FLIGHT: Facility Location Integrating Generalized, Holistic Theory of Welfare

Avyukta Manjunatha Vummintala
International Institute of Information Technology
Hyderabad, India
avyukta.v@research.iiit.ac.in

Shivam Gupta
Indian Institute of Technology
Ropar, India
shivam.20csz0004@iitrpr.ac.in

Shweta Jain
Indian Institute of Technology
Ropar, India
shwetajain@iitrpr.ac.in

Sujit Gujar
International Institute of Information Technology
Hyderabad, India
sujit.gujar@iiit.ac.in

ABSTRACT

The Facility Location Problem (FLP) is a well-studied optimization problem with applications in many real-world scenarios. Past literature has explored the solutions from different perspectives to tackle FLPs. These include investigating FLPs under objective functions such as utilitarian, egalitarian, Nash welfare, etc. We propose a unified framework, FLIGHT, to accommodate a broad class of welfare notions. The framework undergoes rigorous theoretical analysis, and we prove some structural properties of the solution to FLP. Additionally, we provide approximation bounds, which (under certain assumptions) provide insight into an interesting fact– as the number of agents arbitrarily increases, the choice of welfare notion is irrelevant. Furthermore, the paper examines a scenario in which the agents are independently and identically distributed (i.i.d.) according to a given probability distribution. In this setting, we derive results concerning the optimal estimator of the welfare and establish an asymptotic result for welfare functions.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory and mechanism design**.

KEYWORDS

Facility Location; Welfare functions

ACM Reference Format:

Avyukta Manjunatha Vummintala, Shivam Gupta, Shweta Jain, and Sujit Gujar. 2025. FLIGHT: Facility Location Integrating Generalized, Holistic Theory of Welfare. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

The most commonly studied *Facility Location Problem* (FLP) considers the problem of placing a facility on a line segment (typically normalized as $[0, 1]$). Here, agents derive certain utilities from this facility. The goal of a planner is to ensure the welfare of the agents who use this facility is maximized. Traditional approaches typically rely on predefined welfare functions such as *utilitarian welfare* [11], which maximizes the total welfare (e.g., travel distance), or *egalitarian welfare* [26], which maximizes the minimum welfare.

While these approaches are suitable in many scenarios, they may not be sufficient to capture the complexity and nuances of real-world applications, especially when the relationship between agents and the facility involves non-linear or application-specific factors. Factors such as varying environmental conditions and resource constraints can introduce non-linearities that complicate the optimization process. In many cases, using simple distance-based models can lead to suboptimal solutions that fail to reflect the true welfare of the system. Moreover, with the rise of *Artificial Intelligence* (AI) and *Machine Learning* (ML) techniques, there is a growing trend toward *learning* welfare functions from data rather than relying on predefined or assumed models. Recent advancements in machine learning have shown that welfare functions can be inferred directly from historical data, allowing for a more flexible and context-aware approach [1, 24, 34]. This shift towards data-driven models underscores the importance of a generalized framework that can accommodate learned welfare functions, making it adaptable to changing environments. Hence, a *generalized framework* is necessary to accommodate various welfare functions that can adapt to these complexities.

One approach to generalizing welfare functions is through the use of *p-mean functions* [3, 14], which provide a continuous spectrum of solutions ranging from utilitarian welfare (when $p = 1$) to egalitarian welfare (when $p = \infty$). The p-mean functions allow more control over the system's balance between efficiency and fairness. p-mean functions, while powerful, may still not fully capture the diversity of welfare considerations present in real-world applications. In this paper, we seek to go beyond p-mean functions by introducing a *generalized framework* that allows for the inclusion of a wide variety of welfare functions, including but not limited to p-mean functions. Our *generalized welfare framework* – FLIGHT, is designed to handle welfare functions that are learned from data, or defined based on specific application needs. FLIGHT framework is flexible enough to incorporate traditional welfare functions, such as utilitarian and egalitarian welfare, and more complex welfare functions that arise in modern applications as long as the welfare function is non-increasing for each agent from its location.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

We begin by establishing several key properties of generalized welfare functions, such as *concavity* and *location invariance*, which are essential for ensuring tractable optimization. Our goal is to study the structural properties of the solution to FLP modelled in FLIGHT. First, we explore more specific properties of these welfare functions. Next, for a concave positive welfare function α and an arbitrary welfare function β, we provide bound on optimal welfare achieved by α with respect to the welfare achieved by β

Our results show that a class of generalized welfare functions can approximate others with a constant approximation ratio, ensuring that our framework remains efficient. Often, a practitioner might be more interested in expected welfare than exact welfare for every instance. Towards this, we investigate probabilistic versions of the facility location problem in which the agents are independently and identically distributed (i.i.d.) according to a given probability distribution. We derive results concerning the optimal estimator of the welfare and provide an asymptotic property of welfare functions. In summary, our contributions are as follows:

1.1 Our Contributions

(1) We propose a unified framework *FLIGHT* that is capable of accommodating classical welfare functions, including utilitarian, egalitarian, and Nash welfare functions (Section 4)
(2) Under the concavity assumption of the utility function, we derive a series of theoretical results concerning the structural properties of generalized welfare functions. (Section 5) Specifically, we prove:
   - **Theorem 1:** Concavity of the welfare function,
   - **Theorem 2:** Location invariance of the welfare function,
   - **Theorem 3:** Behavior under agent shifts, and
   - **Theorem 4:** Maximum shift property.

   We then explore more specialized properties under stronger assumptions, including:
   - **Theorem 5:** Constant approximation bound for concave and positive utility functions,
   - **Theorem 8:** Bounding distance between the peaks based upon the agent location profile.
(3) In Section 6, we extend the analysis to probabilistic versions of the facility location problem. We establish estimation bounds (**Theorem 9** and **Theorem 10**) and derive asymptotic results, including **Theorem 11**, which provides an asymptotic property of welfare functions.

2 RELATED WORK

The facility location problem (FLP) has a long and rich history, with its origins tracing back to 17^{th}-century mathematicians like Pierre de Fermat and Evangelista Torricelli, who studied geometric optimization problems involving the positioning of points to minimize distances to a given set of locations, known as the Fermat-Weber problem [8]. This early work laid the foundation for modern FLP. The field saw significant growth after World War II, spurred by advances in operations research, as facility location became crucial for industrial planning, supply chains, and logistics [15]. During this period, figures such as Harold Kuhn formalized mathematical models that enabled the practical application of FLP to real-world challenges, ranging from public service placement to telecommunications infrastructure [32]. In modern times, the facility location problem has found broad applications in diverse fields such as operations research, computer science, and electronics. With the rise of data-driven decision-making, facility location models are now applied in cloud computing infrastructure, data centers, network design, and even in the placement of sensors in wireless networks [27]. The continued relevance of facility location models underscores their versatility in addressing problems that require optimal resource allocation and spatial planning.

2.1 General Facility Location

The general facility location problem has been widely studied across various fields due to its applications in logistics, urban planning, and operations research. A general overview of the results and variants of FLPs can be found in Chan et al. [5], Farahani and Hekmatfar [10], Melo et al. [21]. Several variants of the FLP have been studied, such as obnoxious facility location [29] and capacitated facility location [33]. Online FLPs are also studied where the agents arrive in an online fashion and a set of facilities is maintained [12, 22].

Additionally, Leitner et al. [19] examines the polytope associated with the asymmetric version of the facility location problem. Fotakis and Tzamos [13] also study facility location with concave welfare functions. However, their focus is on designing algorithms with a constant approximation ratio, whereas our work investigates the structural properties of such a system. Snyder [28] considers a probabilistic view of FLPs. This is relevant as we also perform a probabilistic analysis.

2.2 Facility Location on a Line, Fairness, and Strategyproofness

The facility location problem on a line, where both agents and facilities are confined to a linear domain, has garnered significant attention for its simplicity and traceability. Procaccia and Tennenholtz [25], Procaccia et al. [26] provide approximation guarantees to deterministic and randomized mechanisms that minimize the total cost while maintaining strategyproofness to ensure no agent can manipulate the outcome. These works highlight the need for welfare functions that incorporate fairness, and our framework addresses this requirement.

Recent work on fairness in FLPs has become increasingly relevant as a growing emphasis has been placed on equitable distribution across agents [6, 16, 20]. Moulin [23] introduced the Nash welfare function, establishing its foundational role in welfare economics. Lam et al. [18] further highlights its application in facility location, demonstrating that the Nash welfare function effectively balances fairness and efficiency. This is particularly important for our generalized welfare framework, which aims to extend beyond specific functions like Nash welfare. Furthermore, Chen et al. [6] introduce algorithms for 2-facility location that ensure envy-freeness, reinforcing the importance of fairness in our work. Lam et al. [17] examine the problem of *proportional fairness* in obnoxious facility location, where facilities are undesirable to agents and fairness becomes a key concern. Wang et al. [31] introduce the concept of *positive intra-group externalities* in FLP, focusing on how intra-group dynamics affect utility and strategyproof mechanisms [31].

2.3 Welfare Functions and p-mean Functions

Welfare functions have long been central to decision-making and resource allocation in facility location. Balcan et al. [1] presents a method for learning welfare functions from revealed preferences, which is critical as our generalized framework aims to accommodate complex and dynamically evolving welfare functions. Pardeshi et al. [24] explores the theoretical front of learning welfares or preferences through the context of generalization bounds. In the context of Nash welfare, Caragiannis et al. [4] demonstrates its use in allocation problems, reinforcing the importance of designing flexible welfare functions that balance fairness and efficiency.

Researchers have also explored generalizations of utilitarian and egalitarian welfare through p-mean functions [14], which can be viewed as a parameterized family of welfare functions where varying the parameter p adjusts the balance between fairness and efficiency [7]. For instance, $p = 1$ corresponds to utilitarian welfare, $p = \infty$ corresponds to egalitarian welfare, and intermediate values of p provide trade-offs between these extremes. Our work builds on these concepts by integrating p-mean functions into a broader framework for generalized welfare functions.

Barman and Suzuki [2] and Lam et al. [18] contribute to the growing body of work on Nash welfare, focusing on balancing fairness and efficiency. Our framework expands on these ideas by allowing for general welfare functions that can capture more complex and non-linear utility structures, as noted by Drezner and Scott [9] in facility location. The increasing need for *learned* generalized welfare functions [1, 24, 34] to accommodate engineering applications and other real-world complexities further motivates our research. In the next section, we will introduce the formal problem setup and explain the notations.

3 PRELIMINARIES

3.1 Facility Location Problem Setup

We consider a scenario in which a set of n agents, denoted by $N = \{1, \dots, n\}$, are positioned along the interval[1] $[0, 1]$. Each agent $i \in N$ is located at a specific point $x_i \in [0, 1]$, and the collective set of agent locations is represented by the vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$. Without loss of generality (w.l.o.g.), we assume that the agent positions are ordered such that $x_1 \leq x_2 \leq \dots \leq x_n$.

The social planner's problem is placing a single facility that serves these agents. Let the mechanism of this mapping be $f : [0, 1]^n \rightarrow [0, 1]$, which takes the vector of agent locations $\mathbf{x}$ as input and returns a location $y \in [0, 1]$ for the facility. For a facility at location y, agent i need to travel $\mid y - x_i \mid$. Thus, $\mid y - x_i \mid$ indicates the cost to it or in some contexts, $1 - \mid y - x_i \mid$ indicates the utility to agent i. The most prominently studied welfare functions are computed as follows.

Definition 1 (Utilitarian Welfare). *For the agents located at* $\mathbf{x}$ *and the facility located at* y*, the* Utilitarian Welfare *is*

$$W_{Utilitarian}(y, \mathbf{x}) = \sum_i (1 - \mid y - x_i \mid)$$

Definition 2 (Egalitarian Welfare). *For the agents located at* $\mathbf{x}$ *and the facility located at* y*, the* Egalitarian Welfare *is*

$$W_{Egalitarian}(y, \mathbf{x}) = \min_i (1 - \mid y - x_i \mid)$$

Definition 3 (Nash Welfare). *For the agents located at* $\mathbf{x}$ *and the facility located at* y*, the* Nash Welfare *is*

$$W_{Nash}(y, \mathbf{x}) = \prod_i (1 - \mid y - x_i \mid)$$

Typically, the social planer aims to place the facility at a location y that maximizes $W_{\text{Utilitarian}}$ or W_{Nash} or $W_{\text{Egalitarian}}$. There are closed-form solutions for Utilitarianism and Egalitarianism.

3.2 Key Important Mechanisms

In facility location problems (FLP), different welfare optimization criteria lead to distinct placement strategies for the facility.

The solution that maximizes *utilitarian welfare*—defined as the total sum of utilities—is the median of the agent locations. Formally, this position is given by $x_{\lfloor n/2 \rfloor}$, where n represents the total number of agents. This placement has the additional advantage of being *strategyproof*, meaning agents cannot benefit from misreporting their locations.

In contrast, the solution that maximizes *egalitarian welfare* (focused on maximizing the minimum utility for any agent)—is the midpoint between the extreme agents. This solution can be expressed as $\frac{x_1 + x_n}{2}$, where x_1 and x_n represent the positions of the agents at the two extremes.

Finally, the solution that maximizes *Nash welfare*—a balance between utilitarian and egalitarian objectives—is more complex. The Nash welfare function is the product of individual utilities, and finding its maximization in FLP is known to be both difficult to compute and interpret in practice [23].

While these solutions maximize different welfare objectives, many interesting properties emerge from their comparative analysis. However, these properties have traditionally been studied separately for each welfare function. This paper proposes a unifying framework that allows for the study of these properties in a more general, abstract manner.

As stated previously (Section 1), there are scenarios where one must go beyond the three classical welfare functions. Rather than developing new solutions for each emerging welfare criterion, a more holistic approach can be adopted. Specifically, we want to study facility location as an abstract problem, agnostic to the specific welfare function, by focusing on common properties shared by many of these functions. One such approach involves the use of p-mean functions, which we explain in the next section.

3.3 p-mean Welfare Functions

Consider the facility location problem where the L_p-norm is used as the distance metric between agents and the facility. The solution y_{Pmean} to the facility location problem under the L_p-norm is defined as the facility location that minimizes the p-mean distance to all agents, given by:

$$y_{Pmean} = \arg \min_{y \in [0,1]} \left(\sum_{i \in N} |y - x_i|^p \right)^{1/p} \tag{1}$$

[1] Note that the [0,1] domain can be extended and translated to be any closed interval.

Since the p-th root is a monotonically increasing function, we can simplify the optimization problem to:

$$y_{Pmean} = \arg\min_{y\in[0,1]} \sum_{i\in N} |y - x_i|^p \quad (2)$$

Next section proposes a more general framework, *FLIGHT* – **Facility Location Integrating Generalized, Holistic Theory of Welfare**. FLIGHT generalizes the concept of welfare functions and provides a unified approach to solving facility location problems.

4 A UNIFIED PERSPECTIVE

We propose FLIGHT and demonstrate how all well-studied welfare functions, including p-mean functions, can be incorporated into it. We show that the Nash Welfare function can also be integrated within the FLIGHT framework, thereby highlighting the versatility and generality of our approach in encompassing a wide range of welfare functions.

4.1 FLIGHT Framework

Utility for an agent at location x_i when the facility is located at y is a function of $y - x_i$. Let the utility function for each agent be $\alpha : \mathbb{R} \to \mathbb{R}$. Specifically, α takes the distance from the facility as input and returns the corresponding utility for the agent as output. The function α encapsulates how the agent's utility diminishes with increasing distance from the facility[2].

Next, we define the **total welfare** $W_\alpha(y, \mathbf{x})$ as the aggregate of individual utilities across all agents. Formally, it is expressed as:

$$W_\alpha(y, \mathbf{x}) = \sum_{i\in N} \alpha(y - x_i)$$

where $y \in [0, 1]$ represents the location of the facility, and $\mathbf{x} = (x_1, \ldots, x_n)$ denotes the vector of agent locations.

Given this total welfare function, the social planner's goal is to determine a location that maximizes global welfare. We denote it as $P_\alpha(\mathbf{x})$. Formally, this can be expressed as:

$$P_\alpha(\mathbf{x}) = \arg\max_{y\in[0,1]} W_\alpha(y, \mathbf{x})$$

For Utilitarian welfare, as stated in Section 3.2, $P_\alpha(x) = x_{\frac{n}{2}}$ and for Egalitarian welfare, $P_\alpha(x) = \frac{x_1+x_n}{2}$. In the next section, we show that p-mean welfare functions are special cases of our framework.

4.2 Incorporating p-mean Welfare Functions into Our Framework

In this section, we demonstrate that *p-Mean utility functions* are fully accommodated by our framework. Since *utilitarian welfare* and *egalitarian welfare* are special cases of p-mean utility functions, which naturally fit within our general framework.

4.2.1 Generalizing to p-Mean Utility Functions. The idea is to align optimization from Equation 2 with FLIGHT, we can express it as a maximization problem as:

$$y_{Pmean} = \arg\max_{y\in[0,1]} -\sum_{i\in N} |y - x_i|^p \quad (3)$$

[2]Note: α could be an asymmetric function as well, meaning it lacks symmetry about the y-axis.

To incorporate p-mean utility functions into our framework, we define a utility function $\alpha : \mathbb{R} \to \mathbb{R}$, where $\alpha(x) = -|x|^p$. Using this definition, the total welfare function $W_\alpha(y, \mathbf{x})$ becomes:

$$W_\alpha(y, \mathbf{x}) = \sum_{i\in N} \alpha(y - x_i) = \sum_{i\in N} -|y - x_i|^p \quad (4)$$

4.2.2 Special Cases: Utilitarian and Egalitarian Welfare. Both *utilitarian welfare* and *egalitarian welfare* are special cases of the p-mean utility functions, fitting naturally within our framework. The utilitarian welfare function corresponds to the case where $p = 1$. Similarly, The egalitarian welfare function corresponds to the limiting case as $p \to \infty$.

4.3 Nash Welfare

In this section, we demonstrate that the *Nash welfare function* [18] is also fully compatible with our framework. To formalize this, let y_{Nash} denote the facility location that maximizes the Nash welfare. We express this as:

$$y_{Nash} = \arg\max_{y\in[0,1]} \prod_{i\in N} (1 - |y - x_i|) \quad (5)$$

We can simplify this by applying the logarithmic transformation. Since the logarithmic function is monotonic, it preserves the location of the maximum. Therefore, we have:

$$y_{Nash} = \arg\max_{y\in[0,1]} \log\left(\prod_{i\in N} (1 - |y - x_i|)\right) \quad (6)$$

Or equivalently:

$$y_{Nash} = \arg\max_{y\in[0,1]} \sum_{i\in N} \log(1 - |y - x_i|) \quad (7)$$

Thus, by defining the utility function as $\alpha(x) = \log(1 - |x|)$, the Nash welfare problem is equivalent to maximizing the total welfare function $W_\alpha(y, \mathbf{x})$.

It is worth noting that our framework naturally accommodates asymmetric welfare functions, an area that has been explored in a limited number of facility location studies [9, 19]. While the study of asymmetry in FLPs remains relatively sparse, our framework offers a convenient approach to incorporating such welfare functions.

Having established that the p-mean functions—along with the utilitarian and egalitarian welfare functions—and the Nash welfare function can be effectively incorporated into our *FLIGHT* framework, we now turn our attention to studying the properties of this generalized welfare formulation. In the following section, we examine key structural properties of the generalized welfare function, with particular emphasis on concavity, which is a natural assumption in many practical contexts since utility typically decreases with increasing distance.

We choose to use *utility* over *cost* in our exposition, primarily for conceptual clarity. Note that our framework, FLIGHT, is inherently flexible and capable of unifying both cost and utility perspectives under the broader notion of *agent single-peaked preferences.* Within this formulation, the α-welfare function can be interpreted as an aggregation of these preferences, ensuring a cohesive and generalized approach to welfare optimization.

We proceed by proving several theorems related to these properties, thereby further elucidating the theoretical foundation of the FLIGHT framework.

5 α-WELFARE: PROPERTIES AND COMPUTATION

In the previous section, we demonstrated that a wide variety of existing welfare notions can be incorporated into our framework. Here, we delve into the general structural properties of α-Welfare functions, focusing on how these properties relate to the computation and approximation of various welfare functions within the framework. Notably, many of these properties echo results found in the literature, thus highlighting the unifying power of our framework. Moreover, several proofs become simplified when viewed through the lens of the generalized framework. We have provided proof sketches wherever possible. The complete proofs are present in Appendix B of the full version of the paper [30].

5.1 Assumptions

We begin by assuming that the utility function $\alpha(x)$ is concave with respect to x, which captures the phenomenon of diminishing returns as the distance between the facility and an agent increases. This assumption of concavity serves as the foundation for the theorems presented in the subsequent sections. Additionally, we assume that $\alpha(x)$ attains its maximum at $x = 0$, reflecting the highest utility when the agent is located at the facility.

The following properties arise naturally from a fundamental assumption of concave utility functions. Note that these do not need to be imposed as design choices.

5.2 Properties of α-Welfare

THEOREM 1. *The total welfare function $W_\alpha(y, \mathbf{x})$, is concave in y.*

Proof. The total welfare is defined as the sum of individual utility functions:

$$W_\alpha(y, \mathbf{x}) = \sum_{i \in N} \alpha(y - x_i)$$

Since $\alpha(x)$ is concave, and the sum of finitely many concave functions is also concave, it follows that $W_\alpha(y, \mathbf{x})$ is concave in y. □

The significance of Theorem 1 lies in the fact that the concavity of the total welfare function $W_\alpha(y, \mathbf{x})$ implies it is *single-peaked* with respect to y, a property that is analogous to the behavior observed in Nash welfare functions [18]. Furthermore, the structural properties established in Theorems 2, 3, and 4 exhibit similar characteristics to those studied in [18], reinforcing the parallels between our framework and Nash welfare-based approaches in facility location. . Single-peakedness is crucial in optimization. It allows the use of efficient convex or concave optimization algorithms to locate the maximum welfare point, facilitating computational approaches to solving the facility location problem. Furthermore, the concavity guarantees that local maxima are also global maxima, simplifying the analysis and solution of the problem. Our next theorem proves the location invariance property of FLIGHT.

THEOREM 2. *Let $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ and $\mathbf{x}' = (x_1 + c, x_2 + c, \ldots, x_n + c)$ be two location profiles, where $c \in \mathbb{R}$ represents a constant shift. Then, the following holds:*

$$W_\alpha(y, \mathbf{x}') = W_\alpha(y - c, \mathbf{x}),$$

and consequently,

$$P_\alpha(\mathbf{x}') = P_\alpha(\mathbf{x}) + c.$$

Proof. The key idea behind this result is that shifting all the agents' locations by a constant c results in a corresponding shift in the facility location by c, without affecting the total welfare function.

To see this, consider the welfare function $W_\alpha(y, \mathbf{x}) = \sum_{i \in N} \alpha(y - x_i)$, where $\alpha(z)$ is the utility function. When the agent profile $\mathbf{x}$ is shifted by a constant c, the welfare function for the shifted profile becomes:

$$W_\alpha(y, \mathbf{x}') = \sum_{i \in N} \alpha((y - c) - x_i).$$

This is equivalent to the original welfare function with the facility location adjusted by c, i.e., $W_\alpha(y, \mathbf{x}') = W_\alpha(y - c, \mathbf{x})$.

Thus, the optimal facility location for the shifted profile, $P_\alpha(\mathbf{x}')$, is simply the original location shifted by c, i.e., $P_\alpha(\mathbf{x}') = P_\alpha(\mathbf{x}) + c$. This completes the proof. □

The next theorem shows how the movement of a single agent x_i by a constant affects the peak P_α with utility function α.

THEOREM 3. *Let $\mathbf{x} = (x_1, \ldots, x_n)$ be the agent location profile. If agent x_i is shifted left by a constant $c \in (0, x_i]$, resulting in a new profile $\mathbf{x}' = (x_1, \ldots, x_i - c, \ldots, x_n)$, then:*

$$P_\alpha(\mathbf{x}') \leq P_\alpha(\mathbf{x})$$

Proof. Since $W_\alpha(y, \mathbf{x}')$ is concave and single-peaked, to prove that $P_\alpha(\mathbf{x}') \leq P_\alpha(\mathbf{x})$, it suffices to show that $W_\alpha(y, \mathbf{x}')$ is strictly decreasing in the interval $(P_\alpha(\mathbf{x}), 1)$. This is because $P_\alpha(\mathbf{x})$ is the maximum of $W_\alpha(y, \mathbf{x})$, and proving the function decreases beyond this point implies that the new maximum $P_\alpha(\mathbf{x}')$ must occur at a lower value.

We begin by expressing the total welfare for the perturbed location profile $\mathbf{x}'$:

$$W_\alpha(y, \mathbf{x}') = \alpha(y - x_1) + \cdots + \alpha(y - (x_i - c)) + \cdots + \alpha(y - x_n)$$

This can be rewritten as:

$$W_\alpha(y, \mathbf{x}') = W_\alpha(y, \mathbf{x}) + \alpha(y - (x_i - c)) - \alpha(y - x_i)$$

Next, we differentiate this expression with respect to y:

$$\frac{dW_\alpha(y, \mathbf{x}')}{dy} = \frac{dW_\alpha(y, \mathbf{x})}{dy} + \frac{d\alpha}{dy}\bigg|_{y-(x_i-c)} - \frac{d\alpha}{dy}\bigg|_{y-x_i}$$

Now, consider the terms on the right-hand side:

1. Since $W_\alpha(y, \mathbf{x})$ is concave and $P_\alpha(\mathbf{x})$ is its maximum, we know that $\frac{dW_\alpha(y,\mathbf{x})}{dy} < 0$ for $y \in (P_\alpha(\mathbf{x}), 1)$.

2. Additionally, because $\alpha(x)$ is concave, we have $\frac{d\alpha}{dy}\big|_{y-(x_i-c)} \leq \frac{d\alpha}{dy}\big|_{y-x_i}$. This follows from the fact that the derivative of a concave function decreases as the input increases.

3. As a result, $\frac{d\alpha}{dy}\big|_{y-(x_i-c)} - \frac{d\alpha}{dy}\big|_{y-x_i} \leq 0$.

Thus, the total derivative $\frac{dW_\alpha(y,\mathbf{x}')}{dy}$ remains negative in the interval $(P_\alpha(\mathbf{x}), 1)$. Since the function is decreasing beyond $P_\alpha(\mathbf{x})$, it follows that: $P_\alpha(\mathbf{x}') \leq P_\alpha(\mathbf{x})$ □

It is important to observe that when an agent moves to the right, i.e., $x_i \longrightarrow x_i + c$ with $c \geq 0$, we can prove by symmetry arguments that $P_\alpha(\mathbf{x}) \leq P_\alpha(\mathbf{x}')$, where $\mathbf{x}'$ represents the updated location profile after the shift.

This result follows by considering $\mathbf{x}'$ as the initial location profile and $\mathbf{x}$ as the deviated profile. Then, by applying Theorem 3, which states that the optimal facility location shifts in the direction of the agent's movement, we conclude that $P_\alpha(\mathbf{x}) \leq P_\alpha(\mathbf{x}')$.

The importance of this theorem lies in the fact that it establishes a directional monotonicity property. Specifically, assuming the positions of all other agents remain fixed, if a single agent shifts to the right, the peak P_α will not move to the left, and similarly, if the agent shifts to the left, the peak will not move to the right.

We now address a broader question: how does the optimal facility location P_α change when all agents move from one location profile $\mathbf{x}$ to a new profile $\mathbf{x}'$? Specifically, we aim to understand the relationship between the shifts in individual agent positions and the resulting change in the welfare-maximizing facility location.

Theorem 4. *For any two agent location profiles $\mathbf{x} = (x_1, x_2 \ldots, x_n)$ and $\mathbf{x}' = (x'_1, x'_2, \ldots, x'_n)$, the following inequality holds:*

$$|P_\alpha(\mathbf{x}) - P_\alpha(\mathbf{x}')| \leq \max_{i\in[n]} |x_i - x'_i|$$

Proof.[3] Define two new location profiles, $\mathbf{x}_{+c}$ and $\mathbf{x}_{-c}$, as follows:

$$\mathbf{x}_{+c} \triangleq (x_1 + c, x_2 + c, \ldots, x_n + c)$$

$$\mathbf{x}_{-c} \triangleq (x_1 - c, x_2 - c, \ldots, x_n - c)$$

By Theorem 3, we know:

$$P_\alpha(\mathbf{x}_{-c}) \leq P_\alpha(\mathbf{x}') \leq P_\alpha(\mathbf{x}_{+c})$$

Additionally, by Theorem 2, we have:

$$P_\alpha(\mathbf{x}) - c \leq P_\alpha(\mathbf{x}') \leq P_\alpha(\mathbf{x}) + c$$

This implies:

$$-c \leq P_\alpha(\mathbf{x}') - P_\alpha(\mathbf{x}) \leq +c$$

Therefore, we conclude:

$$|P_\alpha(\mathbf{x}') - P_\alpha(\mathbf{x})| \leq c = \max_{i\in N} |x'_i - x_i|$$

□

5.3 Approximation Bounds: Comparison to Other Welfare Metrics

We now turn to the question of how well different welfare functions approximate one another under this framework.

Theorem 5. *Let α be a utility function with an upper Lipschitz constant λ^u_α and a lower Lipschitz constant λ^d_α. Additionally, assume that $\alpha(x) > 0$ for all $x \in [-1, 1]$, even though agent locations are restricted to $[0, 1]$. Define:*

$$D_\alpha = \min\{\alpha(0) + (n-1)\alpha(-1), \alpha(0) + (n-1)\alpha(1)\}.$$

Then, for all $y \in [0, 1]$, the following inequality holds:

$$e^{\frac{\lambda^d_\alpha(P_\alpha(\mathbf{x})-y)}{\alpha(0)}} \leq \frac{W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})}{W_\alpha(y,\mathbf{x})} \leq e^{\frac{n\lambda^u_\alpha(P_\alpha(\mathbf{x})-y)}{D_\alpha}}.$$

Proof sketch. To establish the desired bounds, we leverage the Lipschitz properties of the logarithmic function $\ln(x)$, as well as the Lipschitz continuity of $\alpha(x)$, to derive a bound on $\ln\left(\frac{W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})}{W_\alpha(y,\mathbf{x})}\right)$.

The total welfare function $W_\alpha(y, \mathbf{x})$ is the sum of individual utilities based on the location of the facility at y. By applying the Lipschitz properties of $\alpha(x)$, we can bound the ratio of welfare functions by exponentiating the bound on their logarithmic difference. Formally, we write:

$$\ln\left(\frac{W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})}{W_\alpha(y,\mathbf{x})}\right) = \ln(W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})) - \ln(W_\alpha(y,\mathbf{x})).$$

Using the Lipschitz property of $\ln(x)$ and the fact that $\alpha(x)$ is concave with its maximum at $x = 0$, we can apply bounds on this difference. Specifically, since $\alpha(x)$ is Lipschitz continuous with upper and lower bounds given by λ^u_α and λ^d_α, we obtain the bounds for the ratio of the welfare functions.

Finally, applying the exponent to both sides of the inequality provides the result, where the upper and lower bounds depend on the constants λ^u_α, λ^d_α, and the behavior of $\alpha(x)$ at the extremes of its domain. □

Given that the exponent is a rational function with polynomials of the same degree in both the numerator and the denominator, we can conclude (as we prove in Lemma 6) that as $n \longrightarrow \infty$, the exponent is asymptotically bounded by a constant. Furthermore, we can assert that the ratio remains sub-exponential even for small values of n.

Lemma 6. *As the number of agents n increases, the upper bound on the approximation ratio converges to a constant, i.e., at the most $e^{\frac{\lambda^u_\alpha}{\min\{\alpha(-1),\alpha(1)\}}}$.*

Proof. From the previous theorem, the upper bound on the approximation ratio is given by:

$$\frac{W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})}{W_\alpha(y,\mathbf{x})} \leq e^{\frac{n\lambda^u_\alpha(P_\alpha(\mathbf{x})-y)}{D_\alpha}}.$$

As $n \longrightarrow \infty$, the denominator behaves as[4]:

$$D_\alpha \sim (n-1)\min\{\alpha(-1), \alpha(1)\}.$$

Thus, the approximation ratio becomes:

$$\lim_{n\to\infty} e^{\frac{n\lambda^u_\alpha(P_\alpha(\mathbf{x})-y)}{(n-1)\min\{\alpha(-1),\alpha(1)\}}} = e^{\frac{\lambda^u_\alpha(P_\alpha(\mathbf{x})-y)}{\min\{\alpha(-1),\alpha(1)\}}}.$$

However, since $P_\alpha(\mathbf{x}) - y \leq 1$, we have:

$$\lim_{n\to\infty} e^{\frac{n\lambda^u_\alpha(P_\alpha(\mathbf{x})-y)}{(n-1)\min\{\alpha(-1),\alpha(1)\}}} \leq e^{\frac{\lambda^u_\alpha}{\min\{\alpha(-1),\alpha(1)\}}}.$$

Lemma 7. *Let α, β be two utility functions with corresponding welfares $W_\beta(y, \mathbf{x}), W_\alpha(y, \mathbf{x})$ and maximizers $P_\beta(\mathbf{x}), P_\alpha(\mathbf{x})$. Then,* ***Theorem 5*** *yields the approximation ratio between utility functions α and β:*

$$\frac{W_\alpha(P_\alpha(\mathbf{x}),\mathbf{x})}{W_\alpha(P_\beta(\mathbf{x}),\mathbf{x})} \leq e^{\frac{n\lambda_\alpha(P_\alpha(\mathbf{x})-P_\beta(\mathbf{x}))}{D_\alpha}} \leq e^{\frac{n\lambda_\alpha}{D_\alpha}}$$

[3] Note that (except for the FLIGHT setting) this proof is almost exactly equivalent to the one in Lam et al. [18].

[4] Note that $\alpha(x) > 0 \ \forall x \in [-1, 1]$.

By Lemma 6, as $n \to \infty$, this approximation ratio becomes constant.

Proof: The proof follows directly from Lemma 6. The final part of the inequality is true because $P_\alpha(\mathbf{x}) - P_\beta(\mathbf{x}) \leq 1$. □

5.4 Bounding the Distance Between Peaks

The next theorem provides a bound on the distance between the peaks P_α and *median*, based on the configuration of agent locations.

Theorem 8. *Let med denote the median of the location profile* $\mathbf{x} = (x_1, x_2, \ldots, x_n)$. *If n is even, define the median as:*

$$med = \frac{x_{\frac{n}{2}} + x_{\frac{n}{2}+1}}{2}.$$

Then, the following inequality holds:

$$|med - P_\alpha(\mathbf{x})| \leq \frac{1}{2} \max_{i=1}^{\lfloor n/2 \rfloor} |d_i^+ - d_i^-|,$$

where $d_i^+ = |med - x_{\lfloor n/2 \rfloor + i}|$ *and* $d_i^- = |med - x_{\lceil n/2 \rceil - i}|$, *and* α *is a symmetric utility function.*

Proof sketch. The central idea behind this proof is recognizing that if the location profile $\mathbf{x}$ were perfectly symmetric around the median *med*, then the peak of the welfare function, $P_\alpha(\mathbf{x})$, would coincide with the median due to the symmetry of α.

Thus, the deviation of $P_\alpha(\mathbf{x})$ from the median arises solely due to the asymmetry in the distribution of agents around the median. The key question becomes: how far must the agents be shifted to transform the location profile $\mathbf{x}$ into a symmetric configuration?

To quantify this, we compare the distances of the agents on either side of the median. For each agent positioned at $x_{\lfloor n/2 \rfloor + i}$ (to the right of the median), we consider the distance d_i^+ from the median, and similarly, for each agent at $x_{\lceil n/2 \rceil - i}$ (to the left of the median), we consider the distance d_i^-.

The difference $|d_i^+ - d_i^-|$ measures how far these corresponding agents deviate from a symmetric configuration. The maximum of these deviations for all pairs of agents gives a measure of the total asymmetry in the distribution of the agents around the median.

Since the location of $P_\alpha(\mathbf{x})$ is influenced by the overall symmetry of the location profile, the deviation of $P_\alpha(\mathbf{x})$ from the median is bounded by the total asymmetry, which is expressed as:

$$|med - P_\alpha(\mathbf{x})| \leq \frac{1}{2} \max_{i=1}^{\lfloor n/2 \rfloor} |d_i^+ - d_i^-|.$$

This completes the sketch of the proof. □

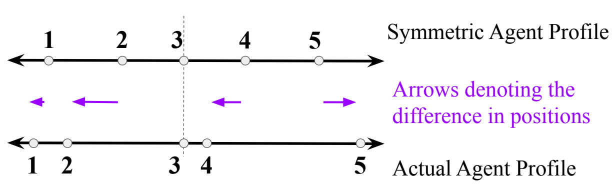

Figure 1: Illustration of Theorem 8: Bounding the Deviation of the Welfare Peak from the Median. The longest arrow, representing the maximum deviation from the symmetric agent profile, provides the required upper bound.

In the following section, we will conduct a probabilistic analysis of our FLIGHT framework. Specifically, we will consider scenarios in which the agent locations are drawn from a probability distribution. By incorporating this probabilistic perspective, we aim to provide a more comprehensive understanding of how the distribution of agents influences the welfare outcomes within our framework.

6 PROBABILISTIC ANALYSIS OF α-WELFARE

In this section, we extend our analysis by introducing a probabilistic framework in which agent locations are treated as random variables. Specifically, we assume that the agents' preferred locations x_i are independently and identically distributed (i.i.d.) samples from a probability distribution $\mathcal{P}$, i.e., $x_i \sim \mathcal{P}$. This formulation allows us to examine the behavior of welfare functions when agent positions are drawn from a probabilistic distribution, which is particularly useful in real-world scenarios where exact agent locations may be uncertain.

6.1 Expected Welfare Function

We define[5] the expected total welfare, $\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x})$, as the expected welfare at location y when agents are sampled from the probability distribution $\mathcal{P}$. This expected welfare provides a probabilistic generalization of our previously defined deterministic welfare functions.

Definition 4 (Empirical Welfare). *The empirical welfare function* $W_\alpha(y, \mathbf{x})$ *is defined as the sum of individual utilities for agents located at* $\mathbf{x} = (x_1, x_2, \ldots, x_n)$, *where the locations* x_i *are sampled from the probability distribution* $\mathcal{P}$. *Formally:*

$$W_\alpha(y, \mathbf{x}) = \sum_{i=1}^{n} \alpha(y - x_i).$$

Definition 5 (Expected Welfare). *Let* $\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x})$ *represent the expected welfare for agents sampled i.i.d. from the distribution* $\mathcal{P}$. *Formally, this is given by:*

$$\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}^n}\left[W_\alpha(y, \mathbf{x})\right],$$

where α *is the individual utility function,* $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ *represents the agent locations, and n is the number of agents.*

Theorem 9. *The expected welfare function* $\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x})$ *is given by:*

$$\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x}) = n \times [\alpha \circledast \mathcal{P}](y),$$

where $[\alpha \circledast \mathcal{P}](y)$ *represents the convolution of the utility function* α *and the probability distribution* $\mathcal{P}$, *evaluated at location* y.

Proof sketch. The result follows from the linearity of expectation and the fact that the sum of n independent random variables sampled from $\mathcal{P}$ is equivalent to scaling the expected utility by n. Specifically, since each agent's utility is independently and identically distributed (i.i.d.) from $\mathcal{P}$, the expected welfare for n agents can be written as the expected welfare for one agent multiplied by n. This yields the convolution result. □

In many practical applications, especially in large-scale systems, we are often not given the exact locations of all agents but rather

[5] In scenarios where agent locations are probabilistic, the total welfare function becomes a random variable. Accordingly, we extend our prior definitions in Section 4 to accommodate this stochastic setting. Such analysis provides insights into ex-ante performance.

a probability distribution $\mathcal{P}$ that describes their likely positions. When dealing with a large number of agents, directly computing the total welfare based on each individual's location can become computationally expensive. This challenge necessitates the development of more efficient techniques for estimating the welfare, particularly when a probability distribution is available.

The following two theorems provide insights, allowing us to circumvent the need to compute the welfare function through explicit agent positions. Instead, these results show that we can rely on the probability distribution $\mathcal{P}$ to derive robust estimates of the welfare function.

First (Theorem 10), we demonstrate that given only the probability distribution $\mathcal{P}$, the best estimator of the welfare function is the expected welfare. This means that the expected welfare serves as an unbiased estimator with the minimum variance, ensuring the most accurate estimate of the empirical welfare function. This result is highly valuable because it simplifies the computation by focusing on the mean of the distribution rather than requiring the sum over all agents.

Second (Theorem 11), we establish that as the number of agents increases, the empirical welfare converges in probability to the expected welfare. This result, rooted in the weak law of large numbers, guarantees that as the number of agents becomes arbitrarily large, the difference between the empirical and expected welfare diminishes. Hence, for engineering applications where we often deal with large populations, using the expected welfare is not only computationally simpler but also practically equivalent to calculating the empirical welfare.

Together, these results underscore the utility of relying on the expected welfare in large-scale systems. The expected welfare is easier to compute and analyze, and as the number of agents increases, it becomes a reliable proxy for the actual welfare function, making it highly suitable for real-world engineering applications.

6.2 Optimality in Function Space

We next examine how the expected welfare function relates to minimizing the distance between welfare functions in a suitable function space.

Definition 6. *Define the $\mathcal{F}$-distance between two functions f_1 and f_2 as:*

$$||f_1 - f_2||_{\mathcal{F}} \triangleq \int_{\mathbb{R}} (f_1(x) - f_2(x))^2 \, dx$$

Theorem 10. *The function $\mathbb{W}_\alpha^{\mathcal{P}}(y, \mathbf{x})$ minimizes the expected $\mathcal{F}$-distance to the empirical welfare function $W_\alpha(y, \mathbf{x})$, i.e., it is the best approximation of the empirical welfare in the $\mathcal{F}$-distance sense.*

Proof sketch. By definition of expected welfare, $W_\alpha(y, \mathcal{P}, n)$ is the mean of the empirical welfare function $W_\alpha(y, \mathbf{x})$, where $\mathbf{x}$ is the random vector of agent locations. Given that the $\mathcal{F}$-distance is a measure of the difference between two functions, the expected value of the empirical welfare is the function that minimizes this difference. The proof mirrors the discrete case. (The Minimum Variance Unbiased Estimator of a random variable is simply its mean.) □

6.3 Asymptotic result on Welfares

The weak law of large numbers is a classical result in probability theory. Here, we derive a variation of the law applied to our welfare framework.

Theorem 11 (Asymptotic result for Welfare Functions). *Let $W_\alpha(y, \mathbf{x})$ be the empirical welfare function for a sample of n agents and $X \sim \mathcal{P}$. Then, as $n \to \infty$, the empirical welfare converges in probability to the expected welfare for a single agent:*

$$\frac{W_\alpha(y, \mathbf{x})}{n} \xrightarrow{p} \mathbb{W}_\alpha^{\mathcal{P}}(y, X)$$

Proof sketch. This is a direct consequence of the weak law of large numbers. As the number of agents increases, the average empirical welfare converges to the expected welfare. The factor of n normalizes the total welfare, ensuring convergence to the expected value for one agent sampled from $\mathcal{P}$. □

7 CONCLUSION

In this paper, we introduced a flexible and unified framework for facility location problems using α-welfare functions, demonstrating that various well-known welfare models, such as *utilitarian*, *egalitarian*, and *Nash welfare*, are special cases within this framework. We established key structural properties of these functions, including *concavity*, *location invariance*, and *monotonicity*, which simplify optimization for facility placement. Additionally, by incorporating a probabilistic perspective, where agent locations are modeled as independent samples from a distribution, we derived the expected welfare function and analyzed how welfare functions approximate each other under uncertainty. Furthermore, we provided *approximation bounds* between different welfare functions and introduced the $\mathcal{F}$-distance as a metric for evaluating discrepancies between empirical and expected welfare functions. This comprehensive framework supports robust decision-making in facility location problems and offers significant potential for future research.

ADDITIONAL RESOURCES

The full version of this paper, including complete proofs and additional discussions, is available in the full version of the paper [30].

REFERENCES

[1] Maria-Florina Balcan, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V. Vazirani. 2014. Learning Economic Parameters from Revealed Preferences. In *Web and Internet Economics*, Tie-Yan Liu, Qi Qi, and Yinyu Ye (Eds.). Springer International Publishing, Cham, 338–353.

[2] Siddharth Barman and Mashbat Suzuki. 2024. Compatibility of Fairness and Nash Welfare under Subadditive Valuations. arXiv:2407.12461 [cs.GT] https://arxiv.org/abs/2407.12461

[3] Nicolas Bourbaki. 2013. *Topological Vector Spaces: Chapters 1–5*. Springer Science & Business Media, Berlin, Heidelberg.

[4] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.* 7, 3, Article 12 (Sept. 2019), 32 pages. https://doi.org/10.1145/3355902

[5] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. 2021. Mechanism Design for Facility Location Problems: A Survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, Vienna, Austria, 4356–4365. https://doi.org/10.24963/ijcai.2021/596 Survey Track.

[6] Xin Chen, Qizhi Fang, Wenjing Liu, and Yuan Ding. 2020. Strategyproof mechanisms for 2-facility location games with minimax envy. In *Algorithmic Aspects in Information and Management: 14th International Conference, AAIM 2020, Jinhua, China, August 10–12, 2020, Proceedings 14*. Springer, Springer International, Berlin, 260–272.

[7] Mark S. Daskin and Kayse Lee Maass. 2015. *The p-Median Problem*. Springer International Publishing, Cham, 21–45. https://doi.org/10.1007/978-3-319-13111-5_2

[8] Zvi Drezner (Ed.). 1995. *Facility Location: A Survey of Applications and Methods* (1 ed.). Springer-Verlag New York, Inc., New York, NY, USA. XVIII, 572 pages. https://doi.org/10.1007/978-1-4612-5357-0

[9] Zvi Drezner and Charles A. Scott. 2001. Facility location in a stochastic environment. *Mathematical and Computer Modelling* 33 (2001), 11–22.

[10] Reza Zanjirani Farahani and Masoud Hekmatfar (Eds.). 2009. *Facility Location: Concepts, Models, Algorithms and Case Studies* (1 ed.). Physica Heidelberg, Heidelberg, Germany. VIII, 549 pages. https://doi.org/10.1007/978-3-7908-2151-2

[11] 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.

[12] Dimitris Fotakis. 2008. On the competitive ratio for online facility location. *Algorithmica* 50, 1 (2008), 1–57.

[13] Dimitris Fotakis and Christos Tzamos. 2013. Strategyproof facility location for concave cost functions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce* (Philadelphia, Pennsylvania, USA) *(EC '13)*. Association for Computing Machinery, New York, NY, USA, 435–452. https://doi.org/10.1145/2492002.2482595

[14] Swati Gupta, Jai Moondra, and Mohit Singh. 2024. Which L_p norm is the fairest? Approximations for fair facility location across all "p". arXiv:2211.14873 [cs.DS] https://arxiv.org/abs/2211.14873

[15] Tjalling C. Koopmans and Martin Beckmann. 1957. Assignment Problems and the Location of Economic Activities. *Econometrica* 25, 1 (1957), 53–76. http://www.jstor.org/stable/1907742

[16] Alexander Lam. 2021. Balancing Fairness, Efficiency and Strategy-Proofness in Voting and Facility Location Problems. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Virtual Event, United Kingdom) *(AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1818–1819.

[17] Alexander Lam, Haris Aziz, Bo Li, Fahimeh Ramezani, and Toby Walsh. 2024. Proportional Fairness in Obnoxious Facility Location. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) *(AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1075–1083.

[18] Alexander Lam, Haris Aziz, and Toby Walsh. 2023. Nash Welfare and Facility Location. arXiv preprint arXiv:2310.04102.

[19] Markus Leitner, Ivana Ljubić, Juan-José Salazar-González, and Markus Sinnl. 2014. On the Asymmetric Connected Facility Location Polytope. In *Combinatorial Optimization*, Pierre Fouilhoux, Luis Eduardo Neves Gouveia, A. Ridha Mahjoub, and Vangelis T. Paschos (Eds.). Springer International Publishing, Cham, 371–383.

[20] Jiaqian Li, Minming Li, and Hau Chan. 2024. Strategyproof Mechanisms for Group-Fair Obnoxious Facility Location Problems. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 9 (Mar. 2024), 9832–9839. https://doi.org/10.1609/aaai.v38i9.28843

[21] M Teresa Melo, Stefan Nickel, and Francisco Saldanha-Da-Gama. 2009. Facility location and supply chain management–A review. *European journal of operational research* 196, 2 (2009), 401–412.

[22] A. Meyerson. 2001. Online Facility Location. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, Washington, DC, USA, 426–431. https://doi.org/10.1109/SFCS.2001.959917

[23] Hervé Moulin. 2004. *Fair Division and Collective Welfare*. MIT Press, Cambridge, MA, USA.

[24] Kanad Pardeshi, Itai Shapira, Ariel D. Procaccia, and Aarti Singh. 2024. Learning Social Welfare Functions. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., Red Hook, NY, USA, 41733–41766. https://proceedings.neurips.cc/paper_files/paper/2024/file/499b12df1531fe8ee0febcf08381f3a4-Paper-Conference.pdf

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

[26] Ariel D. Procaccia, David Wajc, and Hanrui Zhang. 2018. Approximation-Variance Tradeoffs in Facility Location Games. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, Palo Alto, California, USA, 1185–1192. https://doi.org/10.1609/AAAI.V32I1.11451

[27] Kariv O Hakimi SL. 1979. An algorithmic approach to network location problems. II: the p-medians. *SIAM J. Appl. Math* 37, 3 (1979), 539.

[28] Lawrence V Snyder. 2006. Facility location under uncertainty: a review. *IIE transactions* 38, 7 (2006), 547–564.

[29] Arie Tamir. 1991. Obnoxious facility location on graphs. *SIAM Journal on Discrete Mathematics* 4, 4 (1991), 550–567.

[30] Avyukta Manjunatha Vummintala, Shivam Gupta, Shweta Jain, and Sujit Gujar. 2025. FLIGHT: Facility Location Integrating Generalized, Holistic Theory of Welfare. arXiv:2502.14732 [cs.GT] https://arxiv.org/abs/2502.14732

[31] Ying Wang, Houyu Zhou, and Minming Li. 2024. Positive Intra-Group Externalities in Facility Location. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) *(AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1883–1891.

[32] Alfred Weber. 1909. *Über den Standort der Industrien*. Mohr Siebeck, Tübingen, Germany.

[33] Ling-Yun Wu, Xiang-Sun Zhang, and Ju-Liang Zhang. 2006. Capacitated facility location problem with general setup cost. *Computers & Operations Research* 33, 5 (2006), 1226–1241.

[34] Morteza Zadimoghaddam and Aaron Roth. 2012. Efficiently learning from revealed preference. In *Proceedings of the 8th International Conference on Internet and Network Economics* (Liverpool, UK) *(WINE'12)*. Springer-Verlag, Berlin, Heidelberg, 114–127. https://doi.org/10.1007/978-3-642-35311-6_9