Inventory Design for Mobile Food Markets Considering Customer Purchasing Behavior

Final Report

Yiling Zhang

Industrial and Systems Engineering

University of Minnesota

CTS 25-01

CENTER FOR TRANSPORTATION STUDIES
UNIVERSITY OF MINNESOTA

Technical Report Documentation Page

1. Report No.	2.	3. Recipients Accession No.
CTS 25-01		

4. Title and Subtitle	5. Report Date
Inventory Design for Mobile Food Markets Considering Customer Purchasing Behavior	March 2025
	6.
7. Author(s) Zijian Wang, William L. Cooper, Yiling Zhang	8. Performing Organization Report No.
9. Performing Organization Name and Address Industrial and Systems Engineering University of Minnesota 207 Church Street SE, Minneapolis, MN 55455	10. Project/Task/Work Unit No. #2023020
	11. Contract (C) or Grant (G) No.
12. Sponsoring Organization Name and Address Center for Transportation Studies University of Minnesota 440 University Office Plaza 2221 University Avenue SE Minneapolis, MN 55414	13. Type of Report and Period Covered Final Report
	14. Sponsoring Agency Code

15. Supplementary Notes

https://www.cts.umn.edu/publications/researchreports/

16. Abstract (Limit: 250 words)

Food insecurity remains a pressing challenge in the United States, driven in part by limited access to affordable and nutritious food. Mobile markets—vehicles that sell healthy food directly to underserved communities—offer a promising solution. However, many mobile markets face operational and financial challenges due to complex interactions between inventory decisions, customer purchasing behavior, and logistical constraints.

In this project, we propose a bilevel optimization framework to design inventory strategies for mobile markets operating multiple stops on a single trip. The upper-level (leader) problem models the operator's inventory decisions under vehicle capacity constraints, aiming to maximize the total utility of all customers. The lower-level (follower) problem models individual customer purchasing behaviors, which depend on available inventory and budget constraints at the time of their visit. Our model captures the interdependencies between customers, as purchases made at earlier stops affect product availability for later customers.

We provide an exact reformulation and develop a heuristic solution method to address computational complexity. Theoretical performance bounds for the heuristic are established. Numerical experiments demonstrate the efficiency and robustness of our proposed approach and provide insights into how mobile markets can improve customer satisfaction and operational outcomes through strategic inventory planning.

This work lays the foundation for future studies that will incorporate additional operational decisions, such as pricing, routing, and the inclusion of stochastic customer behavior, to further support the financial and social sustainability of mobile market programs.

17. Document Analysis/Descriptors		18. Availability Statement	
Food, Access, Underserved communities, Grocery stores, Inventory control, Consumer behavior, Optimization		No restrictions. Document available from: National Technical Information Services, Alexandria, Virginia 22312	
19. Security Class (this report)	**20. Security Class (this page)**	**21. No. of Pages**	**22. Price**
Unclassified	Unclassified	31	

Inventory Design for Mobile Food Markets Considering Customer Purchasing Behavior

Zijian Wang, William L. Cooper, Yiling Zhang

March 2025

Acknowledgment

The authors would like to thank the Center for Transportation Studies at the University of Minnesota for providing generous funding to support this research. In addition, the authors would like to thank Stephanie Wagner at the Food Group for the extensive discussions that helped shape the ideas.

Contents

List of Figures

List of Tables

Executive Summary

Food insecurity, driven in part by limited access to grocery stores, affects millions across the U.S., including over 9% of Minnesotans. Mobile markets, which bring affordable and healthy foods directly to underserved communities via vehicles, have emerged as a promising solution to alleviate this issue. However, many mobile markets struggle with operational inefficiencies and financial sustainability.

In this study, we investigate how to optimize mobile market operations by focusing on inventory decisions and their relationship to customer purchasing behaviors. We model this as a bilevel optimization problem, where the mobile market operator chooses the inventory to load onto the vehicle, and sequentially visited customers make purchasing decisions based on product availability and budget constraints.

To address the computational challenges of this bilevel model, we develop an exact reformulation and propose a heuristic approach that balances solution quality and efficiency. Our numerical experiments demonstrate the effectiveness of these methods and provide insights into inventory planning for mobile markets. Numerical experiments also demonstrate the robustness of our methods in diverse scenarios.

This work offers a novel decision-making framework for mobile market operators, helping to improve service equity and operational efficiency. Future work will extend this framework to incorporate additional factors, such as routing decisions, pricing strategies, and stochastic customer behavior.

1 Introduction

Food insecurity refers to insufficient or inconsistent access to nutritious food. In 2020, the U.S. witnessed an alarming increase in food insecurity, affecting millions of people, including many children. Food insecurity may give rise to numerous challenges and can be a cause of unhealthy eating habits and diet-related diseases. In 2022, an estimated 537,000 people—more than 9% of the population—suffered from food insecurity in Minnesota (Feeding America, 2024). Unfortunately, this problem is not particular to Minnesota, and many other states have even higher rates of food insecurity. One cause of food insecurity is a lack of nearby full-service grocery stores. This is a particular issue to Minnesota, which has fewer grocery stores per capita than most other states.

Over the past decade, mobile markets have emerged as a means for mitigating food insecurity. A mobile market is a grocery store in the form of a bus, truck, semi-trailer, or other vehicle, typically outfitted with refrigeration equipment, that sells healthy food directly to communities in need at affordable prices. Mobile markets can be particularly helpful for people who do not live close to a full-service brick-and-mortar grocery store and who do not have consistent access to a car. For such individuals, travel to a grocery store may be prohibitively difficult. In Minnesota, mobile markets currently serve communities in the Twin Cities and Duluth. One example is the Twin Cities Mobile Market (TCMM), which is operated by a non-profit organization called the Food Group. TCMM currently makes scheduled stops at more than 20 different locations in the Minneapolis-Saint Paul metropolitan area. Beyond Minnesota, hundreds of mobile market programs collectively serve tens of thousands of people across the U.S. (Mobile Market Coalition, 2022).

Most mobile markets are operated by nonprofit organizations. Unlike traditional retailers, their primary goal is typically not to generate profit but to increase food accessibility and encourage participation from low-income or disadvantaged communities (Kasprzak et al., 2022, 2021). These communities often face financial constraints and are highly sensitive to pricing, which makes affordability a crucial factor in mobile market operations (Zepeda et al., 2014; Weissman et al., 2020). In addition, offering culturally relevant foods can attract participation and foster a sense of belonging, further enhancing engagement and impact. However, many mobile markets struggle with financial sustainability, making it difficult to maintain and expand their services. While there is growing interest in understanding and evaluating the health impact of mobile markets (e.g., Robinson et al., 2016; Horning et al., 2021), there is a need for research on how to strategically and analytically optimize their operations to lower costs and potentially enhance viability.

Currently, pricing and inventory decisions in mobile markets are often made through trial and error or by benchmarking against local grocery retailers. However, these decisions involve complex system dynamics and constraints. For example, mobile markets operate under strict vehicle capacity constraints, requiring careful selection of stocked items to meet customer needs while minimizing food waste and stock shortage on each trip. Importantly, the purchasing behavior of customers is affected by the availability of products at the time of their visit. If high-demand items sell out early, customers at later stops may face limited options, potentially reducing overall sales and satisfaction.

To address these challenges, it is critical to account for the interdependence between inventory design and customer behavior. Some existing studies (Biesinger et al., 2016; Zhang et al., 2012; Aros-Vera et al., 2013) use discrete choice models to predict demand, where customers make a single choice from a fixed set of options. However, such models may not fully capture multi-item purchasing behaviors, which are common in grocery shopping when customers typically select multiple products. In addition, mobile markets typically serve multiple stops on a single trip without any opportunity for restocking. As a result, the order of neighborhood visits also impacts sales dynamics and customer experience. Visiting areas with high demand for certain items early in the trip can lead to rapid inventory depletion, leaving fewer choices for later stops and potentially discouraging participation from these communities. A carefully planned sequence of stops can help balance product availability, enhance revenue, and even ensure more equitable access across neighborhoods.

In this report, we propose a hierarchical sequential decision-making framework that maximizes total customer utility for a single trip while accounting for the interdependence between inventory decisions and customer purchasing behaviors under budget constraints. The problem is modeled as a bilevel program, where the leader represents a mobile market operator, and the followers are customers who are visited sequentially by the mobile market on a single trip. The operator's objective is to maximize the total utility of all customers by determining the initial inventory loaded onto the vehicle. Customers then make purchasing decisions that maximize individual utility while respecting budget constraints and available inventory at the time of their visits. The available inventory at each stop depends on both the operator's initial stocking decisions and the purchases made by earlier customers in the sequence. Most existing literature on bilevel optimization focuses on either single followers or multiple independent followers (Kleinert et al., 2021; Beck et al., 2023; Dempe, 2020). In contrast, our model involves multiple dependent followers, where each customer's decision is affected by both the leader's decisions and the purchasing decisions of other customers.

The main contributions of this research are three-fold:

- A bilevel optimization model is proposed with dependent followers for inventory design while accounting for the interdependence between customer purchasing behaviors and inventory design. A standard Karush-Kuhn-Tucker(KKT) based reformulation is provided to reduce the bilevel problem to a single-level formulation. However, this reformulation introduces complementarity conditions, leading to computational challenges even for medium-sized instances.

- To improve computation, we propose a heuristic approach that provides a suboptimal yet computationally efficient solution to the bilevel problem. This approach involves solving a series of optimization problems, each corresponding to a single customer's decision. We explore the structural properties of these subproblems to develop efficient solution methods. To quantify the optimality gap, theoretical bounds are derived.

- The effectiveness and efficiency of the proposed models and solution methods are demonstrated via numerical experiments, providing insights into inventory planning in mobile markets.

While this report mainly focuses on inventory decisions, future work will build on the models and methods developed here to explore neighborhood visit sequences and pricing strategies. In addition, as discussed in Section 6, we are interested in extending our framework to stochastic settings, incorporating uncertainty in customer behavior and demand patterns.

The remainder of this report is organized as follows. Section 2 reviews related literature. In Section 3, we describe the problem dynamics and present the problem formulation. Section

4 introduces a KKT-based reformulation and a relaxation formulation. In Section 5, we propose a heuristic method, analyze its solution structure, and establish performance bounds. While the previous sections focus on deterministic settings, Section 6 takes an initial step toward incorporating the randomness of customer information, setting up a foundation for future research on stochastic extensions. Finally, Section 7 presents numerical experiments, demonstrating the effectiveness and efficiency of the proposed models and solution methods.

2 Related Literature

Mobile markets have drawn increasing interest due to their significant potential to increase food access and improve health outcomes, particularly for disadvantaged communities (e.g., Hsiao et al., 2019). However, few studies focus on quantitative approaches to addressing the operational challenges faced by mobile markets, including routing, inventory, and customer engagement. Most existing research on food distribution focuses on centralized decision models, where resources are allocated from distribution centers. For example, food banks typically operate through centralized distribution hubs, where food and resources are collected, and then distributed to partner agencies such as community organizations, and shelters in various ways. Specifically, Davis et al. (2014) identify food delivery points (FDPs), where agencies can receive food deliveries. Wishon & Villalobos (2016) solve inventory and routing problems for mobile markets, assuming known customer demands. Liang & Lyu (2022) develop real-time food allocation rules for food banks, balancing fairness and efficiency when serving agencies with sequentially revealed demands. At a county level, Orgut et al. (2018) study equitable food distribution under uncertain county capacity. While these studies provide solution approaches and insights into logistics and resource allocation in centralized food distribution at organization levels, they differ from customer-level food distribution, where demand is affected by individual purchasing behaviors. This project aims to capture these customer behaviors in mobile market operations.

Our work is also related to sequential resource allocation problems, where resources are allocated incrementally in response to sequentially revealed demand. For example, Geng et al. (2014) study a two-customer sequential resource allocation problem with equity constraint, ensuring fair distribution across customers. Lien et al. (2014) extend this concept to nonprofit operations, such as food distribution. Recently, Hassanzadeh et al. (2023) address the sequential allocation of a divisible resource, formulating the problem as a Markov decision process and developing a fairness-driven algorithm. In contrast to these works, our model does not assume that the demand is explicitly revealed. Instead, demand is endogenously determined by customer purchasing behaviors, which depend on available inventory, budget constraints, and individual preferences. We leverage the problem structure to provide efficient solution approaches that optimize inventory allocation while accounting for customer behavior dynamics.

Lastly, some of our results are reminiscent of the price of anarchy (Koutsoupias & Papadimitriou, 1999; Papadimitriou, 2001), which captures the inefficiencies in systems where self-interested agents make decentralized decisions without coordination. In particular, our study explores the trade-off between maximizing total utility and preserving customer choice autonomy. Rather than imposing centralized control, our framework allows for individualized purchasing decisions, capturing the interplay between optimal inventory design and customer freedom in selecting from available stock.

We close this section by noting several empirical studies that evaluate the effectiveness of mobile markets through surveys and/or focus groups to assess the potential impacts and operational challenges of mobile markets. Examples of such work include papers by Zepeda et al. (2014), Hsiao et al. (2018), Weissman et al. (2020), and Horning et al. (2021). In particular, Kasprzak et al. (2022) discuss various high-level operational challenges faced by mobile markets.

3 Problem Formulation

Consider a mobile food market operated on a vehicle with a daily capacity of $C > 0$. This capacity can represent either a physical capacity available for stocking items or a financial budget allocated for purchasing items to put on the vehicle. For simplicity, we consider C to be a one-dimensional quantity for now. Assume that there are n products that the market operator can potentially stock in its vehicle. We denote the stock quantities of the n products by $\boldsymbol{q} = (q_1, q_2, \cdots, q_n)^\top$. Each unit of product j requires $w_j \geq 0$ units of the capacity and is sold at a price of $p_j > 0$. Denote $\boldsymbol{w} = (w_1, w_2, \cdots, w_n)^\top$ and $\boldsymbol{p} = (p_1, p_2, \cdots, p_n)^\top$. We will sometimes refer to w_j as the weight of product j. There are N customers. Each customer i has a budget $y_i > 0$ allocated for purchasing items from the mobile market.

For customer $i = 1, \ldots, N$, the quantity of product $j = 1, \ldots, n$ they purchase, denoted as $x_{ij} \in \mathbb{R}_+$ needs to be determined. Given the purchase decision $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{in})^\top$ for customer i, their total utility is represented by the function $u_i(\boldsymbol{x}_i)$. This utility function can capture customer satisfaction and broader benefits of purchasing items from the mobile market. These benefits may depend on factors such as personal preferences, product variety, and nutrition values, reflecting the varied needs and priorities of different customers. Throughout the remainder of this report we make the following assumption about the utility functions:

Assumption 1. *The utility function $u_i(\boldsymbol{x}_i) : [0, +\infty)^n \to \mathbb{R}$ is non-decreasing, concave, and differentiable for $i = 1, \ldots, N$.*

Assumption 1 arises naturally from the monotonicity of preference and the law of diminishing marginal utility. The non-decreasing property implies that a customer will weakly prefer more of a product to less of that product. The concavity implies that an additional unit of product will bring less incremental value to the customer when the customer has more of the product. The differentiability is needed to derive tractable solution approaches.

We now present a formulation in which each customer individually decides what to purchase. We assume that customers enter the market vehicle sequentially, ordered according to the index i. So, customer i is the ith customer to enter the market. If the mobile market vehicle makes several stops during the day, then the order in which customers enter is determined, in part, by the order of the stops. For instance, if there are three stops with N_1, N_2, and N_3 customers at those steps, then in our formulation, we have $N = N_1 + N_2 + N_3$ and customers $i = 1, \ldots, N_1$ are the customers from stop 1, customers $i = N_1 + 1, \ldots, N_1 + N_2$ are the customers from stop 2, and customers $i = N_1 + N_2 + 1, \ldots, N_1 + N_2 + N_3$ are the customers from stop 3.

Each customer i makes their purchase decision $\boldsymbol{x}_i$ to maximize their own utility, while not exceeding their budget y_i, and not exceeding the inventory available to purchase. The inventory available to customer i is determined by the initial inventory $\boldsymbol{q}$ and the cumulative purchases of previous customers (that is, the purchases of customers 1 through $i-1$). Hence, the inventory level will influence the customer's purchase decision, and the earlier customers' purchase decisions will affect the available inventory levels seen by later customers. Denote customer k's optimal purchase decision under the total inventory level $\boldsymbol{q}$ by $\boldsymbol{x}_k^*(\boldsymbol{q})$. Then the inventory available to customer i is given by $\boldsymbol{q} - \sum_{k=1}^{i-1} \boldsymbol{x_k^*(q)}$.

For each customer $i = 1, \dots, N$, the individual utility maximization problem is formulated as follows:

$$\begin{aligned} v_i^*(\boldsymbol{q}) = \max_{\boldsymbol{x_i}} \quad & u_i(\boldsymbol{x_i}) && (LLP_i) \\ \text{s.t.} \quad & \boldsymbol{p}^\top \boldsymbol{x_i} \le y_i \\ & 0 \le \boldsymbol{x_i} \le \boldsymbol{q} - \sum_{k=1}^{i-1} \boldsymbol{x_k^*(q)}. \end{aligned}$$

The preceding may be viewed as a "Lower-Level Problem" in the sense that solving it requires a value for the initial inventory level $\boldsymbol{q}$ which is set first by the market operator.

Note that the optimal purchase decision of customer 1 can be viewed as a function of the initial inventory $\boldsymbol{q}$. Moreover, the purchase decision $\boldsymbol{x_k}$ of each customer $k > 1$ can be viewed as a function of $\boldsymbol{q}$ and $\boldsymbol{x_1}, \dots, \boldsymbol{x_{k-1}}$. Hence, by arguing inductively, we may view the purchase decision $\boldsymbol{x_i}$ of customer i as well as that customer's associated utility simply as functions of $\boldsymbol{q}$. This is reflected in the notation $v_i^*(\boldsymbol{q})$ in formulation (LLP_i) above.

Notice that here, for any given $\boldsymbol{q}$, we will get a sequence of customer utilities $\{v_i^*(\boldsymbol{q})\}_{i=1}^N$ as well as a sequence of customer purchase vectors $\{\boldsymbol{x_i^*(q)}\}_{i=1}^N$. Although the market operator here cannot directly decide the amounts of products obtained by each customer, it can still affect the overall value received by the customer population through the selection of $\boldsymbol{q}$. So, the utility-sum maximization problem faced by the market operator can be written as:

$$\begin{aligned} z^* = \max_{\boldsymbol{q}} \quad & \sum_{i=1}^N v_i^*(\boldsymbol{q}) && (P) \\ \text{s.t.} \quad & \boldsymbol{w}^\top \boldsymbol{q} \le C \\ & \boldsymbol{q} \ge 0, \end{aligned}$$

where the utilities under the optimal choices of the customers $\{v_i^*(\boldsymbol{q})\}_{i=1}^N$ are defined by the sequence of the lower-level problems (LLP_i). Problem (P) is a bilevel optimization problem with dependent followers, where each follower sequentially solves their problem (LLP_i), making decisions that depend on both the leader's decision and the actions of preceding followers. We will discuss solution approaches for Problem (P) in the next section.

The formulation above assumes that the sequence of customer arrivals is fixed. However, this sequence can be altered by decisions of the mobile market itself. In particular, if the market vehicle alters the order in which it visits various locations or changes the locations that it visits, then the customer sequence will change. A solution to our model that assumes a particular fixed sequence of visits (and hence a fixed customer sequence) can be used in an analysis that decides what sequence of visits to make. In particular, if we are evaluating a set of potential visit sequences, the model above can be used to quantify the desirability of each such sequence with an eye toward picking the best one. Choosing such a visit sequence would also likely include various other considerations such as inventory-related costs and travel times, which are beyond the scope of the present analysis. We expect this to be a topic for future exploration.

4 Solution Approaches

In this section, we describe how to solve the bilevel problem (P) introduced in the previous section. As it is posed, solving it requires optimizing N-dependent lower-level problems (one for each customer). Our basic approach is to use the Karush–Kuhn–Tucker (KKT) conditions (Boyd & Vandenberghe, 2004) from the lower level problems (LLP_i) to reformulate the bilevel problem as a single optimization. This general approach to bilevel optimization problems is described in, for example, Dempe (2002, 2020).

4.1 KKT Formulation

The Lagrangian of lower-level problem (LLP_i) is as follows:

$$\mathcal{L}_i(\boldsymbol{x}_i; \lambda_i, \boldsymbol{r}_i, \boldsymbol{l}_i) = u_i(\boldsymbol{x}_i) + \lambda_i(y_i - \boldsymbol{p}^\top \boldsymbol{x}_i) + \boldsymbol{r}_i(\boldsymbol{q} - \sum_{k=1}^{i} \boldsymbol{x}_k) + \boldsymbol{l}_i \boldsymbol{x}_i,$$

where $\lambda_i \in \mathbb{R}_+$, $\boldsymbol{r}_i \in \mathbb{R}^n_+$, and $\boldsymbol{l}_i \in \mathbb{R}^n_+$ are Lagrangian multipliers associated with the budget constraint, inventory constraint, and nonnegativity constraint in (LLP_i).

The KKT first order condition is:

$$\nabla \mathcal{L}_i = \nabla u_i(\boldsymbol{x}_i) - \lambda_i \boldsymbol{p} - \boldsymbol{r}_i + \boldsymbol{l}_i = 0,$$

where $\nabla u_i(\boldsymbol{x}_i)$ is the gradient of $u_i(\boldsymbol{x}_i)$. The KKT conditions also include the complementary slackness conditions:

$$\lambda_i(y_i - \boldsymbol{p}^\top \boldsymbol{x}_i) = 0$$

$$\boldsymbol{r}_i(\boldsymbol{q} - \sum_{k=1}^{i} \boldsymbol{x}_k) = 0$$

$$\boldsymbol{l}_i \boldsymbol{x}_i = 0.$$

We will plug those conditions along with the feasibility conditions of (LLP_i) into the problem (P). We have the following observation: if any solution $\boldsymbol{q}^*$ is optimal for problem (P), and the related lower-level optimal solutions are $\{\boldsymbol{x}_i^*(\boldsymbol{q}^*)\}_{i=1}^N$, then the solution $\boldsymbol{q} = \sum_{i=1}^N \boldsymbol{x}_i^*(\boldsymbol{q}^*)$ will also be optimal to (P). Hence, in the whole problem with those KKT conditions, we simply let $\boldsymbol{q} = \sum_{i=1}^N \boldsymbol{x}$, and there is no loss of optimality. As a result, we have:

$$\max_{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_N} \quad \sum_{i=1}^{N} u_i(\boldsymbol{x}_i) \tag{P_1}$$

$$\text{s.t.}: \quad \boldsymbol{w}^\top (\sum_{i=1}^{N} \boldsymbol{x}_i) \le C$$

$$\begin{aligned}
&\boldsymbol{p}^\top \boldsymbol{x}_i \le y_i, \forall i = 1,\dots,N\\
&\lambda_i \boldsymbol{p} + \boldsymbol{r}_i - \boldsymbol{l}_i = \nabla u_i(\boldsymbol{x}_i), \forall i = 1,\dots,N\\
&\lambda_i (y_i - \boldsymbol{p}^\top \boldsymbol{x}_i) = 0, \quad \forall i = 1,\dots,N\\
&\boldsymbol{r}_i (\sum_{k=i+1}^{N} \boldsymbol{x}_k) = 0, \quad \forall i = 1,\dots,N\\
&\boldsymbol{l}_i \boldsymbol{x}_i = 0, \quad \forall i = 1,\dots,N\\
&\boldsymbol{x}_i, \lambda_i, \boldsymbol{r}_i, \boldsymbol{l}_i \ge 0, \quad \forall i = 1,\dots,N.
\end{aligned}$$

Under Assumption 1, if $\{\boldsymbol{x_i}\}_{i=1}^N$ is an optimal solution to problem (P_1), then $\boldsymbol{q} = \sum_i \boldsymbol{x_i}$ is an optimal to problem (P).

Note that in the preceding optimization problem (P_1), the constraints from the complementary slackness conditions are non-linear. Using standard approaches, it is possible to re-pose those non-linear constraints as linear constraints through the introduction of binary variables. Doing so yields an equivalent formulation that is amenable to solution with standard optimization algorithms (depending upon our specific assumptions about the customers' utility functions). Please refer to the appendix for the details.

4.2 Formulation Relaxation: Central Decision Maker Problem

In general, (P_1) may be difficult to solve, especially if there are many customers or many products. An alternative is to look into its relaxations. In fact, removing all complementarity constraints in (P_1) results in a relaxation

$$\begin{aligned}
\bar{z} = \max_{\{\boldsymbol{x_i}\}} \quad & \sum_{i=1}^{N} \sum_{j=1}^{n} u_{ij}(x_{ij}) && (P_2)\\
s.t.: \quad & \sum_{i=1}^{N} \boldsymbol{w}^\top \boldsymbol{x_i} \le C\\
& \boldsymbol{p}^\top \boldsymbol{x_i} \le y_i, \forall i = 1,\dots,N\\
& 0 \le \boldsymbol{x_i}, \forall i = 1,\dots,N.
\end{aligned}$$

It is now apparent that Problem (P_2) provides an upper bound on (P). That is,

$$z^* \le \bar{z}\,.$$

Problem (P_2) can be viewed as the market operator providing a personalized bundle of products directly to each customer. The operator does this in a way that maximizes the total utility across all customers. Hence, we will sometimes call it a "central decision maker" problem. Here, we emphasize that customers do not make decisions, but rather the market operator may be viewed as making decisions on their behalf. Note that the sequencing of customers plays no role in Problem (P_2). This is a major difference from Problem (P), where the sequencing of customers has a major role in determining which customers buy which products. The primary reason for our interest in Problem (P_2) is that it provides an upper bound on the optimal objective value of Problem (P). It is important to have an upper bound as a means for evaluating heuristic solutions of (P). If we cannot solve (P) to optimality, then we may need to resort to heuristics. An upper bound allows us bound the extent to which a heuristic is suboptimal.

Observe that in Problem (P_2) there are two groups of constraints (the capacity constraint and budget constraints), both of which are linear. In addition, the objective function is concave. Consequently, this is a well-behaved convex optimization problem, which can be readily solved with the use of commercial solvers.

5 A Computationally Tractable Heuristic

Even with the reformulation developed in the previous section, it is not practical to solve Problem (P_1) to optimality when the problem scale is large. In this section, we propose a scalable heuristic method.

Note that in Problem (P), the function $v_i(\boldsymbol{q})$ is influenced by customer purchasing decisions, which in turn depend on the initial inventory $\boldsymbol{q}$. One intuitive approach is providing customers with items they want. Inspired by this idea, we introduce a heuristic (sub-optimal) solution for Problem (P), which determines inventory decisions by prioritizing the needs of the earliest arriving customers.

Let $\boldsymbol{x_i^*(q)}$ denote the optimal solution of Problem (LLP_i) if the initial inventory at the beginning of the day is $\boldsymbol{q}$. That is, $\boldsymbol{x_i^*(q)}$ is the purchase decision of customer i if the initial inventory decision is $\boldsymbol{q}$. With this notational convention, $\boldsymbol{x_i^*(\infty)}$ represents the optimal purchase decision of customer i if there is no inventory constraint (or, thought of differently, if there is infinite initial inventory). To be more precise, $\boldsymbol{x_i^*(\infty)}$ is the solution to the following problem:

$$\begin{aligned} \max_{\boldsymbol{x_i}} \quad & u_i(\boldsymbol{x_i}) && (LLP_i') \\ s.t.: \quad & \boldsymbol{p}^\top \boldsymbol{x_i} \le y_i \\ & 0 \le \boldsymbol{x_i}. \end{aligned}$$

Observe that the lower level problems (LLP_i) are dependent on the preceding followers through the inventory constraints. If the initial inventory is infinite, then the inventory constraints go away. Consequently, the lower level problems de-couple and each Problem (LLP_i') can be solved in isolation.

In preparation for the ensuing analysis, define

$$N^* = \max\Big\{k \in 1, \dots, N : \sum_{i=1}^{k-1} \boldsymbol{w}^\top \boldsymbol{x_i^*(\infty)} \le C\Big\}.$$

That means, if we continue giving customers their optimal choice when there is no inventory constraint, we will run out of our capacity at N^*. There, we cannot fulfill the N^*th customers' optimal choice. Instead, we are going to solve the following optimization problem:

$$\begin{aligned} \max_{\boldsymbol{x_{N^*}}} \quad & u_{x_{N^*}}(\boldsymbol{x_{N^*}}) && (LLP^*) \\ \text{s.t.}: \quad & \boldsymbol{p}^\top \boldsymbol{x_{N^*}} \le y_i \end{aligned}$$

$$
\begin{aligned}
&\boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{N^*}} \le C - \sum_{i=1}^{N^*-1} \boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{i}}^*(\infty) \\
&0 \le \boldsymbol{x}_{\boldsymbol{N^*}}.
\end{aligned}
$$

Below, we will construct a heuristic algorithm for the initial inventory level under which customers $i = 1, \dots, N^* - 1$ each purchases $\boldsymbol{x}_{\boldsymbol{i}}^*(\infty)$ and customer N^* purchases according to the optimal solution of Problem (LLP^*).

Algorithm 1 Computing the Heuristic

```
q^H ← 0
for i = 1, 2, ··· , N do
    x_i^* ← Solve (LLP'_i)
    cost ← w^T x_i^*
    if C ≥ cost then
        q^H ← q^H + x_i^*
        C ← C − cost
    else
        N^* = i
        x_{N^*}^H ← Solve (LLP^*)
        q^H ← q^H + x_{N^*}^H
        break
    end if
end for
return q^H
```

Algorithm 1 returns $\boldsymbol{q^H}$ where $\boldsymbol{q^H} = \sum_i^{N^*-1} \boldsymbol{x}_{\boldsymbol{i}}^*(\infty) + \boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}}$ and $\boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}}$ is the solution to Problem (LLP^*). By the construction here, we are hoping that the customers will act as we expect: the first $N^* - 1$ customers will choose $\boldsymbol{x}_{\boldsymbol{i}}^*(\infty)$ correspondingly, and customer N^* will choose $\boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}}$. The following lemma guarantees that if the initial inventory level is $\boldsymbol{q^H}$ then the customers will indeed do so.

Lemma 1. *Under Assumption 1, $\boldsymbol{q^H}$ is a feasible solution of Problem (P), and the corresponding lower level optimal solution is:*

$$
\boldsymbol{x}_{\boldsymbol{i}}^{\boldsymbol{H}} = \begin{cases} \boldsymbol{x}_{\boldsymbol{i}}^*(\infty), & if \quad i < N \\ \boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}}, & if \quad i = N^* \\ 0, & if \quad i > N^*. \end{cases} \tag{1}
$$

Proof. Since $\boldsymbol{x}_{\boldsymbol{i}}^*(\infty) \ge 0$ for all $i = 1, \dots, N^* - 1$ and $\boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}} \ge 0$, it is easy to see $\boldsymbol{q^H} = \sum_i^{N^*-1} \boldsymbol{x}_{\boldsymbol{i}}^*(\infty) + \boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}} \ge 0$. Moreover,

$$
\begin{aligned}
\boldsymbol{w}^\top \boldsymbol{q^H} &= \sum_{i=1}^{N^*-1} \boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{i}}^*(\infty) + \boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{N^*}}^{\boldsymbol{H}} \\
&\le \sum_{i=1}^{N^*-1} \boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{i}}^*(\infty) + C - \sum_{i=1}^{N^*-1} \boldsymbol{w}^\top \boldsymbol{x}_{\boldsymbol{i}}^*(\infty)
\end{aligned}
$$

$$= C,$$

which means $\boldsymbol{q^H}$ defined here is feasible for problem (P).

Next, we will show that under initial inventory level $\boldsymbol{q^H}$, the optimal solution to the lower level problem is $\{\boldsymbol{x_i^H}\}_{1\le i\le N}$. Notice that problem (LLP'_i) is a relaxation to problem (LLP_i), thus if the optimal solution to problem (LLP'_i), $\boldsymbol{x_i^*(\infty)}$, is feasible to (LLP_i), then it must also be optimal to (LLP_i). Following this analysis we have $\boldsymbol{x_i^*(q)} = \boldsymbol{x_i^*(\infty)}$ for any $\boldsymbol{q} \geq \sum_{k=1}^{i} \boldsymbol{x_k^*(\infty)}$ (hence $\boldsymbol{x_i^*(\infty)} \leq \boldsymbol{q} - \sum_{k=1}^{i-1} \boldsymbol{x_k^*(\infty)}$ is feasible).

Here $\boldsymbol{q^H} = \sum_{i}^{N^*-1} \boldsymbol{x_i^*(\infty)} + \boldsymbol{x_{N^*}^H} \geq \sum_{i}^{N^*-1} \boldsymbol{x_i^*(\infty)}$. Thus we know that for any $i = 1, 2, \cdots, N^*-1$, $\boldsymbol{x_i^*}(\boldsymbol{q^H})$ must equal to $\boldsymbol{x_i^*}(\infty)$. Next, consider the following problem faced by customer N^*:

$$\begin{aligned} \max_{\boldsymbol{x_{N^*}}} \quad & \sum_{j=1}^{n} u_{N^*j}(\boldsymbol{x_{N^*j}}) \\ \text{s.t. :} \quad & \boldsymbol{p}^\top \boldsymbol{x_{N^*}} \leq y_{N^*} \\ & 0 \leq \boldsymbol{x_{N^*}} \leq \boldsymbol{x_{N^*}^H}. \end{aligned}$$

From problem (LLP^*), we know that $\boldsymbol{x_{N^*}^H} \geq 0$, and $\boldsymbol{p}^\top \boldsymbol{x_{N^*}^H} \leq y_{N^*}$, hence $\boldsymbol{x_{N^*}^H}$ is a feasible solution to the problem above.

According to Assumption 1, for any j, $u_{N^*}(\boldsymbol{x_{N^*}})$ is non-decreasing. Thus we know:

$$u_{N^*}(\boldsymbol{x_{N^*}}) \leq u_{N^*}(\boldsymbol{x_{N^*}^H}),$$

for all $\boldsymbol{x_{N^*}} \leq \boldsymbol{x_{N^*}^H}$. Hence $\boldsymbol{x_{N^*}^H}$ is the optimal solution to the problem above.

For any $i > N^*$, there is no inventory left. So we have the optimal solution to (LLP_i) will be exactly $\mathbf{0}$. □

According to Lemma 1, we know that $\boldsymbol{q^H}$ is a feasible solution to (P). Therefore, the next proposition follows immediately.

Proposition 1. *Define* $\underline{z} := \sum_{i=1}^{N} u_i(\boldsymbol{x_i^H}) = \sum_{i=1}^{N^*-1} u_i(\boldsymbol{x_i^*(\infty)}) + u_{N^*}(\boldsymbol{x_{N^*}^H}) + \sum_{i=N^*+1}^{N} u_i(\mathbf{0})$. *Then under Assumption 1 we have*

$$\underline{z} \leq z^*.$$

6 Introducing Randomness

Up until this point, we have focused exclusively on deterministic problems. In particular, we have assumed that the mobile market operator knows the number of customers, the order of customer arrivals, and also those customers' budgets and utility functions. A direction for future work is to consider versions of the problem where some or all of these quantities are random. It is likely that such problems will be quite complex and difficult to solve to optimality. In such cases, it is important to identify tractable approaches for obtaining initial inventory levels. One such approach is to use the deterministic formulations introduced earlier in this report as a means for generating inventory policies that can be used in the stochastic setting. Prior work in revenue management has shown that in some cases policies based upon deterministic formulations can have provably good performance when used in stochastic settings; see, e.g., Gallego & Van Ryzin (1997) or Cooper (2002). Hence, there is reason to hope that such an approach may work well for mobile market problems.

In this section, we introduce two stochastic models that we plan to further explore in future work. The models include random customer arrivals and random customer types within a stop/neighborhood. Assuming homogeneous neighborhoods at each stop, the first model considers random customer arrivals, introducing uncertainty in the number of customers served at each stop. When the within-neighborhood differences around a stop cannot be ignored, the second model introduces random customer types to account for variability in demographics, purchasing behavior, and product preferences.

6.1 Random Numbers of Arrivals

On a given day, a mobile market vehicle will make a sequence of stops in a pre-defined order. An individual stop will draw customers from a single neighborhood or narrow geographical region. Within one neighborhood people may often have reasonably similar preferences and income levels. Hence, one may make the simplifying assumption that individual customers at a stop will all have the same utility function and the same budget. We can introduce randomness by assuming that the number of customers in each neighborhood is stochastic. To incorporate a setting like this into our model, assume there are L customer types, each corresponding to a neighborhood, indexed $l = 1, \dots, L$. The number of type l customers N_l has a discrete distribution F_l. We can think of each type as corresponding to a vehicle stop. The first stochastic optimization problem could thus be formulated as:

$$\begin{aligned} \tilde{z}^* = \max_{\boldsymbol{q}} \quad & \sum_{l=1}^{L} \mathbb{E}_{\{N_l \sim F_l\}} \left[\sum_{i=1}^{N_l} u_l^*(\boldsymbol{q}) \right] \\ \text{s.t.}: \quad & \boldsymbol{w}^\top \boldsymbol{q} \le C \\ & \boldsymbol{q} \ge 0. \end{aligned} \tag{SP_1}$$

The form of the lower-level problems here stays unchanged from our deterministic setting. Nevertheless, Problem (SP_1) is considerably more complex than what we considered earlier owing to the expectation in the objective function.

6.2 Random Customer Types

Here, we consider another way to introduce randomness into the model. For simplicity, we assume for now that the market visits only a single neighborhood, and that people from that neighborhood have utility functions independently drawn from a common distribution $\mathcal{U}$ over the space of utility functions. A simple version of this could be that utility functions are piecewise linear or have some other parametric form, and there is a joint distribution over the real-valued parameters of the functions. The second stochastic optimization problem is as follows:

$$\begin{aligned} z^* = \max_{\boldsymbol{q}} \quad & \mathbb{E}_{\{u_i^* \sim \mathcal{U}\}} \left[\sum_{i=1}^{N} u_i^*(\boldsymbol{q}) \right] \\ \text{s.t.}: \quad & \boldsymbol{w}^\top \boldsymbol{q} \le C \end{aligned} \tag{SP_2}$$

$$\boldsymbol{q} \geq 0.$$

Although there are numerous ways to introduce randomness, including combining the two previously discussed models, first understanding these two fundamental models provides a foundation for exploring more complex stochastic scenarios.

7 Numerical Experiments

In this section, we provide numerical results to further understand the relationship among the optimal value z^* of Problem (P), the optimal value $\bar{z}$ of the upper bound problem (P_2), and the value $\underline{z}$ associated with the heuristic introduced in Section 5. Following our previous analysis, we know that $\underline{z} \leq z^* \leq \bar{z}$. In the numerical experiments, we focus on problems where customer utility functions are piecewise linear, specifically in the form

$$u(\boldsymbol{x}) = \sum_{j=1}^{n} a_j [x_j \mathbf{1}\{x_j \leq b_j\} + b_j \mathbf{1}\{x_j > b_j\}].$$

This function has two pieces for each product j. When a customer's purchase quantity x_j is no more than the breakpoint b_j, each unit of the product contributes a utility of a_j, resulting in a total utility of $a_j x_j$; when the purchase quantity exceeds b_j, the total utility remains constant at $a_j b_j$. One way to interpret the breakpoint is that it represents the maximum quantity a customer is willing to purchase for product j, given their preferences or the existing stock they already have at home. Beyond this threshold, additional units provide no further value to the customer.

We randomly generate instances with $N = 10$ customers and $n = 20$ products. Let $\mathcal{U}[L, U]$ denote a uniform distribution on the interval $[L, U]$. Other parameters are generated based on the following assumptions.

- Utility coefficients: $a_{ij} \sim \mathcal{U}[0.5, 1.5]$, sampled independently $i = 1, \ldots, N$, $j = 1, \ldots, n$
- Utility breakpoints: $b_{ij} \sim \mathcal{U}[0.5, 1.5]$, sampled independently $i = 1, \ldots, N$, $j = 1, \ldots, n$
- Product prices: $p_j = 1$ for $j = 1, \ldots, n$
- Product weights: $w_j \sim \mathcal{U}[0.5, 1.5]$, sampled independently $j = 1, \ldots, n$
- Customer budgets: $y_i = \tilde{y}_i \sum_{j=1}^{n} b_{ij}$, where $\tilde{y}_i \sim \mathcal{U}[0.5, 1]$, sampled independently for $i = 1, \ldots, N$
- Vehicle capacity: $C = \tilde{C} \sum_{i=1}^{N} y_i$, where $\tilde{C} \sim \mathcal{U}[0.5, 1]$.

7.1 Computational Performance

In this section, we study the computational performance of solving the MILP exact reformulation to obtain z^* and implementing the heuristic method to compute $\underline{z}$. Key factors that influence the performance include the number of customers N and the number of items n. To study the impact of N and n, we generate instances following the parameter generation process described above. We evaluate performance across various values of n and

$N \in \{10, 20, 30, 40\}$. When nN is large, solving the MILP for even a single instance can be slow. Consequently, for each combination of N and n we sampled 10 instances. In Table 1 we report the average computational times in seconds for different values of N and n. A single number in that table shows the average, across 10 instances, of the number of seconds it took to solve the MILP for the corresponding N and n. For example, the average time it took to solve a problem with $N = 20$ customers and $n = 40$ products was 30.45 seconds. The MILP was solved using Gurobi with MATLAB on a desktop with 12th Gen Intel Core i7-12700 2.10GHz CPU, 32GB RAM, and NVIDIA T100 8GB GPU.

Table 1: Computational times to solve the MILP

	$N = 10$	$N = 20$	$N = 30$	$N = 40$
$n = 10$	0.24	1.11	1.86	5.89
$n = 20$	0.87	4.20	12.75	80.90
$n = 30$	2.50	8.10	74.66	217.72
$n = 40$	2.31	30.45	156.74	640.98

Alternative Text Description: Computational times to solve the MILP, while $N = 10$ or $n = 10$ the MILP could be solved in several seconds, while when $N = n = 40$, the MILP takes hundreds of minutes to solve.

Here, we can see that when the scale of the problem is small, solving the MILP is very efficient. However, with increasing n and N, the time cost of the MILP grows rapidly. The increase in computational time is more sensitive to N than to n, because the number of variables and constraints grows faster when N grows. With $n = 50$ and $N = 50$ (not shown in the table), some instances may need more than 10^6 seconds to solve. Another challenge for solving the MILP is the memory required. With $n = 200$ and $N = 100$, the memory needed to store the constraint matrix in the MILP exceeded the maximum that our version of MATLAB could process. The slower computation times for large problems serve to emphasize the importance of computationally practical heuristics, such as the one introduced herein.

To compare the time cost of solving the MILP (which gives us the optimal solution) and the heuristic from Section 5, we fix $nN = 1800$, and allow N to vary from the set $\{10, 20, 30, 60, 90, 180\}$; this makes n change correspondingly. The results are given in Table 2 below.

Table 2: Computational times with $nN = 1800$

	$N = 10$	$N = 20$	$N = 30$	$N = 60$	$N = 90$	$N = 180$
Optimal	47.98	182.71	737.46	2458.23	2817.51	1892.48
Heuristic	0.10	0.21	0.28	0.47	0.83	1.59

Alternative Text Description: Computational times with $nN = 1800$, for each combination of n and N, the time taken by solving to optimal is over hundreds of times more than the time taken by running the heuristic algorithm.

From Table 2, it can be seen that the heuristic solution can be calculated much more quickly than the optimal solution. The difference becomes particularly stark as the size of the problem increases. In the table, the time required to calculate the heuristic solution appears to grow almost linearly in N.

7.2 Fairness

In this section, we examine the fairness of the solutions that emerge from the problem formulations. For this, we consider the optimal and heuristic solutions to the sequential bilevel problem, as well as the central decision maker's problem that generates an upper bound to the bilevel problem. We will measure fairness by comparing the utility levels obtained by different customers. This is reasonable if the customers have similar utility functions. As one can imagine, in a setting where customers enter the market sequentially, one of the most powerful factors influencing a customer's obtained utility is whether that customer comes early or late in the sequence of customer arrivals. Customers late in the sequence can only make purchases when the vehicle capacity is large enough to allow an initial inventory level that can satisfy the demands of the earlier customers. Late customers also may not be able to obtain their preferred products because those have already been purchased by customers earlier in the sequence. This property of the sequential bilevel problem makes the outcome "less fair," compared to the central decision maker setting where the market operator assigns bundles of products to customers and where the order of customer arrival plays no role. We considered different ranges of the capacity C (as reflected in Figure 1 below) and sampled 1000 different instances for each setting. For simplicity, we have normalized the utilities with the first customer's utility and averaged the normalized utilities for each arrival position (arrivals are indexed $i = 1, \ldots, 10$). The results are summarized in Figure 1.

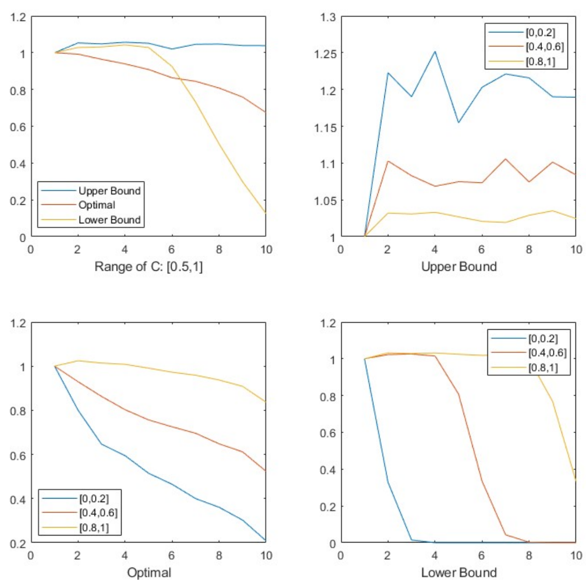

Figure 1: Normalized Utility Received by Customers in Different Arrival Positions

Alternative Text Description: Normalized Utility Received by Customers in Different Arrival Positions, when the capacity gets smaller, compared to the optimal solution, the heuristic solution tends to keep the utility received by the earlier coming customers and cut down the utility received by the later coming customers.

From the figure we can see that, as we discussed above, in our optimal and heuristic solutions to the sequential bilevel problem, the outcome will be less fair compared to the outcome in the central decision maker's problem. Moreover, in comparison to the optimal solution, the heuristic solution tends to be more unfair, because it is constructed in a way such that the initial inventory decision is guided by fully fulfilling the earlier customers' most preferred demands (see Section 5). In the central decision maker's problem, all customers' obtained utilities are nearly the same, because in our setting the parameter distributions of all customers are homogeneous. (The ratio here is not 1, because even if random variables X and Y are identically distributed, it does not necessarily mean that $\mathbb{E}[X/Y] = 1$.) For the optimal and heuristic solutions, we can see that as the capacity decreases, the later customers will receive less and less utility. The heuristic solution is more sensitive to this kind of change.

7.3 Evaluating the Heuristic and the Upper Bound

In this section, we study numerically how well $\bar{z}$ performs as an upper bound for z^* and how close the total utility $\underline{z}$ from the heuristic is to the optimal utility z^*. We do this by examining the ratios $\bar{z}/z^*$ and $\underline{z}/z^*$. As detailed earlier, these ratios are guaranteed to be at least 1 and at most 1 respectively. A value of $\bar{z}/z^*$ just slightly above 1 indicates that $\bar{z}$ is a good upper bound (i.e., it is not too much larger than the quantity that it is bounding). Likewise, a value of $\underline{z}/z^*$ just slightly below 1 indicates that our heuristic performs well (i.e., it yields total utility not too far below the maximum possible total utility). We sampled 200 instances as described above. Figure 2 plots the two ratios for the 200 instances. The plot labeled "Upper Bound" in the figure shows $\bar{z}/z^*$, and the plot labeled "Lower bound" in the figure shows $\underline{z}/z^*$.

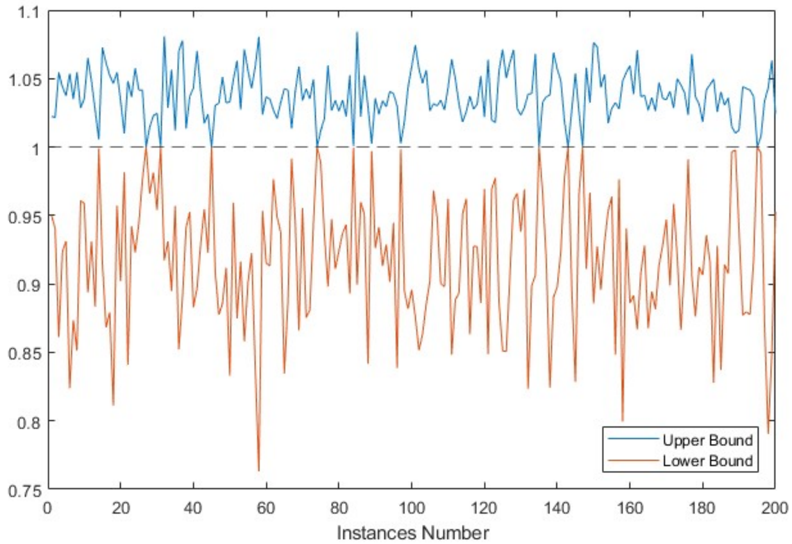

Figure 2: Upper and Lower Bounds in Different Instances

Alternative Text Description: Upper and Lower Bounds in Different Instances, among the instances, most of the upper bounds are less than 1.1 times of the optimal value, while most of the lower bounds are more than 0.8 times of the optimal value.

Here, we can see that the ratios are not very far from 1, indicating reasonably good performance for the instances shown. The ratios $\bar{z}/z^*$ all lie in the interval $[1, 1.1]$, and the ratios $\underline{z}/z^*$ all lie in the interval $[0.75, 1]$. The ratios coincide at 1 for some problem instances. For such instances, the heuristic yields the same total utility to customers as does a benevolent central decision maker that allocates bundles directly to customers in a socially optimal way.

We expect that if coefficients of customers' utility functions are close to each other (so that the customers all value products similarly), then the gap between $\underline{z}$ and $\bar{z}$ will be smaller. If the vector of product weights is nearly proportional to the vector of product prices (which is fixed to be $\vec{1}$ as explained above), then the objective value associated with the heuristic solution will be very close to the objective of the optimal solution. Also, if we have a smaller vehicle capacity, then solutions of the sequential bilevel problem will satisfy the earlier customers' demands while leaving little or nothing for customers later in the sequence. This issue does not arise for the central decision maker formulation wherein the sequencing of arrivals plays no role. In the next three figures in this subsection, we will examine these effects by varying the range of utility parameters, the range of weights, and the capacity. For each parameter setting, we will sample 1000 problem instances.

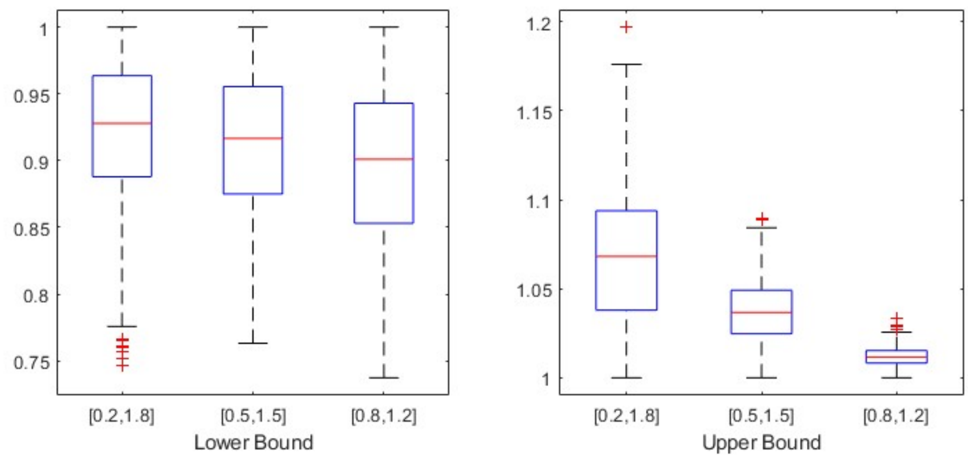

Figure 3: Effect of Variability in Utility Coefficients a_{ij}

Alternative Text Description: Effect of Variability in Utility Coefficients, as the range of utility coefficients shrinks, the upper bound gets much closer to the optimal value, while the lower bound is not affected much.

In Figure 3 we examine the effect of variability in the utility coefficients $\{a_{ij}\}$. The left panel in the figure depicts box plots of the ratios $\underline{z}/z^*$. In that panel, on the horizontal axis, there are three different ranges for the uniform distribution of the utility function coefficients. As we move to the right (while remaining in the left panel) the range for the uniform distribution gets narrower, which yields less variability in the utility coefficients. As we can see, when the range becomes narrower, the ratio $\underline{z}/z^*$ decreases very slightly. The right panel in Figure 3 contains box plots of $\bar{z}/z^*$. There, as we move to the right (while remaining in the right panel) the range for the uniform distribution narrows and the ratio $\bar{z}/z^*$ decreases substantially. Overall, this seems to be aligned with the intuition that says the total gap shrinks and the upper bound becomes tighter.

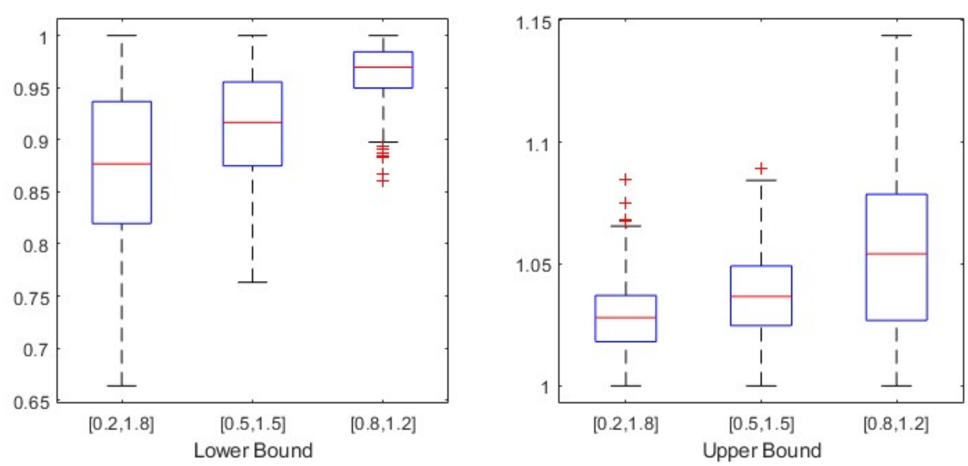

Figure 4: Effect of Variability in Weights w_i

Alternative Text Description: Effect of Variability in Weights, as the range of weights shrinks, the lower bound gets much closer to the optimal value, while the upper bound is not affected much.

Figure 4 shows the effect of variability of product weights $\{w_j\}$. Similar to the previous figure, the left panel of Figure 4 displays boxplots of $\underline{z}/z^*$ and the right panel displays boxplots of $\bar{z}/z^*$. The horizontal axes correspond to less variability in weights as we move rightward within a panel. In the left panel, we see that $\underline{z}/z^*$ is quite close to 1 when weight variability is low. In fact, if we decrease the weight variability even further than shown in the left panel, the ratio will get even closer to 1. This occurs because if the vector $\vec{w}$ is a constant multiple of $\vec{p}$, then $\bar{z} = z^*$. We have a proof of this result, but in the interest of brevity, we do not provide that here. The right panel of the figure shows that $\bar{z}/z^*$ grows slightly as the range of weights becomes smaller. This happens because it becomes more beneficial to allocate some products to customers later in the arrival sequence.

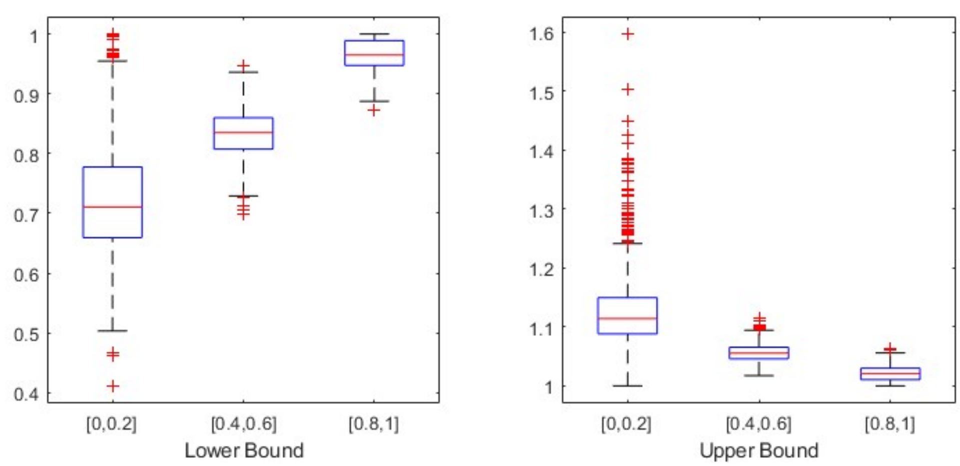

Figure 5: Effect of Variability in Capacity C

Alternative Text Description: Effect of Variability in Capacity, as the range of capacity shrinks, both the lower bound and the upper bound get much closer to the optimal value.

To close this subsection, we seek to understand the influence of capacity C. The left panel of Figure 5 displays boxplots of $\underline{z}/z^*$ and the right panel displays boxplots of $\bar{z}/z^*$. As we move to the right within a panel, the capacity increases. As we can see, the ratios both get closer to 1 as the capacity increases. That is what we expect. As the capacity grows, the related constraint will become less restrictive, which will make the heuristic closer to the optimal value. Also, a high capacity level will enable the customers late in the sequence to get some products (and more preferred products), which mitigates the advantage of the central decision maker who is always able to make favorable product allocations to those customers.

7.4 Sensitivity and Robustness

In the last portion of our numerical experiments, we will study the sensitivity and robustness of our solutions. We will start with sensitivity, which here measures how much the total utility obtained across all customers will change if we allow our instance parameters to change. We sampled 1000 instances and computed the value generated by the optimal solution and by the heuristic. Then, we changed all the parameters in the instance by random percentages within a certain tolerance (5%, 10%, and 15% in our experiments) and re-solved for the optimal solution and heuristic solution using the changed parameter values. Then, we calculated the ratio of the optimal values after and before the change. We did the same for the ratio of the heuristic values after and before the change.

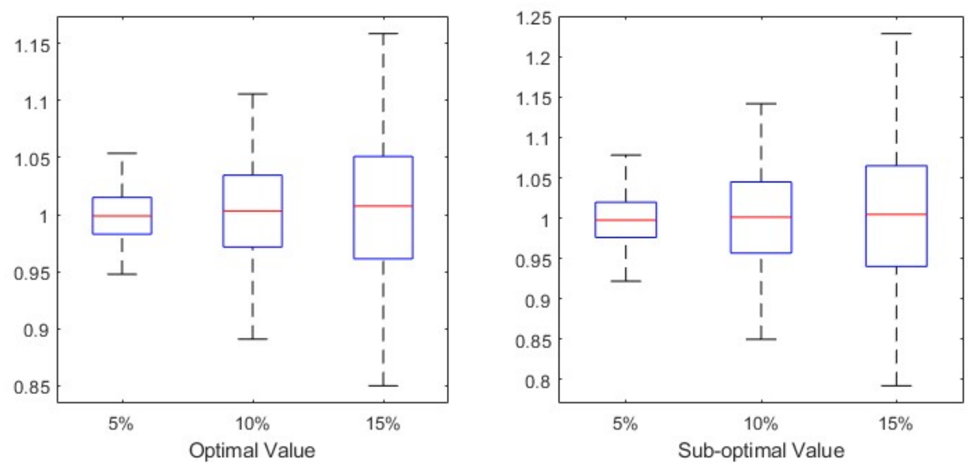

Figure 6: Sensitivity of the Optimal Value and Heuristic (Sub-optimal) Value

Alternative Text Description: Sensitivity of the Optimal Value and Heuristic (Sub-optimal) Value, as the parameters change by at most 5% (10%, 15%), the optimal value will change by at most 5% (resp., 10%, 15%) for most instances, and the heuristic value will change by at most 7.5% (resp., 15%, 22.5%) for most instances.

Boxplots of the ratios are shown in Figure 6. The left panel shows the boxplots for the ratios of optimal solutions. The right panel shows them for the heuristic. In the left panel, we see that if we allow the parameters to change by at most 5% (respectively, 10%, 15%), then the optimal value will change by at most 5% (resp., 10%, 15%) for most instances. The right panel shows the heuristic value will change by at most 7.5% (resp., 15%, 22.5%) for most instances. If we only consider the upper and lower quartiles, we see that almost 50% of the optimal and heuristic values change by less than 5% in all cases.

The sensitivity analysis that we just described compares values in different parameter settings, and requires us to find the optimal or heuristic solution for both cases. However, a more practical scenario is that the market operator first decides the inventory level based on one particular set of parameters, but then that solution is implemented for an actual problem with a different set of parameters. This can occur, for instance, if the operator does not have perfect knowledge of customer budgets and/or the parameters of customer utility functions. Hence, we also want to measure the robustness of our optimal and heuristic solutions in such settings. To do this, for each tolerance in $\{5\%, 10\%, 15\%\}$, we sample 1000 instances, and compute the optimal and heuristic solution. Then, we allow the parameters (utility coefficients $\{a_{ij}\}$, breakpoints $\{b_{ij}\}$ and budgets $\{y_i\}$) to change within the tolerance.

We then compute total utility assuming that we continue using the originally computed optimal and heuristic solutions. To further clarify, let $z(\theta, \boldsymbol{q})$ denote the total customer utility when the instance parameters are θ and the initial inventory is $\boldsymbol{q}$. Let $\boldsymbol{q}^*(\theta)$ denote the optimal initial inventory as a function θ and let $\boldsymbol{q}^{\mathrm{H}}(\theta)$ denote the heuristic initial inventory as a function θ. To evaluate the robustness of the optimal solutions we consider the ratio $z(\theta_1, \boldsymbol{q}^*(\theta_0))/z(\theta_1, \boldsymbol{q}^*(\theta_1))$ where θ_0 is the parameter value under which the system operator computes its policy, and θ_1 is the parameter value of the actual problem setting in which the policy is implemented. To evaluate the robustness of the heuristic solution we consider the ratio $z(\theta_1, \boldsymbol{q}^{\mathrm{H}}(\theta_0))/z(\theta_1, \boldsymbol{q}^{\mathrm{H}}(\theta_1))$.

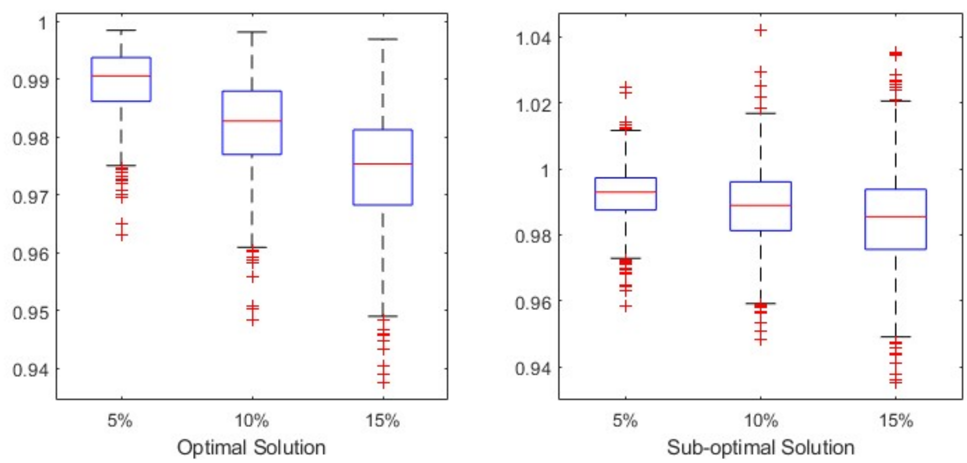

Figure 7: Robustness of Optimal Solution and Heuristic (Sub-optimal) Solutions

Alternative Text Description: Robustness of Optimal Solution and Heuristic (Sub-optimal) Solutions, as the parameters change by at most 5% (10%, 15%), both the optimal solution and the heuristic solution will get a utility loss of at most 3% (resp., 4%, 5%) for most instances.

Figure 7 displays the results of this analysis. The left panel shows boxplots of the ratio $z(\theta_1, \boldsymbol{q}^*(\theta_0))/z(\theta_1, \boldsymbol{q}^*(\theta_1))$ under three different settings of how far θ_1 can be from θ_0. There, it can seen that ratios never exceed 1. This, of course, must be the case because $\boldsymbol{q}^*(\theta_1)$ maximizes $z(\theta_1, \boldsymbol{q})$. It is interesting to note, however, that the ratios are rather close to 1, suggesting that the optimal solutions for these examples are rather robust. The right panel shows boxplots of the ratio $z(\theta_1, \boldsymbol{q}^{\mathrm{H}}(\theta_0))/z(\theta_1, \boldsymbol{q}^{\mathrm{H}}(\theta_1))$ for the heuristic. Here it is possible for the ratio to exceed 1 because $\boldsymbol{q}^{\mathrm{H}}(\theta_1)$ generally does not maximize $z(\theta_1, \boldsymbol{q})$. In most cases, the ratio is slightly below 1, suggesting that the heuristic is fairly robust in these examples.

8 Conclusions

In this research, we investigated mobile markets as a decentralized food distribution strategy in comparison to traditional centralized approaches. Our models capture the interdependencies between product assortment in mobile markets and customer purchasing behaviors. We developed analytical solution approaches to optimize inventory management. The computational experiments demonstrated not only the efficiency but also their robustness. Additionally, the case studies offer valuable insights into the fairness implications of mobile markets.

As discussed in Section 6, uncertainties are inevitable due to random customer behaviors. In future work, we will extend our models to stochastic settings. Beyond inventory decisions, factors such as routing, pricing, and logistics operations (e.g., loading and unloading) also influence customer experiences, program costs, and profitability. Therefore, we will also expand our analysis to incorporate these dimensions in future research.

Appendix

In this appendix, we explain how to transform the non-linear constraints associated with the complementary slackness conditions in (P_1) into linear constraints. For example, consider the non-linear constraint $\lambda_i(y_i - \boldsymbol{p}^\top \boldsymbol{x}_i) = 0$. Suppose that M_1 is a large positive number and $z_i \in \{0, 1\}$, $i = 1, \ldots, N$ is a collection binary variables. We can now replace the preceding non-linear constraint with the following two linear constraints

$$\lambda_i \leq M_1 z_i$$
$$y_i - \boldsymbol{p}^\top \boldsymbol{x}_i \leq M_1(1 - z_i)$$

By applying the same method to the other complementary slackness constraints, we can rewrite (P_1) from Section 4.1 as follows:

$$
\begin{aligned}
\max_{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_N} \quad & \sum_{i=1}^{N} u_i(\boldsymbol{x}_i) \qquad (P_A) \\
s.t.: \quad & \boldsymbol{w}^\top (\sum_{i=1}^{N} \boldsymbol{x}_i) \leq C \\
& \lambda_i \boldsymbol{p} + \boldsymbol{r}_i - \boldsymbol{l}_i = \nabla u_i(\boldsymbol{x}_i), \forall i = 1, \ldots, N \\
& \lambda_i \leq M_1 z_i, \forall i = 1, \ldots, N \\
& 0 \leq y_i - \boldsymbol{p}^\top \boldsymbol{x}_i \leq M_1(1 - z_i), \forall i = 1, \ldots, N \\
& \boldsymbol{r}_i \leq M_2 \boldsymbol{U}_i, \forall i = 1, \ldots, N \\
& \sum_{k=i+1}^{N} \boldsymbol{x}_k \leq M_2(\boldsymbol{1} - \boldsymbol{U}_i), \forall i = 1, \ldots, N \\
& \boldsymbol{l}_i \leq M_3 \boldsymbol{L}_i, \forall i = 1, \ldots, N \\
& 0 \leq \boldsymbol{x}_i \leq M_3(\boldsymbol{1} - \boldsymbol{L}_i), \forall i = 1, \ldots, N \\
& \lambda_i, \boldsymbol{r}_i, \boldsymbol{l}_i \geq 0, \forall i = 1, \ldots, N \\
& z_i \in \{0, 1\}, \forall i = 1, \ldots, N \\
& \boldsymbol{U}_i, \boldsymbol{L}_i \in \{0, 1\}^n, \forall i = 1, \ldots, N
\end{aligned}
$$

References

Aros-Vera, F., Marianov, V., & Mitchell, J. E. (2013). p-hub approach for the optimal park-and-ride facility location problem. *European Journal of Operational Research*, *226*(2), 277–285.

Beck, Y., Ljubić, I., & Schmidt, M. (2023). A survey on bilevel optimization under uncertainty. *European Journal of Operational Research*, *311*(2), 401–426.

Biesinger, B., Hu, B., & Raidl, G. (2016). Models and algorithms for competitive facility location problems with different customer behavior. *Annals of Mathematics and Artificial Intelligence*, *76*(1), 93–119.

Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization.* Cambridge, UK: Cambridge University Press.

Cooper, W. L. (2002). Asymptotic behavior of an allocation policy for revenue management. *Operations Research*, *50*(4), 720–727.

Davis, L. B., Sengul, I., Ivy, J. S., Brock III, L. G., & Miles, L. (2014). Scheduling food bank collections and deliveries to ensure food safety and improve access. *Socio-Economic Planning Sciences*, *48*(3), 175–188.

Dempe, S. (2002). *Foundations of bilevel programming.* Berlin, Heidelberg: Springer Science & Business Media.

Dempe, S. (2020). Bilevel optimization: theory, algorithms, applications and a bibliography. *Bilevel optimization: advances and next challenges*, 581–672.

Feeding America. (2024). Retrieved from `https://map.feedingamerica.org/county/2022/overall/minnesota`.

Gallego, G., & Van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, *45*(1), 24–41.

Geng, X., Huh, W. T., & Nagarajan, M. (2014). Sequential resource allocation with constraints: Two-customer case. *Operations Research Letters*, *42*(1), 70–75.

Hassanzadeh, P., Kreacic, E., Zeng, S., Xiao, Y., & Ganesh, S. (2023). Sequential fair resource allocation under a Markov decision process framework. In *Proceedings of the fourth acm international conference on ai in finance* (pp. 673–680).

Horning, M. L., Alver, B., Porter, L., Lenarz-Coy, S., & Kamdar, N. (2021). Food insecurity, food-related characteristics and behaviors, and fruit and vegetable intake in mobile market customers. *Appetite*, *166*, 105466.

Hsiao, B.-s., Sibeko, L., & Troy, L. M. (2019). A systematic review of mobile produce markets: facilitators and barriers to use, and associations with reported fruit and vegetable intake. *Journal of the Academy of Nutrition and Dietetics*, *119*(1), 76–97.

Hsiao, B.-s., Sibeko, L., Wicks, K., & Troy, L. M. (2018). Mobile produce market influences access to fruits and vegetables in an urban environment. *Public Health Nutrition*, *21*(7), 1332–1344.

Kasprzak, C., Lally, A., Schoonover, J., Gallicchio, D., Haynes-Maslow, L., Vermont, L., ... Leone, L. (2022). Operational challenges that may affect implementation of evidence-based mobile market interventions. *BMC Public Health*, *22*(1), 776.

Kasprzak, C., Schoonover, J., Gallicchio, D., Haynes-Maslow, L., Vermont, L., Ammerman, A., ... Leone, L. (2021). Using common practices to establish a framework for mobile produce markets in the United States. *Journal of Agriculture, Food Systems, and Community Development*, *10*(4), 73–84.

Kleinert, T., Labbé, M., Ljubić, I., & Schmidt, M. (2021, January). A survey on mixed-integer programming techniques in bilevel optimization. *EURO Journal on Computational Optimization*, *9*, 100007.

Koutsoupias, E., & Papadimitriou, C. (1999). Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science* (pp. 404–413).

Liang, J., & Lyu, G. (2022). From efficiency to fairness: Design of allocation rules for food bank operations. In *2022 Winter Simulation Conference* (pp. 1557–1568).

Lien, R. W., Iravani, S. M., & Smilowitz, K. R. (2014). Sequential resource allocation for nonprofit operations. *Operations Research*, *62*(2), 301–317.

Mobile Market Coalition. (2022). *2022 mobile market coalition impact report.* Retrieved from `https://www.mobilemarketcoalition.org/publications--reports.html` (Accessed: February 8, 2025)

Orgut, I. S., Ivy, J. S., Uzsoy, R., & Hale, C. (2018). Robust optimization approaches for the equitable and effective distribution of donated food. *European Journal of Operational Research*, *269*(2), 516–531.

Papadimitriou, C. (2001). Algorithms, games, and the internet. In *Proceedings of the Thirty-Third Annual ACM symposium on Theory of Computing* (pp. 749–753).

Robinson, J. A., Weissman, E., Adair, S., Potteiger, M., & Villanueva, J. (2016). An oasis in the desert? the benefits and constraints of mobile markets operating in Syracuse, New York food deserts. *Agriculture and Human Values*, *33*, 877–893.

Weissman, E., Robinson, J., & Cecio, W. (2020). The promise and pitfalls of mobile markets: An exploratory survey of mobile food retailers in the United States and Canada. *Agriculture and Human Values*, *37*, 895–906.

Wishon, C., & Villalobos, J. R. (2016). Alleviating food disparities with mobile retailers: Dissecting the problem from an OR perspective. *Computers & Industrial Engineering*, *91*, 154–164.

Zepeda, L., Reznickova, A., & Lohr, L. (2014). Overcoming challenges to effectiveness of mobile markets in US food deserts. *Appetite*, *79*, 58–67.

Zhang, Y., Berman, O., & Verter, V. (2012). The impact of client choice on preventive healthcare facility network design. *OR Spectrum*, *34*(2), 349–370.