Extended Ranking Mechanisms for the m-Capacitated Facility Location Problem in Bayesian Mechanism Design

Gennaro Auricchio
University of Bath
Bath, United Kingdom
ga647@bath.ac.uk

Jie Zhang
University of Bath
Bath, United Kingdom
jz2558@bath.ac.uk

Mengxiao Zhang
Durham University
Durham, United Kingdom
mengxiao.zhang@durham.ac.uk

ABSTRACT

In this paper, we initiate the study of the m-Capacitated Facility Location Problem (m-CFLP) within a Bayesian Mechanism Design framework. We consider the case in which every agent's private information is their position on a line and assume that every agent's position is independently drawn from a common and known distribution μ. In this context, we propose the Extended Ranking Mechanisms (ERMs), a truthful generalization of the recently introduced Ranking Mechanisms, that allows to handle problems where the total facility capacity exceeds the number of agents. Our primary results pertain to the study of the efficiency guarantees of the ERMs. In particular, we demonstrate that the limit of the ratio between the expected Social Cost of an ERM and the expected optimal Social Cost is finite. En route to these results, we reveal that the optimal Social Cost and the Social Cost of any ERMs can be expressed as the objective value of a suitable norm minimization problem in the Wasserstein space. We then tackle the problem of determining an optimal ERM tailored to a m-CFLP and a distribution μ. Specifically, we aim to identify an ERM whose limit Bayesian approximation ratio is the lowest compared to all other ERMs. We detail how to retrieve an optimal ERM in two frameworks: (i) when the total facility capacity matches the number of agents and (ii) when μ is the uniform distribution and we have two facilities to place. Lastly, we conduct extensive numerical experiments to compare the performance of the ERMs against other truthful mechanisms and to evaluate the convergence speed of the Bayesian approximation ratio. In summary, all our findings highlight that a well-tuned ERM consistently outperforms all other known mechanisms, making it a valid choice for solving the m-CFLP within a Bayesian framework.

KEYWORDS

Bayesian Mechanism Design; Capacitated Facility Location Problem; Optimal Transport

ACM Reference Format:
Gennaro Auricchio, Jie Zhang, and Mengxiao Zhang. 2024. Extended Ranking Mechanisms for the m-Capacitated Facility Location Problem in Bayesian Mechanism Design. 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.

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.

1 INTRODUCTION

Mechanism Design seeks to establish protocols for aggregating the private information of a set of agents to optimize a global objective. Nonetheless, optimizing a communal goal solely based on reported preferences frequently leads to undesirable manipulation driven by the agents' self-interested behaviour. Hence, a key property that a mechanism should possess is *truthfulness*, which ensures that no agent can gain an advantage by misrepresenting their private information. Unfortunately, this stringent condition often clashes with the goal of optimizing the communal objective, leading to suboptimal outcomes from a truthful mechanism. To quantify the efficiency loss entailed by a truthful mechanism, Nisan and Ronen introduced the concept of *approximation ratio*, which represents the maximum ratio between the objective achieved by the truthful mechanism and the optimal objective attainable across all possible agents' reports [37]. A prominent problem in Mechanism Design is the m-Facility Location Problem (m-FLP) [10, 11]. In its fundamental guise, the m-FLP involves locating m facilities amidst n self-interested agents. Each agent requires access to a facility, making it preferable for a facility to be located as close as possible to their position. Furthermore, each facility can accommodate any number of agents, thus the agents can select which facility to use devoid of any concern about possible overload. A natural extension of the m-FLP is the m-Capacitated Facility Location Problem (m-CFLP), in which every facility has a capacity limit. The capacity limit constraints the amount of agents that the facility can serve. In this case, the solution does not only elicit the positions of the facilities, but it also specifies to which facility every agent is assigned to, ensuring that no facility is overloaded. To the best of our knowledge, there are very few works that analyzed the Mechanism Design aspects of the m-CFLP. Moreover, all the existing results are conducted in the classic worst-case analysis, where the designer has no information on the agents and therefore aims to define mechanisms that work well on every possible input, regardless of the likelihood of the input. This type of analysis is, however, too pessimistic and gives little insight into how to select a mechanism for a specific task. For example, there are currently no mechanisms capable of locating more than two capacitated facilities while achieving a finite approximation ratio. Furthermore, even if we restrict our attention to the case where $m = 2$, the approximation ratio of all the known truthful mechanisms depends linearly on the number of agents [5], making these efficiency guarantees less meaningful as the scale of the problem increases. In this paper, we overcome these issues by studying the m-CFLP from a Bayesian viewpoint. In Bayesian Mechanism Design, every agent's location is a random variable whose law is known to the designer [26] thus the scope of the mechanism is not to minimize the cost in the worst possible case but rather to define truthful routines that work well in expectation.

1.1 Our Contribution

Our contribution is as follows:

(1) We define and study the Extended Ranking Mechanisms (ERMs), a generalization of the class of Ranking Mechanisms introduced in [5]. Our extension allows to tackle a broader framework in which the combined capacity of the facilities surpasses the total amount of agents. Moreover, a Ranking Mechanism is truthful if and only if it places all the facilities at, at most, two locations, while a ERM does not necessarily require such constraint to ensure truthfulness. See Section 3.

(2) We then delve into a Bayesian Framework to examine the asymptotic performances of the ERMs. We consider the case in which every agent's position is represented by a set of i.i.d. random variables. Our investigation reveals that m-CFLP is equivalent to a norm minimization problem over a subset of the Wasserstein space [41]. By leveraging the properties of the Wasserstein distance, we then establish the convergence of the Bayesian approximation ratio, i.e., the ratio between the expected cost incurred by the mechanism and the expected optimal cost, as the number of agents tends toward infinity. This is our primary result, as seen in Theorem 4.2. We also compute this limit and show that it depends solely on the specifics of the problem, namely, the probability distribution μ describing the agents, the vector $\vec{q}$ determining the capacities of the facilities, and the characteristics of the mechanism. See Section 4.

(3) After that, we tackle the problem of identifying an optimal ERM tailored to the m-CFLP with agents distributed according to μ given the capacities determined by $\vec{q}$. First, we show that an optimal ERM always exists and characterize it as a solution to a minimization problem. Subsequently, we narrow our focus on two specific scenarios: (i) *no-spare capacity* framework in which the total capacity of the facilities matches the total number of agents, and (ii) the 2-CFLP for a population of agents distributed according to a uniform distribution. See Section 5.

(4) Finally, we validate our findings through extensive numerical experiments in Section 6. In particular, we compare the performances of the ERMs with the performances of other truthful mechanisms, such as the InnerGap Mechanism [42] and the Extended Endpoint Mechanism [5]. From our experiments, we observe that a well-tuned ERM outperforms all the other mechanisms whenever $n \geq 20$. From these comparisons, we also observe that the limit Bayesian approximation ratio is a reliable estimation of the Bayesian approximation ratio of the ERM when the number of agents is greater than 20.

Our outcomes shed light on how to outline a mechanism depending on the problem, contributing to a better understanding of the strategic aspects of the m-CFLP in a Bayesian framework. All the missing proofs and missing materials are available in the full version of the paper (see https://arxiv.org/pdf/2312.16034.pdf, [4]).

1.2 Related Work

The m-Facility Location Problem (m-FLP) and its variations are significant issues in various practical domains, such as disaster relief [7], supply chain management [36], healthcare [1], clustering [28], and public facilities accessibility [8]. Procaccia and Tennenholtz initially delved into the Mechanism Design study of the m-FLP, laying the groundwork for this field in their pioneering work [38]. Following that, a range of mechanisms with constant approximation ratios for placing one or two facilities on trees [21, 22], circles [31, 31, 32], general graphs [2, 17], and metric spaces [35, 40] were introduced. Despite the generality of the underlying space, it is important to stress that all these positive results are confined to the case in which we have at most two facilities to place and/or the number of agents is limited. Moreover, different works tried to generalize the initial framework proposed in [38], by considering different agents' preferences [15, 34], different costs [19], and additional constraints [20, 23].

In this paper, we analyze the m-Capacitated Facility Location Problem (m-CFLP), a variant of the m-FLP in which each facility can accommodate a finite number of agents. The Mechanism Design aspects of the m-CFLP have only recently begun to attract attention. Indeed, the game theoretical framework for the m-CFLP that we consider was introduced in [5]. This work defined and studied various truthful mechanisms, like the InnerPoint Mechanism, the Extended Endpoint Mechanism, and the Ranking Mechanisms. A more theoretical study of the problem was then presented in [42], demonstrating that no mechanism can position more than two capacitated facilities while adhering to truthfulness, anonymity, and Pareto optimality. Additionally, another paper dealing with Mechanism Design aspects of the m-CFLP is [6]. However, it explores a different framework where only one facility needs to be placed and is unable to serve all agents. To the best of our knowledge, all the results on the m-CFLP concerned the classic Mechanism Design framework that evaluates the performances of the mechanism based on worst-case analysis. In this paper, we consider an alternative approach: the Bayesian Mechanism Design perspective. Unlike traditional Mechanism Design, where the designer lacks information about agent types, in the Bayesian Mechanism Design framework each agent's type follows a known probability distribution [14, 25]. In this framework, we are able to determine a probability distribution over the set of all the possible inputs of the mechanism, enabling us to consider the expected cost of a mechanism. Bayesian Mechanism Design has been applied to investigate routing games [24], facility location problems [44], combinatorial mechanisms using ϵ-greedy mechanisms [33], and, notably, auction mechanism design [12, 13, 18, 27, 43].

2 PRELIMINARIES

In this section, we fix the notations on the m-CFLP, Bayesian Mechanism Design, and Optimal Transport (OT).

The m-Capacitated Facility Location Problem. Given a set of self-interested agents $[n] := \{1, \dots, n\}$, we denote with $\vec{x} \in \mathbb{R}^n$ the vector containing the agents' positions. Likewise, given $m \in \mathbb{N}$, we denote with $\vec{c} \in \mathbb{N}^m$ the vector containing the capacities of the facilities, namely $\vec{c} = (c_1, \dots, c_m)$. In this setting, a facility location is defined by three objects: (i) a m-dimensional vector $\vec{y} = (y_1, \dots, y_m)$ whose entries are m positions on the line, (ii) a permutation $\pi \in \mathcal{S}_m$ that decides the capacity of the facility built at y_j, so that $c_{\pi(j)}$ is the capacity of the facility built at y_j, and

(iii) a matching $\Gamma \subset [n] \times [m]$ that determines how the agents are assigned to facilities, i.e. $(i,j) \in \Gamma$ if and only if the agent at x_i is assigned to y_j. Due to the capacity constraints, the degree of vertex $j \in [m]$ according to Γ is at most $c_{\pi(j)}$. Similarly, every agent is assigned to only one facility, thus the degree of every $i \in [n]$ according to Γ is 1. Given the positions of the facilities $\vec{y}$ and a matching Γ, we define the cost of an agent positioned in x_i as $d_{i,\Gamma}(x_i, \vec{y}) = |x_i - y_j|$, where (i,j) is the unique edge in Γ adjacent to i. Given a matching Γ and a permutation $\pi \in \mathcal{S}_m$, a cost function is a map $C_\Gamma : \mathbb{R}^n \times \mathbb{R}^m \to [0,+\infty)$ that associates to $(\vec{x}, \vec{y})$ the overall cost of placing the facilities at the positions in $\vec{y}$ following the permutation π and assigning the agents positioned at $\vec{x}$ according to Γ.[1] Given a vector $\vec{x} \in \mathbb{R}^n$ containing the agents' positions, the *m-Capacitated Facility Location Problem* with respect to the cost C, consists in finding the locations for m facilities, a permutation π, and the matching Γ that minimize the function $\vec{y} \to C(\vec{x}, \vec{y})$. The most studied cost function is the *Social Cost* (*SC*), which is defined as the sum of all the agents' costs. Since multiplying the cost function by a constant does not affect the approximation ratio results, throughout the paper we consider the Social Cost rescaled by the total number of agents, that is $SC(\vec{x}, \vec{y}) = \frac{1}{n}\sum_{i\in[n]} |x_i - y_j|$.

Mechanism Design, the Worst-Case Analysis, and the Ranking Mechanisms. An *m-facility location mechanism* is a function f that takes the agents' reports $\vec{x}$ in input and returns a set of m positions $\vec{y}$ on the line, a permutation $\pi \in \mathcal{S}_m$, and a matching Γ to allocate the agents to the facilities. In general, an agent may misreport its position if it would result in a set of facility locations such that the agent's incurred cost is smaller than reporting truthfully. A mechanism f is said to be *truthful* (or *strategy-proof*) if, for every agent, its cost is minimized when it reports its true position. That is, $d_i(x_i, f(\vec{x})) \le d_i(x_i, f(\vec{x}_{-i}, x_i'))$ for any misreport $x_i' \in \mathbb{R}$, where $\vec{x}_{-i}$ is the vector $\vec{x}$ without its i-th component. Although deploying a truthful mechanism prevents agents from getting a benefit by misreporting their positions, this leads to a loss of efficiency. To evaluate this efficiency loss, Nisan and Ronen introduced the notion of approximation ratio [37]. Given a truthful mechanism f, its approximation ratio with respect to the Social Cost is defined as $ar(f) := \sup_{\vec{x}\in\mathbb{R}^n} \frac{SC_f(\vec{x})}{SC_{opt}(\vec{x})}$, where $SC_f(\vec{x})$ is the Social Cost of placing the facilities and assigning the agents to them following the output of f, while $SC_{opt}(\vec{x})$ is the optimal Social Cost achievable when the agents' report is $\vec{x}$. The *Ranking Mechanisms* are a class of mechanisms for the m-CFLP that work under the assumption that the total capacity of the facilities matches the number of agents, [5]. Each Ranking Mechanism is defined by two parameters: a permutation $\pi \in \mathcal{S}_m$ and a vector $\vec{t} = (t_1, \dots, t_m) \in [n]^m$. Given π and $\vec{t}$ the routine of the Ranking Mechanism is as follows: (i) given $\vec{x}$ the vector containing the agents' reports ordered non-decreasingly, then the mechanism places the facility with capacity $c_{\pi(j)}$ at x_{t_j}. (ii) The agents are assigned to the facility from left to right while respecting the capacity constraints. It was shown in [5] that a Ranking Mechanism is truthful if and only if every t_j admits at most two different values. Moreover, the approximation ratio of these mechanisms is bounded only when the number of agents is even, the number of facilities to places is 2, and the two facilities have the same capacity. In this case, the mechanism is also called InnerPoint Mechanism (IM) and its approximation ratio is $\frac{n}{2} - 1$.

Bayesian Mechanism Design. In Bayesian Mechanism Design, every agent's type is described by a random variable X_i. In what follows, we assume that every X_i is identically distributed according to a law μ and independent from the other random variables. In this framework, a mechanism is truthful if, for every agent i, it holds

$$\mathbb{E}_{\vec{X}_{-i}}[d_i(x_i, f(x_i, \vec{X}_{-i}))] \le \mathbb{E}_{\vec{X}_{-i}}[d_i(x_i, f(x_i', \vec{X}_{-i}))],\ \forall x_i \in \mathbb{R}, \quad (1)$$

where x_i agent i's true type, $\vec{X}_{-i}$ is the $(n-1)$-dimensional random vector that describes the other agents' type, and $\mathbb{E}_{\vec{X}_{-i}}$ is the expectation with respect to the joint distribution of $\vec{X}_{-i}$. It is easy to see that if a mechanism is truthful in the classic Mechanism Design framework, it is also truthful in the Bayesian framework. Given $\beta \in \mathbb{R}$, a mechanism f is a β-approximation if it holds $\mathbb{E}[SC_f(\vec{X}_n)] \le \beta\,\mathbb{E}[SC_{opt}(\vec{X}_n)]$, where $\mathbb{E}$ is the expectation with respect to the joint distribution of $\vec{X}_n$. Similarly to what happens for the approximation ratio, the lower β is, the better the mechanism is. To unify the notation, we define the *Bayesian approximation ratio* of a mechanism f as the ratio between the expected Social Cost of a mechanism and the expected optimal Social Cost. More formally, given a mechanism f, its Bayesian approximation ratio with respect to the Social Cost is defined as follows

$$B_{ar}(f) := \frac{\mathbb{E}[SC_f(\vec{X}_n)]}{\mathbb{E}[SC_{opt}(\vec{X}_n)]}, \quad (2)$$

where the expected value is taken over the joint distribution of the vector $\vec{X}_n := (X_1, \dots, X_n)$. Notice that, if $B_{ar}(f) < +\infty$, then f is a $B_{ar}(f)$-approximation.

The Wasserstein Distance. Let us denote with $\mathcal{P}(\mathbb{R})$ the set of probability measures over $\mathbb{R}$. Given $\gamma \in \mathcal{P}(\mathbb{R})$, we denote with $spt(\gamma) \subset \mathbb{R}$ the support of γ and with $med(\gamma)$ the smallest median of γ. We denote with $\mathcal{P}_m(\mathbb{R})$ the set of probability measures over $\mathbb{R}$ whose support contains at most m points, thus a measure $\nu \in \mathcal{P}_m(\mathbb{R})$ is such that $\nu = \sum_{j=1}^m \nu_j \delta_{x_j}$, where $x_j \in \mathbb{R}$ for all $j \in [m]$, ν_j are non-negative values satisfying $\sum_{j=1}^m \nu_j = 1$, and δ_{x_j} is the Dirac delta measure centered at x_j. Given two measures $\alpha, \beta \in \mathcal{P}(\mathbb{R})$, the Wasserstein distance between α and β is defined as $W_1(\alpha,\beta) = \min_{\pi\in\Pi(\alpha,\beta)} \int_{\mathbb{R}\times\mathbb{R}} |x-y| d\pi$, where $\Pi(\alpha,\beta)$ is the set of probability measures over $\mathbb{R}\times\mathbb{R}$ whose first marginal is equal to α and the second marginal is equal to β [29]. For a comprehensive introduction to Optimal Transport theory, we refer to [41] and [39].

Basic Assumptions. Finally, we layout the basic assumptions of our framework. In what follows, we tacitly assume that the underlying distribution μ satisfies all the following properties: (i) The measure μ is absolutely continuous. We denote with ρ_μ its density. (ii) The support of μ is an interval, which can be bounded or not, and that ρ_μ is strictly positive on the interior of the support. (iii) The density function ρ_μ is differentiable on the support of μ. (iv) The probability measure μ has finite first moment, i.e. $\int_{\mathbb{R}} |x| d\mu < +\infty$. Notice that, according to this set of assumptions, the cumulative distribution function (c.d.f.) of μ, namely F_μ, is locally bijective.

[1] In what follows, we omit Γ from the indexes of d and C if it is clear from the context which matching we are considering.

3 THE EXTENDED RANKING MECHANISMS

Since our study aims to examine the behaviour of the mechanism as the number of agents goes to infinity, expressing the problem in terms of absolute capacities is unsuitable. For this reason, we need to rephrase the specifics of the m-CFLP in terms of percentages. In particular, instead of considering a collection of capacities $c_j \in \mathbb{N}$ for each $j \in [m]$, we shift our focus to the percentage capacity vector (p.c.v.) $\vec{q} \in (0,1)^m$. Each value q_j corresponds to the percentage of agents that the j-th facility can accommodate. Given $\vec{q}$ and a number of agents n, we recover the absolute capacity of the j-th facility by setting $c_j = \lfloor q_j(n-1) \rfloor + 1$. Conversely, when we are given the absolute capacities c_j for a given number of agents n, the corresponding p.c.v. is $q_j = \frac{c_j}{n}$. Without loss of generality, we assume that the entries of $\vec{q}$ are ordered non-increasingly.

Mechanism 1 (Extended Ranking Mechanisms (ERMs)). *Let $\vec{q} = (q_1, \dots, q_m)$ be a p.c.v.. Given a permutation $\pi \in \mathcal{S}_m$ and an non-decreasingly ordered vector $\vec{p} = (p_1, \dots, p_m) \in [0,1]^m$, that is $p_j \le p_{j+1}$, the routine of the ERM associated with π and $\vec{p}$, namely $\mathsf{ERM}^{(\pi,\vec{p})}$, is as follows: (i) First, we collect the reports of the agents and order them non-decreasingly. We denote with $\vec{x}$ the ordered vector containing the agents' reports, thus $x_i \le x_{i+1}$. (ii) Second, we elicit m positions on the line by setting $y_j = x_{\lfloor p_j(n-1) \rfloor + 1}$ for every $j \in [m]$ and place the facility with capacity $q_{\pi(j)}$ at y_j. (iii) Finally, we assign every agent to the facility closer to the position they reported (break ties arbitrarily without overloading the facilities).*

Notice that, in the routine of the ERM, the vector $\vec{p}$ plays the role that $\vec{t}$ plays in the routine of a Ranking Mechanism. Indeed, the main difference between the Ranking Mechanisms and our generalization lies in how the mechanism matches the agents to the facilities. While the Ranking Mechanisms assign the agents monotonically from left to right, the ERM assigns every agent to the facility that is closer to their report. Depending on the pair $(\pi, \vec{p})$ and $\vec{q}$, however, the matching returned by the ERM might overload some of the facilities. We say that a couple $(\pi, \vec{p})$ induces a *feasible* ERM if, for every $n \in \mathbb{N}$ and every $\vec{x} \in \mathbb{R}^n$, the output of $\mathsf{ERM}^{(\pi,\vec{p})}$ is a facility location for the m-CFLP induced by $\vec{q}$. Given a p.c.v. $\vec{q}$, the set of parameters $(\pi, \vec{p})$ that induce a feasible ERM is characterizable through a system of inequalities. For the sake of simplicity, we first consider the case in which $\vec{p}$ does not have two equal entries, that is $p_j \neq p_l$ for every $j \neq l$.

Theorem 3.1. *Given $\vec{q}$ and $\vec{p}$ such that $p_j \neq p_i$ for every $j \neq i \in [m]$, then $\mathsf{ERM}^{(\pi,\vec{p})}$ is feasible if and only if the following system of inequalities are satisfied*

$$\begin{cases} q_{\pi(1)} \ge p_2 \\ q_{\pi(2)} \ge p_3 - p_1 \\ \quad\vdots \\ q_{\pi(m-1)} \ge p_m - p_{m-2} \\ q_{\pi(m)} \ge 1 - p_{m-1} \end{cases}. \tag{3}$$

The feasibility of the ERM ensures also its truthfulness.

Theorem 3.2. *Given a p.c.v. $\vec{q}$, any feasible $\mathsf{ERM}^{(\pi,\vec{p})}$ is truthful. Thus it is also truthful in the Bayesian framework.*

The truthfulness of any ERM is bound to the fact that every facility y_j can accommodate all the agents between y_{j-1} and y_{j+1}. This constraint limits the set of $\vec{p}$ for which there exists a $\pi \in \mathcal{S}_m$ such that $\mathsf{ERM}^{(\pi,\vec{p})}$ is feasible.

Corollary 3.2.1. *Given a p.c.v. $\vec{q} \in (0,1)^m$, let us fix $\vec{p} \in [0,1]^m$. Then, if $(\max_{j\in[m]} p_j - \min_{j\in[m]} p_j) > \sum_{j\in[m]} q_j - 1$, the mechanism $\mathsf{ERM}^{(\pi,\vec{p})}$ is not feasible for every $\pi \in \mathcal{S}_m$.*

Finally, we notice that if $\vec{p}$ has at least two equal entries, not all the feasible ERMs necessarily satisfy system (3). For example, the `all-median` mechanism that places all the facilities at the median agent is always feasible and truthful.[2] However, depending on $\vec{q}$, the vector $\vec{p} = (0.5, \dots, 0.5)$ may not satisfy system (3) for any $\pi \in \mathcal{S}_m$. Indeed, if $p_j = p_l$ for some indices $j, l \in [m]$, we need to consider all the facilities placed at $x_{\lfloor p(n-1) \rfloor + 1}$ as if they were a unique facility whose capacity is the total capacity of the facilities placed at $x_{\lfloor p(n-1) \rfloor + 1}$. By doing so, we are able to extend Theorem 3.1 and 3.2 to the case in which $\vec{p}$ has at least two equal entries. In particular, given $\vec{p} \in [0,1]^m$, let $\vec{p}' \in [0,1]^{m'}$ be the vector containing all the different entries of $\vec{p}$, then, the mechanism $\mathsf{ERM}^{(\pi,\vec{p})}$ is feasible if and only if the following system is satisfied

$$\begin{cases} \sum_{l\in[m];\, s.t.\, p_{\pi(l)}=p'_1} q_{\pi(l)} \ge p'_2 \\ \sum_{l\in[m];\, s.t.\, p_{\pi(l)}=p'_2} q_{\pi(l)} \ge p'_3 - p'_1 \\ \quad\vdots \\ \sum_{l\in[m];\, s.t.\, p_{\pi(l)}=p'_{m-1}} q_{\pi(l)} \ge p'_{m'} - p'_{m'-2} \\ \sum_{l\in[m];\, s.t.\, p_{\pi(l)}=p'_{m'}} q_{\pi(l)} \ge 1 - p'_{m'-1} \end{cases}.$$

Lastly, we notice that, except in a few specific cases, the approximation ratio of the Ranking Mechanisms is unbounded. Consequentially, it is impossible to retrieve a bound on the approximation ratio of any ERM for generic m-CFLPs. Since the classic worst-case analysis does not give any insight into the performances of the ERMs, we move our attention to the Bayesian analysis.

4 THE BAYESIAN ANALYSIS OF THE ERMS

In this section, we present our main result, which characterizes the limit of the Bayesian approximation ratio of any feasible ERM as a function of $\vec{p}$, π, μ, and $\vec{q}$. As a preliminary lemma, we relate the m-CFLP to a norm minimization problem in the Wasserstein Space.

Lemma 4.1. *Given a p.c.v. $\vec{q}$, let $\vec{x} \in \mathbb{R}^n$ be the vector containing the agents' reports. Let us fix $\mu_{\vec{x}} := \frac{1}{n}\sum_{i\in[n]} \delta_{x_i}$. Then, it holds*

$$SC_{opt}(\vec{x}) = \min_{\sigma\in\mathcal{S}_m} \min_{\zeta\in\mathcal{P}_{\sigma,\vec{q}}(\mathbb{R})} W_1(\mu_n, \zeta), \tag{4}$$

where $\mathcal{P}_{\sigma,\vec{q}}(\mathbb{R})$ is the set of probability measures such that $\zeta = \sum_{j\in[m]} \zeta_j \delta_{y_j}$, where $y_1 \le y_2 \le \dots \le y_m$ and $\zeta_j \le q_{\sigma(j)}$ for every $j \in [m]$. Similarly, given a permutation $\pi \in \mathcal{S}_m$ and a vector $\vec{p}$ that induce a feasible ERM, it holds

$$SC_{\pi,\vec{p}}(\vec{x}) = \min_{\lambda_j \le q_{\pi(j)}, \sum_{j\in[m]}\lambda_j = 1} W_1(\mu_n, \lambda), \tag{5}$$

where $SC_{\pi,\vec{p}}(\vec{x})$ is the Social Cost attained by $\mathsf{ERM}^{(\pi,\vec{p})}$ on instance $\vec{x}$, $\lambda = \sum_{j\in[m]} \lambda_j \delta_{x_{r_j}}$, and $r_j = \lfloor p_j(n-1) \rfloor + 1$.

[2] In this case, we do not need to specify which agent is served by which facility nor $\pi \in \mathcal{S}_m$, since all the agents and all the facilities are served/located at the same place.

The proof of Lemma 4.1 consists in showing that, for any given instance $\vec{x}$ and given an optimal facility location for the m-CFLP, it is possible to construct a measure ν such that $W_1(\mu_n, \nu) = SC_{opt}(\vec{x})$ and, vice-versa, given a measure ν that minimizes (4), it is possible to retrieve an optimal facility location problem for the m-CFLP. Since the m-CFLP admits a solution, problem (4) is well-posed and admits a solution. By a similar argument, we infer the same conclusions for problem (5). The connection between the m-CFLP and Optimal Transport theory highlighted in Lemma 4.1 enables us to exploit the properties of the Wasserstein distances and to characterize the limit Bayesian approximation ratio of every feasible ERM.

Theorem 4.2. *Given the p.c.v. $\vec{q}$, let $\vec{p} \in (0,1)^m$ and $\pi \in \mathcal{S}_m$ be such that $\mathsf{ERM}^{(\pi,\vec{p})}$ is feasible. Then, it holds*

$$\lim_{n\to+\infty} \frac{\mathbb{E}[SC_{\pi,\vec{p}}(\vec{x})]}{\mathbb{E}[SC_{opt}(\vec{x})]} = \frac{W_1(\mu, \nu_{\vec{p}})}{W_1(\mu, \nu_m)}, \tag{6}$$

where ν_m is a solution to the following minimization problem

$$\min_{\sigma\in\mathcal{S}_m} \min_{\zeta\in\mathcal{P}_{\sigma,\vec{q}}(\mathbb{R})} W_1(\mu, \zeta), \tag{7}$$

and $\nu_{\vec{p}}$ is a solution to

$$\min_{\lambda_j \le q_{\pi(j)}, \sum_{j\in[m]}\lambda_j = 1} W_1\Big(\mu, \sum_{j\in[m]} \lambda_j \delta_{F_\mu^{[-1]}(p_j)}\Big). \tag{8}$$

Sketch of the proof. The proof consists of three steps: (i) First, we show that the expected optimal Social Cost for the m-CFLP converges to the objective value of the minimization problem in (7). (ii) Second, we show that the expected Social Cost of $\mathsf{ERM}^{(\pi,\vec{p})}$ converges to the objective value of the minimization problem in (8). (iii) We combine the two convergence results to retrieve (6).

The limit of the expected optimal Social Cost. Let ν_m be such that $W_1(\mu, \nu_m) = \min_{\sigma\in\mathcal{S}_m} \min_{\zeta\in\mathcal{P}_{\sigma,\vec{q}}(\mathbb{R})} W_1(\mu, \zeta)$. From Lemma 4.1, we know that, for every $n \in \mathbb{N}$ and for every $\vec{x} \in \mathbb{R}^n$, there exists a $\nu_{\vec{x},m}$ such that $SC_{opt}(\vec{x}) = W_1(\mu_n, \nu_{\vec{x},m})$. By definition of ν_m, we have that $W_1(\mu, \nu_m) \le W_1(\mu, \nu_{\vec{x},m})$. Since W_1 is a distance, it holds $W_1(\mu, \nu_m) \le W_1(\mu, \nu_{\vec{x},m}) \le W_1(\mu, \mu_n) + W_1(\mu_n, \nu_{\vec{x},m})$. By rearranging the terms and by taking the expected value with respect to the distribution of $\vec{X}$, we obtain $\mathbb{E}[W_1(\mu, \nu_m)] - \mathbb{E}[W_1(\mu_n, \nu_{\vec{x},m})] \le \mathbb{E}[W_1(\mu, \mu_n)]$. By a similar argument, we have that $W_1(\mu_n, \nu_{\vec{x},m}) \le W_1(\mu_n, \nu_m) \le W_1(\mu, \mu_n) + W_1(\mu, \nu_m)$, hence $\mathbb{E}[W_1(\mu_n, \nu_{\vec{x},m})] - \mathbb{E}[W_1(\mu, \nu_m)] \le \mathbb{E}[W_1(\mu, \mu_n)]$. We then infer that

$$|\mathbb{E}[W_1(\mu, \nu_m)] - \mathbb{E}[W_1(\mu_n, \nu_{\vec{x},m})]| \le \mathbb{E}[W_1(\mu, \mu_n)].$$

Since the right handside of this inequality converges to 0 as n goes to $+\infty$ (see [9]), we infer that $\lim_{n\to\infty} \mathbb{E}[W_1(\mu_n, \nu_{\vec{x},m})] = W_1(\mu, \nu_m)$, which concludes the first part of the proof.

The limit of the expected Social Cost of the Mechanism. The argument used for this part is similar to the one used for the limit expected optimal Social Cost but more delicate. Indeed, in this case, the set on which we minimize the Wasserstein distance does depend on the agents' report $\vec{x}$. In particular, the sets on which are formulated problem (8) and (5) are, in general, different. To overcome this issue, we need to define two auxiliary probability measures, namely ϕ and ψ. Given $\vec{x} \in \mathbb{R}^n$, let $\nu_{\vec{p}}$ and $\nu_{\vec{x},\vec{p}}$ be the solutions to problem (8) and (5), respectively. We define the measures $\phi = \sum_{j\in[m]} (\nu_{\vec{x},\vec{p}})_j \delta_{y_j}$, where y_j is the support of the measure $\nu_{\vec{p}}$. For every $n \in \mathbb{N}$ and every $\vec{x} \in \mathbb{R}^n$, we have that $W_1(\mu, \nu_{\vec{p}}) \le W_1(\mu, \phi) \le W_1(\mu, \mu_n) + W_1(\mu_n, \nu_{\vec{x},\vec{p}}) + W_1(\nu_{\vec{x},\vec{p}}, \phi)$. We therefore infer

$$W_1(\mu, \nu_{\vec{p}}) - W_1(\mu_n, \nu_{\vec{x},\vec{p}}) \le W_1(\mu, \mu_n) + W_1(\nu_{\vec{x},\vec{p}}, \phi). \tag{9}$$

Similarly, given $\vec{x} \in \mathbb{R}^n$, we define $\psi = \sum_{j\in[m]} (\nu_{\vec{p}})_j \delta_{y_{\vec{x},j}}$, where $\{y_{\vec{x},j}\}_{j\in[m]}$ is the support of $\nu_{\vec{x},\vec{p}}$. We then have

$$W_1(\mu_n, \nu_{\vec{x},\vec{p}}) - W_1(\mu, \nu_{\vec{p}}) \le W_1(\mu, \mu_n) + W_1(\nu_{\vec{p}}, \psi). \tag{10}$$

Since the Wasserstein distance is always non negative, we can combine the estimations in (9) and (10), to obtain

$$|W_1(\mu_n, \nu_{\vec{x},\vec{p}}) - W_1(\mu, \nu_{\vec{p}})| \le W_1(\mu, \mu_n) + W_1(\nu_{\vec{p}}, \psi) + W_1(\nu_{\vec{x},\vec{p}}, \phi).$$

If we take the expectation of both sides of the inequality the inequality still holds. Thus, if we show that $\lim_{n\to\infty} \mathbb{E}[W_1(\nu_{\vec{p}}, \psi)] = \lim_{n\to\infty} \mathbb{E}[W_1(\nu_{\vec{x},\vec{p}}, \phi)] = 0$, we conclude this second step of the proof since, by [9], we have that $\lim_{n\to\infty} \mathbb{E}[W_1(\mu_n, \mu)] = 0$. Let us consider $\mathbb{E}[W_1(\nu_{\vec{p}}, \psi)]$, the convergence of $\mathbb{E}[W_1(\nu_{\vec{x},\vec{p}}, \phi)]$ follows by a similar argument. We notice that ψ and $\nu_{\vec{p}}$ have different supports, but $\psi_j = (\nu_{\vec{p}})_j$, thus it holds $\mathbb{E}[W_1(\nu_{\vec{p}}, \psi)] \le \sum_{j\in[m]} \psi_j \mathbb{E}[|y_j - y_{\vec{x},j}|]$, where $y_{\vec{x},j}$ is the j-th point in the support of $\nu_{\vec{x},\vec{p}}$. By definition of ERM, it holds $y_{\vec{x},j} = x_{\lfloor p_j(n-1)\rfloor+1}$. Since the $(\lfloor p_j(n-1)\rfloor+1)$-th order statistics converges to the p_j-th quantile of μ [16, 30], we have that $\mathbb{E}[|y_j - y_{\vec{x},j}|] \to 0$ as $n \to \infty$, which concludes the second part of the proof.

Characterizing the Bayesian approximation ratio. To conclude, notice that the distance between an absolutely continuous measure and a discrete measure is always greater than zero, thus $\lim_{n\to\infty} \mathbb{E}[SC_{opt}(\vec{X})] > 0$. For this reason, we have that the limit of the ratio is equal to the ratio of the limits, which proves (6). □

Notice that our result applies only to feasible $\mathsf{ERM}^{(\pi,\vec{p})}$, such that $\vec{p} \in (0,1)^m$, since, for general measures μ, the values $F_\mu^{[-1]}(0)$ and $F_\mu^{[-1]}(1)$ might not be finite. Finally, we notice that the Bayesian approximation ratio is invariant to positive affine transformation of μ. In particular, the limit of the Bayesian approximation ratio remains the same across all the Gaussian-distributed populations.

Corollary 4.2.1. *Let $\vec{q}$ be a p.c.v. and let X be the random variable associated with μ. Given $\alpha > 0$ and $\beta \in \mathbb{R}$, let $\mu_{\alpha,\beta}$ be the probability distribution associated with $\alpha X + \beta$. Then, the asymptotical Bayesian approximation ratio of any feasible ERM is the same regardless of whether the agent type is distributed according to μ or $\mu_{\alpha,\beta}$.*

5 HOW TO SELECT AN OPTIMAL ERM

As shown in Theorem 4.2, the limit Bayesian approximation ratio of any ERM hinges upon μ, $\vec{q}$, π, and $\vec{p}$. While μ and $\vec{q}$ are beyond the control of the mechanism designer, both π and $\vec{p}$ serve as parameters that can be tuned depending upon μ and $\vec{q}$. In this section, we study how to determine an optimal ERM tailored to μ and $\vec{q}$. Given μ and $\vec{q}$, we say that a feasible $\mathsf{ERM}^{(\pi,\vec{p})}$ is optimal if

$$\lim_{n\to\infty} B_{ar}\big(\mathsf{ERM}^{(\pi,\vec{p})}\big) \le \lim_{n\to\infty} B_{ar}\big(\mathsf{ERM}^{(\pi',\vec{p}')}\big)$$

for any other feasible $\mathsf{ERM}^{(\pi',\vec{p}')}$. The main result of the section assesses that an optimal ERM exists and characterizes its parameters $(\pi, \vec{p})$ as the solution to a suitable minimization problem.

Theorem 5.1. *Given μ and a p.c.v. $\vec{q}$, there always exist a tuple $(\pi, \vec{p})$ whose associated ERM, that is $\mathsf{ERM}^{(\pi,\vec{p})}$, is optimal. Moreover, the couple $(\pi, \vec{p})$ is a solution to the following minimization problem*

$$\min_{\pi \in \mathcal{S}_m} \min_{\vec{p} \in (0,1)^m} W_1\Big(\mu, \sum_{j \in [m]} \eta_j \delta_{F_\mu^{[-1]}(p_j)}\Big) \tag{11}$$

$$\text{such that } p_{j+1} - p_{j-1} \le q_{\pi(j)} \quad \text{for} \quad j = 2, \dots, m-1,$$

$$p_1 \le q_{\pi(q)}, \text{ and } p_m \le 1 - q_{\pi(m)},$$

where $\eta_1 = F_\mu\Big(\frac{F_\mu^{[-1]}(p_2) + F_\mu^{[-1]}(p_1)}{2}\Big)$, $\eta_j = F_\mu\Big(\frac{F_\mu^{[-1]}(p_{j+1}) + F_\mu^{[-1]}(p_j)}{2}\Big) - F_\mu\Big(\frac{F_\mu^{[-1]}(p_j) + F_\mu^{[-1]}(p_{j-1})}{2}\Big)$ *for* $j = 2, \dots, m-1$*, and* $\eta_m = 1 - \eta_{m-1}$.

Notice that the limit Bayesian approximation ratio of the optimal ERM does not necessarily converge to 1, as the next example shows.

Example 1. *Let μ be the uniform distribution $\mathcal{U}[0, 1]$ and let $\vec{q} = (0.8, 0.4)$ be a p.c.v.. Since μ is symmetric, one of the solutions to problem (7) is $\nu_2 = (0.4)\delta_{0.2} + (0.6)\delta_{0.7}$ (the other one is $\nu_2' = (0.6)\delta_{0.3} + (0.4)\delta_{0.8}$). However, by Corollary 3.2.1, there does not exist a permutation $\pi \in \mathcal{S}_m$ such that $\mathsf{ERM}^{((0.2,0.7),\pi)}$ is feasible. Thus, by Theorem 4.2, no feasible ERM is such that $\lim_{n\to\infty} B_{ar}(\mathsf{ERM}^{(\pi,\vec{p})}) = 1$.*

We now characterize the optimal ERM in two specific cases. In the first one, the total capacity of the facilities is the same as the number of agents. In the second one, we need to place two capacitated facilities and μ is a uniform distribution.

5.1 The No-Spare Capacity Case

In the no-spare capacity case, the total capacity of the facilities matches the number of agents, thus $\sum_{j \in [m]} q_j = 1$. Due to Corollary 3.2.1, we have that the only feasible ERMs are the ones for which it holds $p_j = p \in [0, 1]$ for every $j \in [m]$. Thus, owing to the properties of the W_1 distance and to Lemma 4.1, the optimal ERM is the `all-median` mechanism, i.e. $\mathsf{ERM}^{(Id,(0.5,\dots,0.5))}$.

Theorem 5.2. *In the no-spare capacity case, the optimal ERM is unique and is the* `all-median` *mechanism.*

When the agents are distributed following a uniform distribution, that is $\mu = \mathcal{U}[0, 1]$, we can express the limit Bayesian approximation ratio of the `all-median` mechanism as a function of $\vec{q}$. Indeed, since $\sum_{j \in [m]} q_j = 1$, any solution to problem (7) separates the interval $[0, 1]$ into m intervals whose length is q_j. Notice that the order in which $[0, 1]$ is divided is irrelevant. Furthermore, the facility is placed at the median of such interval, thus the objective value of (7) is $\frac{1}{4}\sum_{j \in [m]} q_j^2$. By Theorem 4.2, the limit Bayesian approximation ratio of the `all-median` mechanism is $\lim_{n\to\infty} B_{ar}(\texttt{all-median}) = (\sum_{j \in [m]} q_j^2)^{-1}$, since the asymptotic cost of placing all the facilities at the median point is $\int_0^1 |x - 0.5| dx = \frac{1}{4}$. Notice that the limit Bayesian approximation ratio gets closer to 1 as the values of $\vec{q}$ become concentrated at one index. Conversely, if all the facilities have the same capacity $\frac{1}{m}$, we have that the Bayesian approximation ratio of the `all-median` becomes the largest possible, that is $\lim_{n\to\infty} B_{ar}(\mathsf{ERM}^{(\pi,\vec{p})}) = m$.

Comparing the Ranking Mechanisms with the ERMs. To the best of our knowledge, the only other truthful mechanisms capable of placing more than 2 facilities in the no-spare capacity case are the Ranking Mechanisms [5]. To close the section, we show that the `all-median` mechanism is also the asymptotically best possible Ranking Mechanism. Indeed, by [5], we know that a truthful Ranking Mechanism either (i) puts all the facilities at the same position or (ii) places all the facilities at two adjacent agents' reports, namely x_t and x_{t+1}. However, not all the values of t are feasible, as it must exist $J' \subset [m]$ such that $t = \sum_{j \in J'} c_j$. Owing to Theorem 5.2, in case (i), the best possible mechanism is the `all-median` mechanism. In case (ii) the limit Bayesian approximation ratio of any mechanism is either equal or higher (see Appendix C). Indeed, consider a population distributed as a Uniform distribution, i.e. $\mu = \mathcal{U}[0, 1]$, and $\vec{q} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and the Ranking Mechanism that places q_1 at $x_{\lfloor \frac{n-1}{3} \rfloor + 1}$ and the other two facilities at $x_{\lfloor \frac{n-1}{3} \rfloor + 2}$. Since every agent is assigned to its closest facility, it is possible to adapt Lemma 4.1 and Theorem 4.2 and show that the limit of the expected Social Cost of the mechanism is $W_1(\mu, \delta_{\frac{1}{3}}) = \frac{5}{18}$. Hence, the limit Bayesian approximation ratio of the considered Ranking Mechanism is ~ 3.33, while $\lim_{n\to\infty} B_{ar}(\texttt{all-median}) = 3$.

5.2 Placing two capacitated facilities for a uniform population.

In this section, we retrieve the optimal ERMs when $\mu = \mathcal{U}[0, 1]$ and $m = 2$. For the sake of simplicity, we divide the discussion into three steps. First, we demonstrate that both permutations in $\mathcal{S}_2$ yield equally optimal ERMs. Second, we explicit the objective function of problem (11) as a function of $\vec{p}$ and compute its derivatives. Lastly, we retrieve the optimal ERM for any given $\vec{q}$ using Lagrangian multipliers. In Appendix B, we show how to generalize this process to the case in which μ is symmetric and other relevant frameworks.

First, notice that, since $F^{[-1]}_{\mathcal{U}[0,1]}(p) = p$ for every $p \in [0, 1]$, the objective value of problem (11) boils down to

$$\mathcal{W}(\vec{p}) := W_1\Big(\mu, \frac{p_1 + p_2}{2}\delta_{p_1} + \big(1 - \frac{p_1 + p_2}{2}\big)\delta_{p_2}\Big).$$

Step 1. Since $m = 2$, $\mathcal{S}_2 = \{Id, \theta\}$, where the permutation Id is such that $Id(i) = i$, and θ switches 1 and 2. First, we show if $(Id, \vec{p})$ satisfies system (3), there exists a vector $\vec{p}'$ such that $(\theta, \vec{p}')$ satisfies (3) and $\mathcal{W}(\vec{p}) = \mathcal{W}(\vec{p}')$. Given $\vec{p} = (p_1, p_2)$, let us define $\vec{p}' = (p_1', p_2') = (1 - p_2, 1 - p_1)$. If $(Id, \vec{p})$ satisfies system (3), we have that $p_2 \le q_{Id(1)} = q_1$ and $1 - p_1 \le q_{Id(2)} = q_2$. It is then easy to see that $p_2' = 1 - p_1 \le q_2 = q_{\theta(1)}$ and, likewise $1 - p_1' \le q_{\theta(2)}$. Thus, $(\theta, \vec{p}')$ satisfies (3). Finally, due to the symmetry of μ with respect to 0.5, we have that $\mathcal{W}(\vec{p}) = \mathcal{W}(\vec{p}')$. Therefore, for both Id and θ, there exists a vector $\vec{p}$ that induces an optimal ERM.

Step 2. Let us fix $\pi = Id$. Due to the properties of the optimal transportation plan on the line [39], have that $\mathcal{W}(\vec{p}) = \int_0^1 \min\{|x - p_1|, |x - p_2|\} dx$, thus $\mathcal{W}(\vec{p}) = \frac{p_1^2}{2} + \frac{(1-p_2)^2}{2} + \frac{(p_2 - p_1)^2}{4}$. From a simple computation, we infer that the objective function $\mathcal{W}(\vec{p})$ is differentiable and retrieve the formula of its gradient as $\nabla\mathcal{W}(\vec{p}) = \frac{1}{2}\big(3p_1 - p_2, 3p_2 - 2 - p_1\big)$. Notice that the set of points on which the first derivative of $\mathcal{W}$ nullifies is the line $p_2 = 3p_1$, while the set of points on which the second derivative of $\mathcal{W}$ nullifies is $p_2 = \frac{1}{3}(p_1 + 2)$, thus $\nabla\mathcal{W}(\vec{p}) = (0, 0)$ if and only if $\vec{p} = (0.25, 0.75)$.

Step 3. Lastly, to detect an optimal ERM, we need to implement the feasibility constraints described in Theorem 3.1. From Step 1, we focus on the set of constraints induced by $Id \in \mathcal{S}_m$. In this case, we have that $\vec{p}$ must satisfy the following constraints (i) $p_2 \leq q_1$, (ii) $1-p_1 \leq q_2$ that is $1-q_2 \leq p_1$, and (iii) $p_1 \leq p_2$. Thus the set of feasible $\vec{p}$ is a triangle, namely $T(\vec{q})$, whose vertexes are (q_1, q_1), $(1-q_2, 1-q_2)$, and $(1-q_2, q_1)$. Finally, we retrieve the best vector $\vec{p}$ depending on the value of $\vec{q}$. First, notice that if $(0,25,0.75) \in T(\vec{q})$, i.e. if and only if $q_1, q_2 \geq 0.75$, then the minimum is attained in $\vec{p} = (0,25,0.75)$ as it is the only point in which the gradient $\nabla\mathcal{W}$ nullifies and the Hessian matrix is positive definite. For all the other values of $\vec{q}$, we need to search the minimum over the boundary of $T(\vec{q})$, since the gradient is never null on the interior of $T(\vec{q})$. First, notice that the minimum cannot lay in the segment connecting the vertexes $(1-q_2, 1-q_2)$ and (q_1, q_1). In this case, we have that $\nabla\mathcal{W}(\vec{p}) = \frac{1}{2}(2p, 2p-2)$, hence the gradient is never perpendicular to the line $p_1 = p_2$. We then need to search for the minimum in the sets $\{(1-q_2, t) \text{ s.t. } t \in [1-q_2, q_1]\}$ and $\{(t, q_1) \text{ s.t. } t \in [1-q_2, q_1]\}$.

Theorem 5.3. *Let $\vec{q}$ be a p.c.v. such that $q_2 \leq q_1$, then an optimal ERM for a uniformly distributed population is induced by $(Id, \vec{p})$ where (i) $\vec{p} = (0.25, 0.75)$ if $q_2 \geq 0.75$, (ii) $\vec{p} = (1-q_2, q_1)$ if $3q_2 - 2 - q_1 \leq 0$, and (iii) $\vec{p} = (1-q_2, 1-\frac{q_2}{3})$ otherwise.*

Lastly, we compare the performances of an optimal ERM and the Extended Endpoint Mechanism (EEM), introduced in [5].

Bayesian Analysis of the EEM. The EEM is a truthful mechanism that can locate any two capacitated facilities. The routine of the EEM is as follows. Given $\vec{x} = (x_1, \dots, x_n)$ a vector containing the agents' reports ordered non-decreasingly, we define $A_1 = \left\{x_i \text{ s.t. } |x_i - x_1| \leq \frac{|x_1 - x_n|}{2}\right\}$ and $A_2 = \left\{x_i \text{ s.t. } |x_i - x_n| > \frac{|x_1 - x_n|}{2}\right\}$. Depending on the cardinality of the sets A_i and the capacities of the facilities, the EEM determines the positions of the facilities following one of six possible routines (see Appendix C). To analyze the convergence of the Bayesian approximation ratio of the EEM we need to study all the costs of all the six possible outcomes, weight them by the likelihood of them occurring, and take the limit for n goes to $+\infty$.[3] For instance, let us consider Example 1. In this case, we have that, as the number of agents increases, the EEM will either (1) place the facility with capacity q_2 at $y_1 = x_1$ and the other facility at $y_2 = 2x_{n-(\lfloor q_1(n-1)\rfloor+1)} - x_1$ with probability 0.5 or (2) place the facility with capacity q_2 at $y_2 = x_n$ and the other facility at $y_1 = 2x_{\lfloor q_1(n-1)\rfloor+2} - x_n$ with probability 0.5. Since $\mu = \mathcal{U}[0,1]$, we have that $x_1 \to 0, x_n \to 1, x_{n-\lfloor q_2(n-1)\rfloor+2} \to F_\mu^{[-1]}(1-q_2) = 1-q_2$, and $x_{\lfloor q_2(n-1)\rfloor+1} \to F_\mu^{[-1]}(q_2) = q_2$. In particular, the limit expected Social Cost of the mechanism is then $\frac{1}{2}W_1(\mu, q_2\delta_{2q_2-1} + (1-q_2)\delta_1) + \frac{1}{2}W_1(\mu, (1-q_2)\delta_0 + q_2\delta_{2-2q_2})$. Since μ is symmetric with respect to $\frac{1}{2}$, we have that $W_1(\mu, q_2\delta_{2q_2-1} + (1-q_2)\delta_1) = W_1(\mu, (1-q_2)\delta_0 + q_2\delta_{2-2q_2}) = \frac{(1-q_2)^2}{2} + \frac{q_2(2-3q_2)}{2} = 0.34$, hence the limit Bayesian approximation ratio of the EEM is ~ 2.62. By Theorem 5.3, the optimal ERM is the one induced by $\vec{p} = (0.6, 0.8)$, thus its limit Bayesian approximation ratio is ~ 1.62. The sub-efficiency of the EEM with respect to an optimal ERM is due to the fact that the EEM places both the facilities outside the interval (x_1, x_n), which, in the limit, leads to a loss of efficiency. Moreover, it is worth to notice that the convergence of the expected Social Cost of the EEM depends on the convergence of the first or last order statistic.

[3] Since the EEM cannot be phrased as an ERM, we cannot rely on Theorem 4.2.

6 NUMERICAL EXPERIMENTS

In this section, we complement our theoretical study of the ERMs by running several numerical experiments. Indeed, most of our results pertain to the limit analysis of the mechanism. For this reason, we want to test two aspects of the ERMs. First, we want to compare the Bayesian approximation ratio of the ERMs with the Bayesian approximation ratios of other truthful mechanisms when the number of agents n is small. Since the Ranking Mechanisms are a subset of the ERMs and have been discussed in Section 5.1, we only consider the EEM and, when possible, the IG. Second, we want to evaluate the convergence speed of the Bayesian approximation ratio of $\mathsf{ERM}^{(\pi,\vec{p})}$. More specifically, we want to assess how close the Bayesian approximation ratio of an ERM and the limit detected in Theorem 4.2 are when the number of agents is small. We run our experiments for different distributions μ and percentage capacity vector $\vec{q}$. All the experiments are performed in Matlab 2023a on macOS Monterey system with Apple M1 Pro CPU and 16GB RAM. The code is available on https://anonymous.4open.science/r/Bayesian-CFLP-38D5/.

Experiment setup. Throughout our experiments, we sample the agents' positions from three probability distributions: the uniform distribution $\mathcal{U}[0,1]$, the standard normal distribution $\mathcal{N}(0,1)$, and the exponential distribution $\text{Exp}(1)$. Owing to Corollary 4.2.1, we do not consider other parameter choices since testing the mechanisms over the standard Gaussian distribution is the same as testing over $\mathcal{N}(\mathfrak{m}, \sigma^2)$. To compare the ERMs to other mechanisms, we limit our tests to cases in which $m = 2$, as all the known truthful mechanisms different from the Ranking Mechanisms operate only under this restriction. We consider different percentage capacity vectors $\vec{q} \in (0,1)^2$. Specifically, we consider balanced capacities $\vec{q} = (q, q)$ and unbalanced capacities $\vec{q} = (q_1, q_2), q_1 \neq q_2$. For the case of balanced capacities, we set $q = 0.7, 0.8$ and 0.9. For the case of unbalanced capacities, we consider the slightly unbalanced capacities i.e. $\vec{q} = (0.85, 0.75)$, and extremely unbalanced capacities i.e. $\vec{q} = (0.8, 0.4)$, $(0.85, 0.35)$. As benchmark mechanisms, we consider the EEM [5] and, when possible, the IG [42].

Experiment results – Comparison with the EEM and the IG. We first consider the case of balanced capacities $\vec{q} = (q, q)$, in which we are able to compare the Bayesian approximation ratio of three different mechanisms for the m-CFLP: the EEM, the IG, and the ERMs. Regardless of the distribution, we consider the optimal ERM with respect to the uniform distribution, i.e. $\mathsf{ERM}^{(\pi,\vec{p})}$, with $\pi = Id$ and $\vec{p} = (\max\{0.25, 1-q\}, \min\{0.75, q\})$. Figure 1 (and Table 1 in Appendix D) shows the average and the 95% confidence interval (CI) of Bayesian approximation ratio for $n = 10, 20, 30, 40, 50$. Each average is computed over 500 instances. We observe that, in most cases, the ERM achieves the lowest Bayesian approximation ratio comparing to the other two mechanisms. When $q = 0.7$, the ERM is still better than the EEM but slightly worse than the IG. However, the empirical Bayesian approximation ratio of the ERM and IG converges to the same value as the number of agents increases.

Next, we consider the case of unbalanced capacities, where $q_1 \neq q_2$, specifically $\vec{q} = (0.85, 0.75), (0.8, 0.4), (0.85, 0.35)$. Since the IG requires the two capacities to be identical, we compare only the ERM and the EEM. Amongst the possible ERMs, we select the one optimal with respect to the Uniform distribution, obtained via

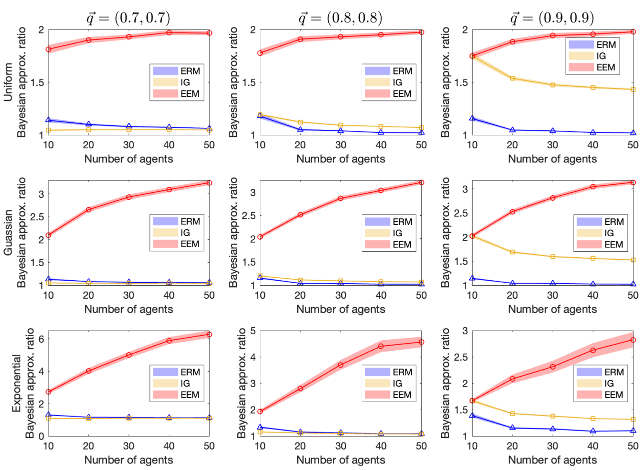

Figure 1: The Bayesian approximation ratio of ERM, IG, and EEM in the balanced case, i.e. $q_1 = q_2$ for $n = 10, 20, \ldots, 50$. Every column contains the results for different vector $\vec{q}$. The first row contains the results for the Uniform distribution, the second row the results for the Gaussian distribution, and the third one the results for the Exponential distribution.

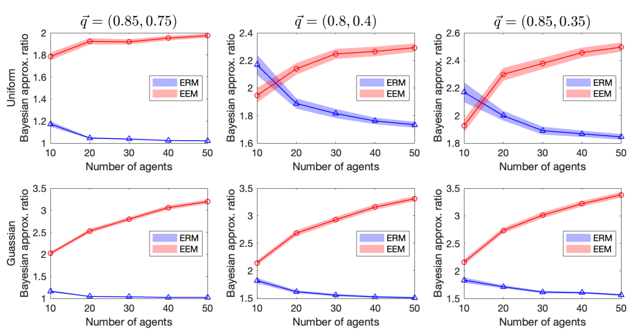

Figure 2: The Bayesian approximation ratio of ERM and EEM in the unbalanced case, i.e. $q_1 \neq q_2$ for $n = 10, 20, \ldots, 50$. Every column contains the results for a different vector $\vec{q}$. The first row contains the results for the Uniform distribution. The second row contains the results for the Gaussian distribution.

Theorem 5.3. Thus the parameters of $\mathsf{ERM}^{(\pi,\vec{p})}$ are (i) $\pi = Id$ for every $\vec{q}$ and (ii) $\vec{p} = (0.25, 0.75)$ for $\vec{q} = (0.85, 0.75)$, $\vec{p} = (0.6, 0.8)$ for $\vec{q} = (0.8, 0.4)$, and $\vec{p} = (0.65, 0.85)$ for $\vec{q} = (0.85, 0.35)$. In this case, we consider only symmetric probability distributions, i.e. the Gaussian and the Uniform distribution. Figure 2 (and Table 2 in Appendix D) shows the average and the 95% CI of Bayesian approximation ratio computed over 500 instances. Whenever $n \geq 20$, the ERM has a much lower approximation ratio. Notice that when $\vec{q} = (0.75, 0.85)$, the ERM is optimal in both cases, and its limit Bayesian approximation ratio is 1. Indeed, the Bayesian approximation ratio of the ERM is almost equal to 1 for every $n \geq 10$ and gets closer as n increases, however the Bayesian approximation ratio of the EEM is always ≥ 1.79 and gets worse as n increases.

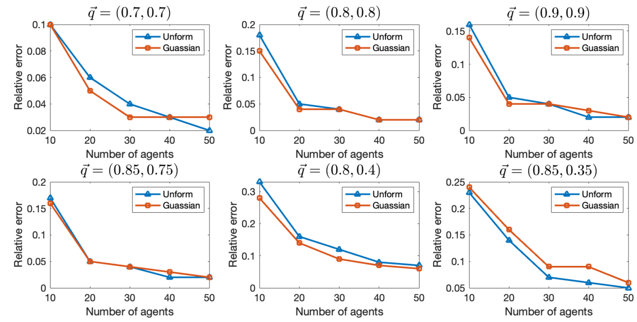

Figure 3: The relative error of ERM for $n = 10, 20, \ldots, 50$ for the Uniform and Gaussian distributions. The first row shows the results for the balanced case while the second row shows the results for unbalanced case.

Experiment results – Convergence speed of the limit Bayesian approximation ratio. Lastly, we test how close the Bayesian approximation ratio of the ERM is to the limit detected in Theorem 4.2. That is, we calculate the relative error as $err_{rel} = \frac{\text{empirical } B_{ar} - \text{limit of } B_{ar}}{\text{limit of } B_{ar}}$. Figure 3 (and Table 3 in Appendix D) show the relative error for the six cases. Each average is computed over 500 instances. We observe that the relative error decreases as the number of agents increases, regardless of the distribution or the p.c.v. $\vec{q}$. Moreover, in all cases but $\vec{q} = (0.8, 0.4), (0.85, 0.35)$, the relative errors of the ERM are less than 0.05 as long as $n \geq 20$, which validates Theorem 4.2 as a tool to predict the Bayesian approximation ratio for small values of n. In other cases, we observe a slower convergence, as the percentage error is slightly larger and its highest value is 0.33.

7 CONCLUSION AND FUTURE WORKS

In this paper, we introduced the Extended Ranking Mechanisms, a generalization of Ranking Mechanisms introduced in [5]. After establishing the conditions under which ERMs remain truthful, we characterized the limit Bayesian approximation ratio of truthful ERM in terms of the probability distribution μ, the p.c.v. $\vec{q}$, and the mechanism's parameters, namely π and $\vec{p}$. We have shown that, given μ and $\vec{q}$, there exists an optimal ERM and characterized it via a minimization problem, which we solved in two relevant frameworks. Lastly, we conducted extensive numerical experiments to validate our findings, from which we inferred that a well-tuned ERM consistently outperforms all other known mechanisms.

For future works, we aim to extend our studies to include other relevant metrics, such as the Maximum Cost or the l_p-costs. Another interesting extension would be to extend the ERMs to handle problems in higher dimensions using the decomposition proposed in [3]. Lastly, we plan to study the connections between the Optimal Transportation problem and other Mechanism Design problems.

ACKNOWLEDGMENTS

Gennaro Auricchio and Jie Zhang are partially supported by a Leverhulme Trust Research Project Grant (2021 – 2024). Jie Zhang is also supported by the EPSRC grant (EP/W014912/1). Mengxiao Zhang is supported by EPSRC Grant (EP/X01357X/1).

REFERENCES

[1] Amir Ahmadi-Javid, Pardis Seyedi, and Siddhartha S Syam. 2017. A survey of healthcare facility location. *Computers & Operations Research* 79 (2017), 223–263.

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

[3] Gennaro Auricchio, Federico Bassetti, Stefano Gualandi, and Marco Veneroni. 2018. Computing Kantorovich-Wasserstein Distances on d-dimensional histograms using $(d+1)$-partite graphs. *Advances in Neural Information Processing Systems* 31 (2018), 2441–2469.

[4] Gennaro Auricchio, Jie Zhang, and Mengxiao Zhang. 2023. Extended Ranking Mechanisms for the m-Capacitated Facility Location Problem in Bayesian Mechanism Design. *arXiv preprint arXiv:2312.16034* (2023).

[5] Haris Aziz, Hau Chan, Barton Lee, Bo Li, and Toby Walsh. 2020. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. MIT Press, New York, 1806–1813.

[6] Haris Aziz, Hau Chan, Barton E Lee, and David C Parkes. 2020. The capacity constrained facility location problem. *Games and Economic Behavior* 124 (2020), 478–490.

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

[8] O Haluk Barda, Joseph Dupuis, and Pierre Lencioni. 1990. Multicriteria location of thermal power plants. *European Journal of Operational Research* 45, 2-3 (1990), 332–346.

[9] Sergey Bobkov and Michel Ledoux. 2019. *One-dimensional empirical measures, order statistics, and Kantorovich transport distances*. Vol. 261. American Mathematical Society, Providence, Rhode Island.

[10] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.

[11] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. 2021. Mechanism Design for Facility Location Problems: A Survey. In *30th International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 4356–4365.

[12] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. ACM, New York, 243–251.

[13] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. *Multi-parameter mechanism design and sequential posted pricing*. ACM, New York. 311–320 pages.

[14] Shuchi Chawla and Balasubramanian Sivan. 2014. Bayesian algorithmic mechanism design. *ACM SIGecom Exchanges* 13, 1 (2014), 5–49.

[15] Richard L Church and Robert S Garfinkel. 1978. Locating an obnoxious facility on a network. *Transportation science* 12, 2 (1978), 107–118.

[16] Laurens De Haan and Elselien Taconis-Haantjes. 1979. On Bahadur's representation of sample quantiles. *Ann. Inst. Statist. Math* 31, Part A (1979), 299–308.

[17] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. 2012. Mechanism design on discrete lines and cycles. In *EC*. ACM, New York, 423–440.

[18] Ludwig Ensthaler and Thomas Giebe. 2014. Bayesian optimal knapsack procurement. *European Journal of Operational Research* 234, 3 (2014), 774–779.

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

[20] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, New York, 269–286.

[21] Michal Feldman and Yoav Wilf. 2013. Strategyproof facility location and the least squares objective. In *EC*. ACM, New York, 873–890.

[22] Alina Filimonov and Reshef Meir. 2021. Strategyproof Facility Location Mechanisms on Discrete Trees. In *AAMAS*. ACM, New York, 510–518.

[23] Aris Filos-Ratsikas, Minming Li, Jie Zhang, and Qiang Zhang. 2015. Facility Location with Double-Peaked Preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 893–899. https://doi.org/10.1609/AAAI.V29I1.9297

[24] Martin Gairing, Burkhard Monien, and Karsten Tiemann. 2005. Selfish routing with incomplete information. In *Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*. ACM, New York, 203–212.

[25] Jason D Hartline et al. 2013. Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science* 8, 3 (2013), 143–263.

[26] Jason D Hartline and Brendan Lucier. 2010. Bayesian algorithmic mechanism design. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, New York, 301–310.

[27] Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, New York, 225–234.

[28] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. Vol. 2. Springer, Berlin.

[29] Leonid V Kantorovich. 2006. On the translocation of masses. *Journal of mathematical sciences* 133, 4 (2006), 1381–1382.

[30] J Kiefer. 1967. On Bahadur's representation of sample quantiles. *The Annals of Mathematical Statistics* 38, 5 (1967), 1323–1342.

[31] Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *EC*. ACM, New York, 315–324.

[32] Pinyan Lu, Yajun Wang, and Yuan Zhou. 2009. Tighter Bounds for Facility Games. In *WINE (Lecture Notes in Computer Science, Vol. 5929)*. Springer, Berlin, 137–148.

[33] Brendan Lucier and Allan Borodin. 2010. Price of anarchy for greedy auctions. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, Society for Industrial and Applied Mathematics, Philadelphia, 537–553.

[34] Lili Mei, Deshi Ye, and Guochuan Zhang. 2016. Mechanism Design for One-Facility Location Game with Obnoxious Effects. In *International Workshop on Frontiers in Algorithmics*. Springer, Berlin, 172–182.

[35] Reshef Meir. 2019. Strategyproof Facility Location for Three Agents on a Circle. In *SAGT (Lecture Notes in Computer Science, Vol. 11801)*. Springer, Berlin, 18–33.

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

[37] Noam Nisan and Amir Ronen. 1999. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*. ACM, New York, 129–140.

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

[39] Filippo Santambrogio. 2015. Optimal transport for applied mathematicians. *Birkäuser, NY* 55, 58-63 (2015), 94.

[40] Pingzhong Tang, Dingli Yu, and Shengyu Zhao. 2020. Characterization of Group-strategyproof Mechanisms for Facility Location in Strictly Convex Space. In *EC*. ACM, New York, 133–157.

[41] Cédric Villani. 2009. *Optimal transport: old and new*. Vol. 338. Springer, Berlin.

[42] Toby Walsh. 2022. Strategy Proof Mechanisms for Facility Location with Capacity Limits. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. International Joint Conferences on Artificial Intelligence Organization, Massachusetts, 527–533.

[43] Qiqi Yan. 2011. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, San Francisco, 710–719.

[44] Emmanouil Zampetakis and Fred Zhang. 2023. Bayesian Strategy-Proof Facility Location via Robust Estimation. In *International Conference on Artificial Intelligence and Statistics*. PMLR, London, 4196–4208.