s_0 = M, ..., s_i, ..., s_{N_2} = F_μ(1 − q_2/2)}, where s_i < s_{i+1} and |s_i − s_{i+1}| ≤ δ
    else
        S = {M}
    end if
    for s ∈ S do
        val = W(t, s)
        if val < m then
            m ← Int, y_1 ← t, y_2 ← s
        end if
    end for
end for
Return m, y_1, y_2
```

Furthermore, notice that the best percentile mechanism associated with most of the other standard probability distributions supported over $[0, 1]$ (e.g. the uniform distribution or the triangular distribution) can be obtained from either Theorem 4, 5, or 6, thus we do not consider them in our experiments. When agents are not i.d., we consider a family of uniform distributions $\left\{\frac{1}{\theta}\mathbb{I}_{[0,\theta]}\right\}_{\theta\in[0,1]}$, where $\mathbb{I}_A(x)$ is the indicator function of A, which is equal to 1 if $x \in A$ and equal to 0 otherwise. The agent type θ has density $3\theta^2$, thus $\Theta = [0, 1]$ and η is the probability measure induced by $f_\eta(\theta) = 3\theta^2$.

Computing the Best Percentile Mechanism. First, we compute the optimal percentile mechanism associated with different Beta distributions when we need to locate one or two facilities. When $m = 1$, we use a simple bisection method to find the zeros of $\mathcal{W}'$. When $m = 2$, we run Algorithm 1 with $\delta = 0.001$. We ran our code with various choices of δ. Values of $\delta > 0.001$ did not yield significant improvements in solution quality, while setting $\delta \ll 0.001$ led to a sharp decline in quality without notable gains in computation time. In Table 1, we report our findings for $\mathcal{B}(\alpha, \beta)$ with $\alpha, \beta \in \{2, 3, 4, 5, 6\}$ and (i) $q = 0.5$ when $m = 1$, and (ii) $\vec{q} = (0.2, 0.2)$ when $m = 2$. Our findings are in line with Theorem 5: when $m = 1$ and $\alpha = \beta$ the optimal percentile mechanism is the median mechanism.

The Bayesian Approximation Ratio. We now assess the quality of the mechanisms outlined by our theoretical results or by Algorithm 1. We first consider agents whose positions are $n = 20, \ldots, 100$ i.i.d. samples of a Beta distribution. We then compute the expected optimal SW and the expected SW induced by the optimal mechanism as the mean of 10000 instances. In Figure 2, we plot the Bayesian approximation ratio for $\mathcal{B}(6, 2)$ when $m = 1, 2$. Since the Confidence Interval of all our results is less than 0.005, we omit it from our plots. For $\mathcal{B}(6, 2)$, the Bayesian approximation ratio is at most 1.02 in all the instances considered and, for $m = 1$,

α \ β	2	3	4	5	6
2	0.5	0.42	0.39	0.37	0.35
3	0.58	0.5	0.46	0.44	0.43
4	0.61	0.54	0.5	0.48	0.46
5	0.63	0.56	0.52	0.5	0.48
6	0.65	0.57	0.54	0.52	0.5

α \ β	2	3	4	5	6
2	(0.2,0.8)	(0.84,0.24)	(0.86,0.24)	(0.86,0.26)	(0.87,0.27)
3	(0.24,0.84)	(0.2,0.8)	(0.82,0.22)	(0.82,0.22)	(0.84,0.24)
4	(0.26,0.86)	(0.22,0.82)	(0.2,0.8)	(0.81,0.21)	(0.82,0.22)
5	(0.26,0.86)	(0.22,0.82)	(0.21,0.81)	(0.2,0.8)	(0.8,0.21)
6	(0.27,0.87)	(0.24,0.84)	(0.22,0.82)	(0.21,0.8)	(0.2,0.8)

Table 1: The optimal percentiles associated to several Beta distributions. In the left table, we report the optimal percentile when $m = 1$ and $q = 0.5$. In the right table, we report the best percentile vectors when $m = 2$ and $q_1 = q_2 = 0.2$.

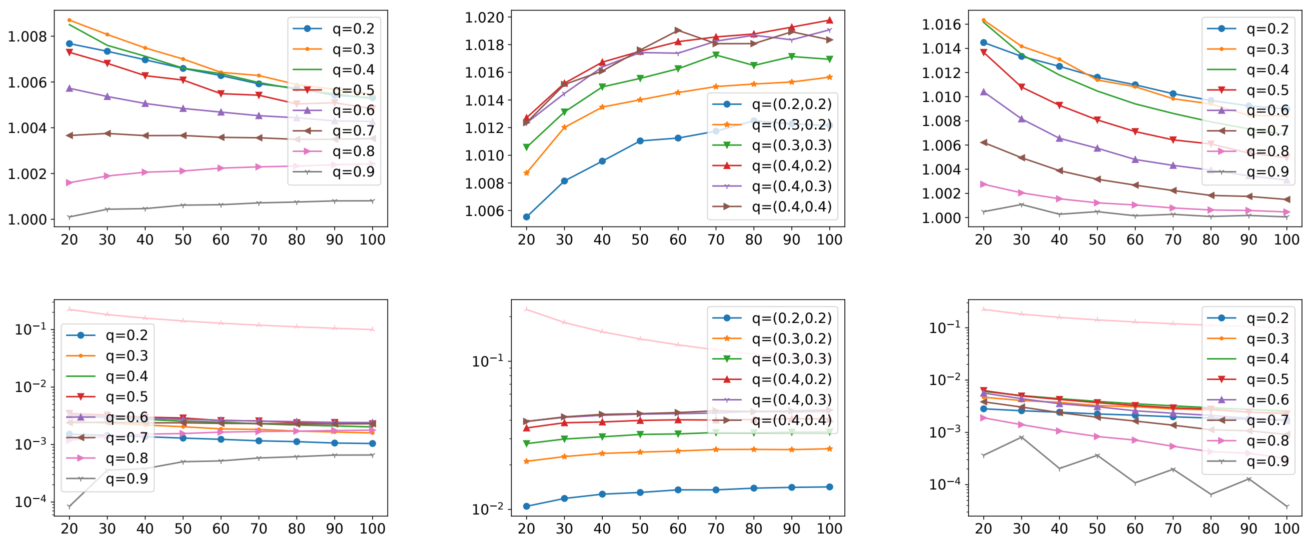

Figure 2: In the first row, we report the Bayesian approximation ratio (y-axis) depending on the number of agents (x-axis) of the mechanism considered. The leftmost and central plot, reports the Bayesian approximation ratio for $\mu \sim \mathcal{B}(6, 2)$, $m = 1$, and $m = 2$, respectively. The rightmost plot reports the Bayesian approximation ratio when the agents are not i.d.. In the second row, we report a logarithmic plot of the absolute error (y-axis) depending on the number of agents (x-axis). In the leftmost and central plot, we report the log-absolute error for $\mu \sim \mathcal{B}(6, 2)$, $m = 1$, and $m = 2$, respectively. The rightmost plot reports the log-absolute error when the agents are not i.d.. In all three figures, we plot the function $n^{-0.5}$ (in pink) for comparison.

it consistently decreases as n increases. This proves that the mechanism found using the bisection method or Algorithm 1 is almost optimal even for small values of n. In the appendix, we report the missing results, which are in line with the analysis presented so far. We then consider the case of non-i.d. agents. We set $\Theta = [0, 1]$, assume that an agent whose type is θ is distributed uniformly on $[0, \theta]$, and set $f_\eta(\theta) = 3\theta^2$. From Theorem 7, we get $f_\mu(x) = \frac{3}{2}(1 - x^2)$. Since $f_{\bar{\mu}}$ is non-increasing, the optimal percentile mechanism is $\frac{q}{2}$ (see Theorem 4). In Figure 2 we plot the Bayesian approximation ratio of the optimal percentile mechanism for non-i.d. agents. In line with Theorem 7, we observe no difference between the results presented for the i.i.d. case.

Convergence Speed. Lastly, we evaluate the speed at which the SW attained by the mechanism converges to its limit, by computing the absolute error of the mechanism $err_{abs} = |\mathbb{E}[SW_{\vec{p}}(\vec{X})] - \mathbb{E}[SW_{opt}(\vec{X})]|$. In Figure 2, we plot the log-scale of the best mechanism's absolute error for $m = 1, 2$ and $\mu \sim \mathcal{B}(6, 2)$, or when the agents are not i.d..

Conclusions and Future Works

In this paper, we studied the Facility Location Problem with Scarce Resources from a Bayesian mechanism design perspective. We introduced a criterion for retrieving optimal mechanisms from the agents' distribution and facility capacities. We then identified the best mechanisms—both analytically and algorithmically—for locating one or two facilities when agents are independent and or not identically distributed. We validated our findings through numerical experiments, proving that mechanisms derived from our theoretical results or routines achieve a low Bayesian approximation ratio that quickly converges to its theoretical limit.

Future work includes extending our search routine to more than two facilities, improving its efficiency, and adapting it to alternative objectives such as Nash welfare. We also aim to identify when the independence assumption on agents' distributions can be relaxed. Finally, exploring a broader class of Equilibrium Stable mechanisms may further improve performance with few agents.

Acknowledgments

Jie Zhang was supported by a Leverhulme Trust Research Project Grant and the EPSRC grant (EP/W014912/1).

References

Aardal, K.; van den Berg, P. L.; Gijswijt, D.; and Li, S. 2015. Approximation algorithms for hard capacitated k-facility location problems. *European Journal of Operational Research*, 242(2): 358–368.

Ahmadi-Javid, A.; Seyedi, P.; and Syam, S. S. 2017. A survey of healthcare facility location. *Computers & Operations Research*, 79: 223–263.

Auricchio, G.; Clough, H. J.; and Zhang, J. 2024. On the capacitated facility location problem with scarce resources. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, UAI '24.

Auricchio, G.; and Zhang, J. 2024. The k-Facility Location Problem via Optimal Transport: A Bayesian Study of the Percentile Mechanisms. In *Algorithmic Game Theory*, 147–164. Springer Nature Switzerland. ISBN 978-3-031-71033-9.

Auricchio, G.; and Zhang, J. 2026. Leveraging Optimal Transport to Design Optimal Mechanisms for the Facility Location Problem. *ACM Transactions on Economics and Computation*.

Auricchio, G.; Zhang, J.; and Zhang, M. 2024. Extended Ranking Mechanisms for the m-Capacitated Facility Location Problem in Bayesian Mechanism Design. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 87–95. IFAAMAS).

Aziz, H.; Chan, H.; Lee, B.; Li, B.; and Walsh, T. 2020a. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34.

Aziz, H.; Chan, H.; Lee, B. E.; and Parkes, D. C. 2020b. The capacity constrained facility location problem. *Games and Economic Behavior*, 124: 478–490.

Balcik, B.; and Beamon, B. M. 2008. Facility location in humanitarian relief. *International Journal of Logistics Research and Applications*, 11(2): 101–121.

Barda, O. H.; Dupuis, J.; and Lencioni, P. 1990. Multicriteria location of thermal power plants. *European Journal of Operational Research*, 45(2-3): 332–346.

Brimberg, J.; Korach, E.; Eben-Chaim, M.; and Mehrez, A. 2001. The Capacitated p-facility Location Problem on the Real Line. *International Transactions in Operational Research*, 8(6): 727–738.

Chawla, S.; and Sivan, B. 2014. Bayesian algorithmic mechanism design. *ACM SIGecom Exchanges*, 13(1): 5–49.

Cuturi, M.; and Doucet, A. 2014. Fast Computation of Wasserstein Barycenters. *Proceedings of Machine Learning Research*, 32(2): 685–693.

Daskalakis, C.; Deckelbaum, A.; and Tzamos, C. 2013. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*.

De Haan, L.; and Taconis-Haantjes, E. 1979. On Bahadur's representation of sample quantiles. *Ann. Inst. Statist. Math*, 31(Part A): 299–308.

Feldman, M.; and Wilf, Y. 2013. Strategyproof facility location and the least squares objective. In *EC*, 873–890. ACM.

Filos-Ratsikas, A.; Li, M.; Zhang, J.; and Zhang, Q. 2017. Facility location with double-peaked preferences. *Auton. Agents Multi Agent Syst.*, 31(6): 1209–1235.

Frogner, C.; Zhang, C.; Mobahi, H.; Araya, M.; and Poggio, T. A. 2015. Learning with a Wasserstein loss. *Advances in neural information processing systems*, 28.

Gairing, M.; Monien, B.; and Tiemann, K. 2005. Selfish routing with incomplete information. In *Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*, 203–212. New York: ACM.

Hartline, J. D.; and Lucier, B. 2010. Bayesian algorithmic mechanism design. In *Proceedings of the forty-second ACM symposium on Theory of computing*, 301–310.

Hartline, J. D.; and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234. New York: ACM.

Levina, E.; and Bickel, P. 2001. The Earth Mover's Distance is the Mallows distance: Some insights from statistics. *Proceedings of the IEEE International Conference on Computer Vision*, 2: 251 – 256 vol.2.

Lucier, B.; and Borodin, A. 2010. Price of anarchy for greedy auctions. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, 537–553.

Meir, R. 2019. Strategyproof Facility Location for Three Agents on a Circle. In *SAGT*, volume 11801 of *Lecture Notes in Computer Science*, 18–33. Berlin: Springer.

Melo, M.; Nickel, S.; and da Gama, F. S. 2009. Facility location and supply chain management – A review. *European Journal of Operational Research*, 196(2): 401–412.

Nisan, N.; and Ronen, A. 1999. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, 129–140. New York: ACM.

Pal, M.; Tardos, T.; and Wexler, T. 2001. Facility location with nonuniform hard capacities. In *Proceedings 42nd IEEE symposium on foundations of computer science*, 329–338.

Pasupathy, R.; and Kim, S. 2011. The stochastic root-finding problem: Overview, solutions, and open questions. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 21(3): 1–23.

Pele, O.; and Werman, M. 2009. Fast and robust Earth Mover's Distances. In *Computer vision, 2009 IEEE 12th international conference on*, 460–467. New Jersey: IEEE.

Procaccia, A. D.; and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)*, 1(4): 1–26.

Rubner, Y.; Tomasi, C.; and Guibas, L. J. 1998. Metric for distributions with applications to image databases. *Proceedings of the IEEE International Conference on Computer Vision*.

Scagliotti, A. 2023. Deep Learning approximation of diffeomorphisms via linear-control systems. *Mathematical Control and Related Fields*, 13(3): 1226–1257.

Sui, X.; Boutilier, C.; and Sandholm, T. 2013. Analysis and Optimization of Multi-Dimensional Percentile Mechanisms. In *IJCAI*, 367–374. Citeseer, Washington: AAAI press.

Villani, C. 2009. *Optimal transport: old and new*, volume 338. Berlin: Springer.