Les Cahiers du GERAD ISSN: 0711–2440

Beyond one-size-fits-all: personalized delivery and fulfillment optimization

Q. Zhou, M. Gümüş, S. Miao

G–2025–06

January 2025

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Citation suggérée : Q. Zhou, M. Gümüş, S. Miao (Janvier 2025). Beyond one-size-fits-all: personalized delivery and fulfillment optimization, Rapport technique, Les Cahiers du GERAD G–2025–06, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2025-06) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Suggested citation: Q. Zhou, M. Gümüş, S. Miao (January 2025). Beyond one-size-fits-all: personalized delivery and fulfillment optimization, Technical report, Les Cahiers du GERAD G–2025–06, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2025-06) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2025
– Bibliothèque et Archives Canada, 2025

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2025
– Library and Archives Canada, 2025

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Beyond one-size-fits-all: personalized delivery and fulfillment optimization

Quan Zhou [a]

Mehmet Gümüş [a, b]

Sentao Miao [c]

[a] *Desautels Faculty of Management, McGill University, Montréal (Qc), Canada, H3A 1G5*

[b] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[c] *Leeds School of Business, University of Colorado Boulder, Boulder, Colorado, USA, 80309*

quan.zhou3@mail.mcgill.ca
mehmet.gumus@mcgill.ca
sentao.miao@colorado.edu

January 2025
Les Cahiers du GERAD
G–2025–06

Copyright © 2025 Zhou, Gümüş, Miao

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : Motivated by our collaboration with an online platform operating in North America, we explore the joint optimization of the order fulfillment process with personalized delivery options in the context of e-commerce. Customers can choose from personalized fulfillment options to proceed with the purchase or leave with no purchase. The retailer periodically makes fulfillment decisions and relies on multiple logistic providers to perform the fulfillment operations. We model customer behavior with a general discrete choice model and formulate the joint optimization as a stochastic dynamic program. We propose a tractable deterministic approximation and develop a computationally efficient heuristic with a provable performance guarantee. We also extend the proposed heuristic to scenarios when customer behaviors are more complex and affected by fulfillment speed, cost, and order value. Using real datasets collected from our industrial partner, we demonstrate the value of personalizing fulfillment options for the customers and jointly optimizing the options with fulfillment assignments. Our results show that demand management via personalized fulfillment options is prominent when customers favor quicker fulfillment and when the fulfillment capacity is limited. However, an optimized fulfillment operation becomes more critical when customers are more willing to wait.

Keywords: Fulfillment optimization, personalized delivery, e-commerce

Résumé : Motivés par notre collaboration avec une plateforme en ligne opérant en Amérique du Nord, nous explorons l'optimisation conjointe du processus d'exécution des commandes avec des options de livraison personnalisées dans le contexte du commerce électronique. Les clients peuvent choisir parmi des options d'exécution personnalisées de procéder à l'achat ou de repartir sans avoir acheté. Le détaillant prend périodiquement des décisions d'exécution et fait appel à plusieurs fournisseurs logistiques pour effectuer les opérations d'exécution. Nous modélisons le comportement du client à l'aide d'un modèle général de choix discret et formulons l'optimisation conjointe sous la forme d'un programme dynamique stochastique. Nous proposons une approximation déterministe traçable et développons une heuristique efficace en termes de calcul avec une garantie de performance prouvée. Nous étendons également l'heuristique proposée à des scénarios où les comportements des clients sont plus complexes et affectés par la vitesse d'exécution, le coût et la valeur de la commande. En utilisant des ensembles de données réelles collectées auprès de notre partenaire industriel, nous démontrons la valeur de la personnalisation des options d'exécution pour les clients et de l'optimisation conjointe des options avec les affectations d'exécution. Nos résultats montrent que la gestion de la demande par le biais d'options d'exécution personnalisées est importante lorsque les clients préfèrent une exécution plus rapide et lorsque la capacité d'exécution est limitée. Cependant, une opération d'exécution optimisée devient plus critique lorsque les clients sont plus disposés à attendre.

Mots clés : Optimisation de l'exécution, livraison personnalisée, commerce électronique

Acknowledgements: The authors thank GERAD for funding the presentation of this work at 2024 INFORMS Annual Meeting.

1 Introduction

Customer shopping behaviors have been significantly reshaped since the COVID-19 pandemic. There is a persistent trend of customers favoring online shopping, leading to e-commerce accounting for nearly 19% of global retail sales, reaching approximately $5.2 trillion (eMarketer 2023). This trend also catalyzes the adoption of online shopping in previously unexpected sectors, such as grocery shopping. A recent industry report valued the global online grocery market at $50.28 billion in 2022, estimating it to reach a value of $305.13 billion by 2030, with an annual growth rate of 26.8% throughout the forecast period (Grand View Research 2022).

As customers increasingly gravitate towards online shopping, they are drawn to retailers providing personalized delivery services. While fast and cheap fulfillment remains preferable for most consumers, a rising number of shoppers now express heterogeneous preferences for receiving their orders. Speed, cost, convenience, and eco-friendliness are among the key factors that impact customer preference for fulfillment options (Digital Commerce 360 2023). Thus, offering diversified options such as standard shipping, express shipping, attended home delivery, and out-of-home delivery has become increasingly popular among practitioners. However, although various options can help increase conversions, they complicate the fulfillment operations and require seamless integration of multiple fulfillment methods to cater to the varied needs of the customers (Wolfe 2024). As a result, online retailers are urged to pair diversified customer expectations with complex and expensive fulfillment operations. It is crucial for online retailers to ensure that they not only meet these expectations but also optimize the delivery process. The primary objective of this paper is to address these challenges by jointly optimizing the personalized fulfillment options and the fulfillment process within an e-commerce context.

The setting of this paper is motivated by our collaboration with one of the online grocery platforms in North America. This company enables online shoppers to purchase groceries from various suppliers, offering them personalized fulfillment options at checkout. As depicted in Figure 1, customers are presented with a three-level choice. They have to choose either home or out-of-home delivery options, then specify a delivery date and (if needed) a pickup location such as a convenience store, a restaurant, or a warehouse. In addition, they also need to indicate a time slot for the pickup. To facilitate these services, the company partners with third-party trucking companies and local convenience stores, restaurants, etc.

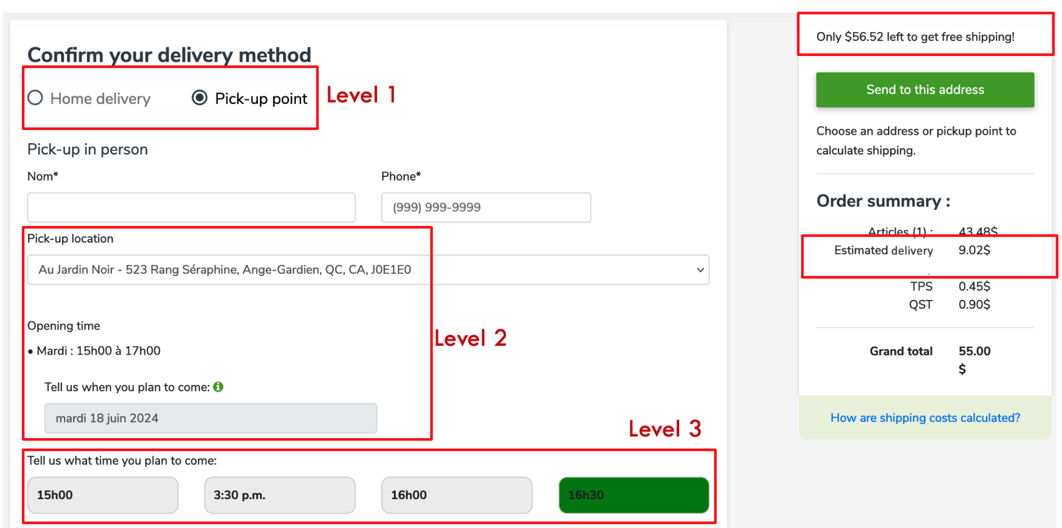

Figure 1: The personalized delivery options page from our industrial partner

The core challenge faced by our partner involves efficiently managing these varied fulfillment options to optimize customer satisfaction and operational efficiency. They encounter difficulties in predicting customer preferences for different fulfillment options and coordinating the logistic capacities of various partners. The complexity is further amplified by fluctuating demand patterns and the need to minimize costs while maintaining service quality.

To address these issues, we adopt a general discrete choice model, which allows us to incorporate various customer behaviors when they shop online. Additionally, we develop a stochastic model that jointly optimizes the display of fulfillment options and the operations of fulfilling orders. This model aims to balance the trade-offs between customer satisfaction, characterized by the availability of preferred fulfillment options, and operational efficiency, which involves minimizing logistics costs and optimizing resource allocation. The integration of these models provides a comprehensive approach to enhance the decision-making process in e-commerce fulfillment, contributing to both theoretical understanding and practical improvements in e-commerce supply chain management.

1.1 Main contributions

We consider a multi-period joint optimization of personalized fulfillment options and assignments for companies involved in delivery and fulfillment services. Such companies include e-commerce platforms like Amazon or delivery services like Instacart, Shipt, and our industrial partner, specializing in local grocery delivery. The decisions include what fulfillment options to show customers and how to fulfill orders, aiming to maximize total expected profits. To the best of our knowledge, this is among the first studies in the literature to address the dynamic aspect of fulfillment flexibility and assignment problems. This focus is notable, given the significant impact of the problem in enabling companies to satisfy customers and maximize profits in e-commerce contexts. Our results and contributions are summarized as follows:

1. We formulate the problem as a joint stochastic program incorporating a general discrete choice model, where each fulfillment option specifies the latest fulfillment time and a fixed surcharge. We propose a tractable deterministic (fluid) approximation, replacing random variables with their expected values and constraining fulfillment option display decisions to be static and randomized. We show that this deterministic approximation serves as an upper bound for our original problem. Furthermore, focusing on the deterministic setting, we show that an optimal list of fulfillment options must be *efficient* in that it maximizes expected revenue while optimizing fulfillment resource utilization. However, this result cannot be extended to the original stochastic setting, which motivates to adjust an option list dynamically.

2. We develop a two-phase algorithm called the "Economic option, Rounded and Threshold fulfillment" policy (**ERT**) as depicted in Figure 2. In the initial (offline) phase, **ERT** generates a set of efficient assortments and solves the deterministic relaxation to form referenced assortment and fulfillment decisions. In the subsequent (online) phase, **ERT** iteratively engages in personalization and fulfillment stages throughout the planning cycle. During the delivery personalization stage, the algorithm starts with efficient assortments. It dynamically excludes *economically nonviable* fulfillment options based on customer location and the current state of orders to be fulfilled. In the fulfillment stage, the algorithm consolidates all orders. Then, it allocates them to the appropriate warehouses according to a threshold-based policy determined based on a fulfillment cost function.

3. We provide a performance guarantee for the **ERT** algorithm. Our analysis shows that the optimality gap of **ERT** increases at a rate of $\mathcal{O}(TMK\sqrt{K\Lambda N})$, with T representing the number of fulfillment cycles, M the number of warehouses, N the number of customer locations, K the maximum length of a fulfillment window offered to customers, and Λ the largest customer arrival rate. Specifically, the *sub-linear* growth in N and Λ suggests that **ERT** achieves asymptotic optimality as the number of customer locations or the number of arriving customers per location increases, a scenario commonly encountered by numerous online retailers, including our industrial partner.

4. Beyond these methodological contributions, our study generates managerial insights for an effective joint strategy in managing demand (through personalized fulfillment options) and capacity

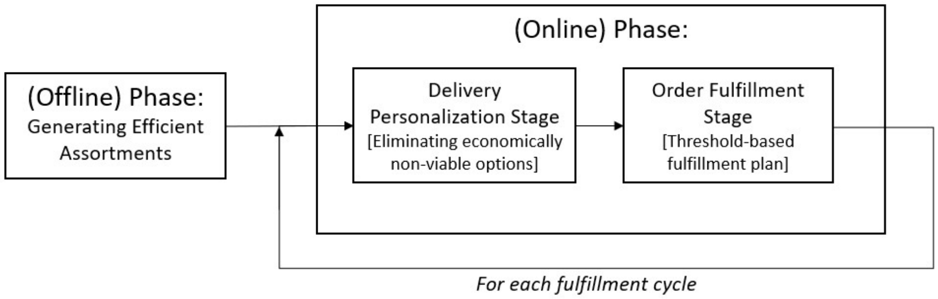

Figure 2: The overall diagram of our proposed (ERT) algorithm

(via fulfillment assignments), substantiated with real datasets collected from our industrial partner. Our numerical study indicates that while optimizing personalized fulfillment options and assignments is crucial, their roles vary in different scenarios. More specifically, when customers favor quicker fulfillment options, demand management is prominent. Conversely, capacity management becomes more critical when customers prefer more cost-effective (but slower) fulfillment options. Furthermore, demand management is more pivotal when fulfillment capacity is limited, but as capacity increases, the focus shifts toward capacity management.

1.2 Organization

The remainder of this paper is organized as follows. In §2, we review the related literature. We present the detailed model formulation in §3 and introduce choice-based deterministic relaxation in §4. Utilizing the solution to the relaxation, we develop a heuristic and show its theoretical performance guarantee in §5. We conduct extensive numerical studies of our proposed heuristic policy and present managerial insights in §6. In §7, we extend the formulation to the scenario where customers are sensitive to both the speed and the cost of a fulfillment option and demonstrate the consistency and robustness of the proposed heuristic with additional numerical studies. §8 concludes the paper.

2 Literature review

Our paper falls within the general theme of fulfillment optimization (refer to reviews by Acimovic and Farias 2019 and Qi et al. 2020). However, to the best of our knowledge, ours is among the first to explore fulfillment optimization with personalized delivery options. There are three streams of research closely related to our paper.

E-commerce order fulfillment. The first stream investigates e-commerce fulfillment. Prior work in this stream has examined order consolidation (Xu et al. 2009), single- and multi-item fulfillment under limited inventory conditions (Acimovic and Graves 2015, Jasin and Sinha 2015), and omnichannel retailing (Andrews et al. 2019).

Two specific sub-streams are particularly relevant to our study. The first involves the integration of pricing and fulfillment issues. In this context, Lei et al. (2018) and Harsha et al. (2019) examined the joint pricing and fulfillment problem, while Lei et al. (2022) expanded the scope to include pricing, display, and fulfillment. Notably, Lei et al. (2018) approached the problem as a stochastic program, devising a two-tier approximation method. Initially, they utilized a deterministic approximation, which, due to its complexity and non-linearity, was further approximated by a solvable linear program through discrete price adjustments. This facilitated the development of heuristic policies based on the model's outputs. Our paper differs from existing literature as we manage demand realization more subtly to take advantage of strategically postponing fulfillment.

The second sub-stream focuses on the strategic benefits of delaying fulfillment decisions. For instance, Mahar and Wright (2009) investigated a quasi-dynamic assignment policy in an omnichannel retail setting with a linear cost structure. Wei et al. (2021) analyzed optimal order consolidation policies, providing heuristic solutions based on structural properties. Finally, Xie et al. (2023) studied online resource allocation with delayed decision-making to leverage real-time demand data. They demonstrated that the performance gap between their online algorithm with delays and the offline optimal policy reduces exponentially as the delay period increases. We add to this literature stream by incorporating personalized delivery options.

Choice-based revenue management. The second stream is related to revenue management (RM) and assortment optimization. Comprehensive reviews are available in Kök et al. (2009) and Gallego and Topaloglu (2019). Our focus is on the subset of research utilizing choice-based demand models in RM, as our paper explores demand realization through such models. Talluri and Van Ryzin (2004) examined a single-resource RM issue employing a general discrete choice model. They introduced two pivotal concepts: efficient sets and nested-by-fare order. By formulating the problem as a dynamic program (DP), they demonstrated that optimal assortments are efficient sets. Furthermore, they established that optimal assortments adhere to a nested-by-fare order under several specific choice models. Liu and Van Ryzin (2008) expanded the efficient set concept to network revenue management (NRM) within a general discrete choice model framework. They proved that only efficient sets form optimal assortments in the choice-based deterministic relaxation formulation. Unlike in the single-resource scenario, this principle does not universally apply to our DP formulation. Specific choice models provide insights into efficient sets, as discussed in Gallego and Li (2017) and Cao et al. (2022). Our paper focuses on developing an implementable heuristic with a provable performance guarantee based on solving the deterministic relaxation with efficient sets.

Flexible products. The third stream is related to network revenue management (NRM) with flexible products. Gallego and Phillips (2004) describe a *flexible product* as a set of alternatives offered to a market, where the seller later assigns a specific alternative to the purchaser. In our context, an e-commerce company accepts orders encompassing various fulfillment options (akin to products) and later fulfills these orders from a selection of warehouses (comparable to alternatives). Gallego et al. (2004) explored NRM with flexible products under both independent and dependent demand scenarios, employing a general choice model. They demonstrated the asymptotic optimality of deterministic relaxation for both cases. Notably, they used a column generation algorithm to resolve the deterministic relaxation for demand influenced by customer choices. Gönsch et al. (2014) extended the work and introduced heuristic policies derived from linear programming approximations. Subsequently, Cheung and Simchi-Levi (2016) concentrated on an approximate yet efficient solution for the choice-based deterministic relaxation, proposing a Potential-based algorithm with proven efficiency: a polynomial-time method for solving the deterministic relaxation to any desired level of accuracy.

More recently, Ma et al. (2020) developed an approximate dynamic programming (DP) algorithm for solving the choice-based network RM problem. They utilized an *availability tracking basis* for value function approximation, showing that their heuristic's total expected revenue reaches at least a $1/(1+L)$ fraction of the optimal total expected revenue, where L represents the maximum number of resources a product consumes. Zhu and Topaloglu (2023) expanded this methodology to flexible product scenarios. The primary challenge in such settings is the availability of remaining capacity information. To overcome this, they introduced an auxiliary variable to track remaining resource capacities, optimistically approximating the value function with this variable. They established that the approximate DP method retains the same constant performance bound even in the context of flexible products. Our paper differs from the previous research in this stream by proposing an asymptotically optimal heuristic based on deterministic relaxation that can be solved efficiently with a commercial convex optimization solver.

3 The model

We consider an online retailing company, referred to as the DM, that aims to improve its fulfillment system over a finite horizon. In particular, the DM utilizes a combination of (both in-house and third-party) logistic carriers to carry out fulfillment operations. To cover the total fulfillment cost, as is typical in the industry, the DM relies on two revenue sources. The first one is a fixed percentage of the basket value (i.e., the *fulfillment reserve*), and the second is an additional fare customers pay for shipping (i.e., the *fulfillment surcharge*). The DM's objective is to balance the revenue and cost terms associated with the delivery process.

We adopt the following notation. The DM sells products to customers in N distinct locations, each of which can be conceptualized as a combination of a specific geographic location and a specific combination of products that generates some fulfillment reserve. The DM aggregates orders and ships them in batches from M warehouses over a fixed period called a *fulfillment cycle*. The planning horizon is divided into T fulfillment cycles. Each order has a specific time frame, referred to as *fulfillment window*, spanning k fulfillment cycles (where $k = 1, \ldots, K$), within which the order must be fulfilled. Customers are offered a menu of fulfillment options when they check out their shipping carts. Each option specifies a fulfillment window and a fulfillment surcharge. Customers may proceed with the purchase or leave without the purchase. The DM must assign the submitted orders to one of its M warehouses and deliver them within the specified fulfillment window. *To maximize the expected profit associated with the fulfillment process, the DM decides which fulfillment options to show to each customer and when and from which warehouse to fulfill each order within its fulfillment window.*

Throughout this paper, we define $[n]$ as the set $\{1, 2, \cdots, n\}$ and denote $\mathbb{R}_+$ denote $[0, +\infty)$. We use $i \in [M]$ to denote the i-th warehouse, $j \in [N]$ for the j-th location, $\tau \in [T]$ for the τ-th fulfillment cycle, and $k \in [K]$ to denote the remaining cycles in the fulfillment window.

3.1 Fulfillment options

A fulfillment option consists of a guaranteed latest fulfillment window and a surcharge. We initially impose the following assumption regarding different fulfillment options for ease of exposition. This assumption will be relaxed in §7.

Assumption 1. There are K fulfillment options in total, $\mathcal{S} = \{(1, r_1), (2, r_2), \cdots, (K, r_K)\}$.

In alignment with the standard notation used in the assortment literature, we denote an *assortment* as $S \subseteq \mathcal{S}$, representing a list of specific fulfillment options. We utilize the notation $k \in S$ to indicate a specific option within an assortment, characterized by a fulfillment window of k and a surcharge of r_k. This notation should not lead to confusion within the context of this paper: when we mention $k \in [K]$, it refers to a fulfillment window, whereas $k \in S$ denotes a fulfillment option. The set of all possible assortments is represented by $\mathcal{N}$, and $|\mathcal{N}| = 2^K$ under Assumption 1.

A customer decides whether to proceed with a purchase by selecting from the available fulfillment options based on a known discrete choice model denoted as $\{\pi_k(\cdot) : k \in \mathcal{S}\}$. Specifically, if the DM presents an assortment $S \in \mathcal{N}$, the customer will choose a fulfillment option $k \in S$ and purchase with a probability $\pi_k(S)$. Additionally, the customer may opt not to proceed with the checkout with a probability $\pi_\varnothing(S)$. This scenario, known as *cart abandonment*, occurs when the customer is dissatisfied with the provided fulfillment options.

3.2 Demand

We assume that customers arrive from location j (or customer j for short) during a fulfillment cycle τ follow a Poisson process with a time-dependent arrival rate of $\lambda_{j,\tau}(t)$, where t represents the elapsed time since the start of the current fulfillment cycle. This arrival process is independent across locations and need not be homogeneous in τ and t. Without loss of generality, we normalize the duration of each fulfillment cycle to the interval $[0, 1]$, with $t = 0$ marking the beginning of cycle τ and $t = 1$ indicating its closure (just before the fulfillment decision is made). We denote $\tilde{D}_{j,\tau}(t)$ the cumulative arrivals of customer j by time t during cycle τ. Furthermore, we let $\tilde{D}_{j,\tau} := \tilde{D}_{j,\tau}(1)$ represent the total number of customers from location j arrive during cycle τ. Consequently, $\tilde{D}_{j,\tau}$ follows a Poisson distribution with parameter $\Lambda_{j,\tau} := \int_0^1 \lambda_{j,\tau}(t)dt$. That is,

$$\mathbb{P}\{\tilde{D}_{j,\tau} = n\} = \frac{(\Lambda_{j,\tau})^n}{n!} e^{-\Lambda_{j,\tau}}$$

When a customer arrives, the probability that she will choose fulfillment option k, given an assortment S is offered, is $\pi_k(S)$. We determine the assortment based on a decision rule $S_{j,\tau}(t)$. This decision rule may vary depending on the customer location j, the fulfillment cycle τ, and the elapsed time t within the cycle. We denote $D_{k,j,\tau}(t)$ as the cumulative number of customers from location j who choose fulfillment option k by the elapsed time t in cycle τ, representing the *realized* demand.

It is important to note that while the cumulative arrival process $\tilde{D}_{j,\tau}(t)$ follows a Poisson distribution, the realized demand $D_{k,j,\tau}(t)$ for each fulfillment option $k \in [K]$ may not necessarily follow a Poisson distribution. This deviation arises because even though $D_{k,j,\tau}(t)$ is derived from $\tilde{D}_{j,\tau}(t)$, it is influenced by the selected fulfillment option, which in turn is influenced by the decision rule $S_{j,\tau}(t)$. Since the decision rule $S_{j,\tau}(t)$ is dynamic and may depend on the history of realized demand, $D_{k,j,\tau}(t)$ also becomes dependent on past events, losing the memoryless property, which is characteristic of a Poisson process. This intricate interplay between assortment rules and realized demand poses significant computational challenges. To tackle these complexities, in the next section, we propose employing a heuristic approach based on deterministic approximation, which captures the impact of decision rules on the demand realizations in a tractable fashion.

3.3 System dynamics

We define the "pre-assortment" state, denoted as $x_{k,j,\tau}(t)$, as the cumulative unfulfilled orders received from location j that have k remaining cycles left in the fulfillment window at the beginning of time t during fulfillment cycle τ. Similarly, the "post-assortment" state, denoted as $y_{k,j,\tau}(t)$, represents the cumulative unfulfilled orders at the end of time t. When a new customer arrives at time t and an assortment S is displayed, if customer j and selects fulfillment option k, then $y_{k,j,\tau}(t) = x_{k,j,\tau}(t) + 1$. Otherwise, $y_{k,j,\tau}(t) = x_{k,j,\tau}(t)$. Additionally, we define $x_{k,j,\tau} := x_{k,j,\tau}(0)$ as the pre-assortment state of fulfillment cycle τ and $y_{k,j,\tau} := y_{k,j,\tau}(1)$ as the post-assortment state of that cycle. Thus, we have the following relationship

$$y_{k,j,\tau} = x_{k,j,\tau} + D_{k,j,\tau}$$

We point out that all the states $x_{k,j,\tau}$ and $y_{k,j,\tau}$ are subject to two sources of randomness: the arrivals of customers are stochastic, and the fulfillment options they choose, conditioned on some assortments offered, are also stochastic.

At the end of the fulfillment cycle τ, the DM makes a fulfillment decision denoted by $u_\tau = (u_{k,j,i,\tau} \geq 0 : k \in [K], j \in [N], i \in [M])$. In particular, $u_{k,j,i,\tau}$ corresponds to the number of orders shipped by fulfillment cycle τ from warehouse i to demand location j to satisfy the unfulfilled orders with k cycles left in fulfillment window. The DM cannot ship products in advance to locations for potential future demand realizations. To ensure that, we impose the following feasibility constraints for all $k \in [K]$, $j \in [N]$, and $\tau \in [T]$:

$$\sum_{i \in [M]} u_{k,j,i,\tau} \leq y_{k,j,\tau}$$

Next, we define the state update equations. Following the fulfillment decision u_τ, unfulfilled orders remain in the system with their remaining cycles in the fulfillment window reduced by one. Among these unfulfilled orders, if any have only one fulfillment cycle remaining (i.e., $y_{1,j,\tau} - \sum_{i \in [M]} u_{1,j,i,\tau}$), they exit the system and incur penalty costs. The system dynamics can be expressed as follows:

$$x_{k,j,\tau+1} = \begin{cases} y_{k+1,j,\tau} - \sum_{i\in[M]} u_{k+1,j,i,\tau} & k < K \\ 0 & k = K \end{cases}$$

3.4 Revenue and cost

Our focus on optimizing the delivery system involves revenue and cost terms associated with the delivery process. The DM offsets the delivery costs by collecting income from two revenue sources for each order. The first source is the fulfillment reserve, a fixed percentage of the order's basket value. Let p_j represent the basket size for a customer at demand location j. An order from that customer has a κp_j reserve to cover the fulfillment cost, where κ specifies the fixed percentage. The second source is the shipping surcharge r_k, which varies based on the delivery option k chosen by the customer. To sum, we denote $r_{k,j}$ as the revenue collected by the DM to cover the cost of delivery and define it as follows: $r_{k,j} := r_k + \kappa p_j$ where k specifies the fulfillment option and j specifies the customer location. This two-part strategy is common in the industry, where the first term is derived from the displayed price, and the second term corresponds to a separate shipping cost. Finally, the DM does not generate revenue if the customer abandons the cart. Therefore, the expected revenue generated during the fulfillment cycle τ can be calculated as follows:

$$\mathbb{E}\left\{\sum_{k,j} r_{k,j} D_{k,j,\tau}\right\}$$

Next, we consider the cost term. Recall that our industry partner utilizes a variety of third-party carriers to fulfill the orders. Each carrier commits a unique capacity in terms of weight and charges a flat rate for each fulfilled order. Initially, our partner prioritizes using a more cost-effective carrier. However, if the logistics capacity is insufficient to meet demand from the cheaper carrier, our partner turns to a carrier with a higher unit cost. This results in a piecewise linear function representing the lower envelope of the total cost curve faced by the retailer, with each breakpoint reflecting cumulative capacity. To model this setup, we assume that the fulfillment cost function is convex and increasing in the total number of products shipped from each warehouse $i \in [M]$. To summarize, we define $C_i : \mathbb{R}_+ \to \mathbb{R}_+$ to denote the fulfillment cost function and make the following assumption:

Assumption 2. $C_i(z)$ is convex and increasing in z and its derivative $C_i'(z) < \overline{C}$ for all $z \in \mathbb{R}_+$, where $0 < \overline{C} < \infty$.

The fulfillment cost for warehouse i, given its fulfillment vector $u_{i,\tau}$, can be obtained as

$$C_i\left(\sum_{k\in[K]}\sum_{j\in[N]} w_j u_{k,j,i,\tau}\right) = C_i\left(w^T u_{i,\tau}\right)$$

where w_j is the billable weight of the order associated with location j, and $w = (w_{k,j} : \forall\, k \in [K], j \in [N])$, $w_{k,j} = w_j$, is the vector of billable weights. We define $\overline{w} := \max_j w_j$ as the largest billable weight.

Note that the assumption $\overline{C} < \infty$ in Assumption 2 is made primarily for the sake of technical analysis. It can be relaxed because the marginal cost can always be capped by the maximum penalty cost $\max_j b_j$ without loss of generality.

Finally, if the DM cannot fulfill the order from location j within its fulfillment period k using the committed logistic capacity, it incurs a unit penalty b_j. This penalty can be considered the shipping rate of an expedited service that can fulfill an order from j with a single fulfillment window. We assume $C_i'(0) < \max_j b_j$ to avoid triviality, which implies there exists a third-party carrier whose shipping rate is lower than the expedited shipping service. Using the state and action defined above, we can calculate the total number of outstanding orders from location j in cycle τ fulfilled with the expedited service by $y_{1,j,\tau} - \sum_{i\in[M]} u_{1,j,i,\tau}$. Multiplying this by location-specific unit penalty cost b_j and summing over all locations $j \in [N]$, we obtain the total penalty cost as follows: $\sum_{j\in[N]} b_j(y_{1,j,\tau} - \sum_{i=1}^{M} u_{1,j,i,\tau})$.

3.5 Joint optimization

The DM's objective is to maximize its expected profit by deciding which fulfillment options to offer at any time and how to fulfill orders at the end of each fulfillment cycle. We can formulate this problem as a joint stochastic optimization called **OPT**.

$$
\begin{aligned}
V^*(x) = \max_{S,u} \mathbb{E}&\left\{\sum_{\tau\in[T]}\sum_{j\in[N]}\sum_{k\in[K]} r_{k,j} D_{k,j,\tau}\right\} \\
&- \mathbb{E}\left\{\sum_{\tau\in[T]}\left[\sum_{i\in[M]} C_i(w^\top u_{i,\tau}) + \sum_{j\in[N]} b_j\left(y_{1,j,\tau} - \sum_{i\in[M]} u_{1,j,i,\tau}\right)\right]\right. \\
&\left. + \sum_{k\in[K]}\sum_{j\in[N]} b_j x_{k,j,T+1}\right\} \\
\text{s.t.}\quad & x_{k,j,1} = x_{k,j}, && \forall\, k,j && \text{(1a)} \\
& y_{k,j,\tau} = x_{k,j,\tau} + D_{k,j,\tau}, && \forall\, k,j,\tau && \text{(1b)} \\
& x_{k,j,\tau+1} = y_{k+1,j,\tau} - \sum_i u_{k+1,j,i,\tau}, && \forall\, k<K,j,\tau && \text{(1c)} \\
& x_{K,j,\tau+1} = 0, && \forall\, j,\tau && \text{(1d)} \\
& \sum_i u_{k,j,i,\tau} \le y_{k,j,\tau}, && \forall\, k,j,\tau && \text{(1e)} \\
& u_{k,j,i,\tau} \ge 0 && \forall\, k,j,i,\tau && \text{(1f)} \\
& S_{j,\tau}(t) \in \mathcal{N}, && \forall\, j,\tau,t\in[0,1] && \text{(1g)}
\end{aligned}
$$

In **OPT**, the first part of the objective function is the expected revenue. The second part is the expected costs, which consist of three components: the first one is the fulfillment cost, the second one is the penalty cost, and the third one is the terminal cost, meaning all unfulfilled orders at the end of the horizon are handled with expedited shipping services. For the constraints, (1a) specifies the initial condition, (1b) to (1d) describe the system dynamics, (1e) and (1f) requires fulfillment decisions to be feasible, and (1g) ensures the assortment must be selected from the possible fulfillment option set. We point out that the random variable $D_{k,j,\tau}$ depends on $S_{j,\tau}(t)$. All constraints must hold almost surely.

Theoretically, **OPT** can be solved using dynamic programming (DP). Formally, let $V_\tau(x_\tau)$ denote the optimal value function at the beginning of the fulfillment cycle τ given the unfulfilled orders x_τ. Define G_τ as the profit function of fulfillment cycle τ, which is formulated as follows:

$$
G_\tau(x_\tau, u_\tau, S_\tau) = \sum_{k,j} r_{k,j} D_{k,j,\tau} - \left(\sum_i C_i(w^\top u_{i,\tau}) + \sum_j b_j\left(x_{1,j,\tau} + D_{1,j,\tau} - \sum_i u_{1,j,i,\tau}\right)\right)
$$

where $D_{k,j,\tau}$ is a random variable that depends on the fulfillment options offered S_τ. The DP formulation can be defined as follows:

$$
V_\tau(x_\tau) = \max_{S_\tau, u_\tau} \mathbb{E}\left\{G_\tau(x_\tau, u_\tau, S_\tau) + V_{\tau+1}(x_{\tau+1})\right\}
$$

with the boundary condition $V_{T+1}(x) = -\sum_{k,j} b_j x_{k,j}$.

The above DP is computationally intractable. First, the distribution of $D_{k,j,\tau}$ remains unknown, as it is contingent on the assortment decisions $S_{j,\tau}(t)$. Even if we restrict our attention to finite and pre-determined assortment decisions, potentially enabling us to identify the family of distributions of $D_{k,j,\tau}$, solving above DP is still computationally intractable due to the curse of dimensionality. Consequently, our emphasis shifts towards devising computationally tractable heuristic policies and demonstrating provably effective performance.

4 Choice-based deterministic approximation

To serve our basis for computationally tractable heuristics, we define deterministic programming as follows: There will be $\Lambda_{j,\tau}$ customers come to the system from location j in cycle τ. For each arrived customer, if the DM offers an assortment S, and a deterministic and fractional proportion of $\pi_k(S)$ of the customer's order is fulfilled using option k, while a fraction $\pi_\varnothing(S)$ is not fulfilled due to cart abandonment. Thus, the DM collects only a partial reward $\sum_{k=1}^{K} r_{k,j}\pi_k(S)$. We consider a strategy such that during cycle τ, we offer assortment S to $\gamma_{j,\tau}(S)$ proportion of customers from location j. Under this strategy, the realized demand indexed with k, j, τ can be calculated as $\sum_S \Lambda_{j,\tau}\gamma_{j,\tau}(S)\pi_k(S)$ The deterministic programming can then be formulated as follows, which we refer to as **DET**.

$$
\begin{aligned}
V^D(x) = \max_{\gamma,\bar{u}} \; & \sum_\tau \sum_j \sum_k \sum_S r_{k,j}\Lambda_{j,\tau}\gamma_{j,\tau}(S)\pi_k(S) \\
& - \sum_\tau \left(\sum_i C_i(w^\top \bar{u}_{i,\tau}) + \sum_j b_j \left(\bar{y}_{1,j,\tau} - \sum_i \bar{u}_{1,j,i,\tau} \right) \right) \\
& - \sum_{k,j} b_j \bar{x}_{k,j,T+1}
\end{aligned}
$$

$$
\begin{aligned}
\text{s.t.} \quad & \bar{x}_{k,j,1} = x_{k,j}, && \forall\, k, j && \text{(2a)} \\
& \bar{y}_{k,j,\tau} = \bar{x}_{k,j,\tau} + \Lambda_{j,\tau} \sum_S \gamma_{j,\tau}(S)\pi_k(S), && \forall\, k, j, \tau && \text{(2b)} \\
& \bar{x}_{k,j,\tau+1} = \bar{y}_{k+1,j,\tau} - \sum_i \bar{u}_{k+1,j,i,\tau}, && \forall\, k < K, j, \tau && \text{(2c)} \\
& \bar{x}_{K,j,\tau+1} = 0, && \forall\, j, \tau && \text{(2d)} \\
& \sum_i \bar{u}_{k,j,i,\tau} \le \bar{y}_{k,j,\tau}, && \forall\, k, j, \tau && \text{(2e)} \\
& \sum_S \gamma_{j,\tau}(S) = 1, && \forall\, j, \tau && \text{(2f)} \\
& \bar{u}_{k,j,i,\tau} \ge 0 && \forall\, k, j, i, \tau && \text{(2g)} \\
& \gamma_{j,\tau}(S) \ge 0, && \forall\, j, \tau, S \in \mathcal{N} && \text{(2h)}
\end{aligned}
$$

First, we establish that the optimal function value $V^D(x)$ for **DET** serves as an upper bound for the optimal expected profit $V^*(x)$ in **OPT**. We defer all the proofs to the electronic companion.

Proposition 1. $V^*(x) \le V^D(x)$

Note that **DET** constitutes a convex optimization problem with linear constraints, which implies that it may be solved by commercial solvers. However, it poses two challenges. Firstly, since we define a decision variable for each assortment, the number of decision variables grows *exponentially* with the number of fulfillment options. A common approach to address this issue is the column generation technique. Nonetheless, implementing this technique in **DET** presents difficulties, primarily because the fulfillment cost function C_i lacks a specific functional form, complicating the derivation of a dual formulation. Secondly, unlike **OPT**, **DET** eliminates all sources of randomness. As discussed in §5, the solution obtained from **DET** does not effectively adapt to demand realizations and dynamic changes in fulfillment capacity, leading to suboptimal performance in practice. By addressing the second issue delegated to the next section, in this section, our focus is to tackle the first challenge by characterizing specific structural properties of an optimal solution to **DET**. Specifically, in what follows, focusing on efficient assortments, we will demonstrate that only a limited number of assortments suffices to be considered in the optimal solution of **DET**.

An assortment is deemed *efficient* when it presents the most advantageous balance between expected revenue and capacity usage. In other words, an assortment S is considered efficient if there does not exist another fulfillment assortment that can yield higher revenue while requiring the same or less capacity. However, the challenge is that the exact capacity consumption for each warehouse remains unknown when an assortment decision has to be made. This is due to the flexibility in the fulfillment operations, which allows decisions to be made at the end of each fulfillment cycle. We resolve that issue by consolidating the available capacity from all warehouses when making an assortment decision. Following the seminar work Talluri and Van Ryzin (2004), we identify a necessary and sufficient condition for an assortment S to be efficient.

To proceed, we introduce the following notation and the formal definition of an efficient assortment. Let $R_j(S)$ denote the expected revenue generated by offering assortment S to customer j and $Q(S)$ denote the probability of purchase when S is offered:

$$R_j(S) = \sum_k r_{k,j}\pi_k(S)$$
$$Q(S) = \sum_k \pi_k(S) = 1 - \pi_\varnothing(S)$$

We can now define an efficient assortment as follows:

Definition 1 (Efficient assortments). An assortment $S \in \mathcal{N}$ is *efficient* for customer j if there does not exist a set of convex weights $\alpha(S)$ satisfying $\sum_S \alpha(S) = 1$ and $\alpha(S) \geq 0$ such that

$$R_j(S) < \sum_{S'} \alpha(S')R_j(S')$$
$$Q(S) \geq \sum_{S'} \alpha(S')Q(S')$$

Otherwise, the assortment S is *inefficient* for customer j.

We denote the set of efficient assortments as $\mathcal{E}_j$. The following result characterizes a necessary and sufficient condition for an assortment S to be efficient:

Lemma 1 (Talluri and Van Ryzin 2004). For customers j, an assortment S is efficient, i.e., $S \in \mathcal{E}_j$ if and only if, there exists some real number $\theta_j \leq 0$ such that S is the optimal solution to

$$\max_{S'} R_j(S') + \theta_j Q(S')$$

The following proposition shows that we can solve **DET** optimally by restricting our attention only to the set of efficient assortments.

Proposition 2. If $\gamma^*_{j,\tau}(S) > 0$ for some j and τ is an optimal solution to **DET**, then S is an efficient assortment.

The proof consists of two steps: First, we show that if $\gamma^*_{j,\tau}(S) > 0$, then S solves the following problem

$$\max_{S'} R_j(S') + \sum_k \theta_{k,j}\pi_k(S')$$

for some $\theta_j = (\theta_{1,j}, \cdots, \theta_{K,j})$. Second, we show that $\theta_{k,j}$ are identical and non-positive for all k. Finally, by Lemma 1, S must be efficient.

Proposition 2 has two critical implications. First, even though the number of assortments in $\mathcal{N}$ grows exponentially in K, it suffices to consider only a subset of $\mathcal{N}$ to solve **DET**, namely, the set of efficient assortments $\mathcal{E}_j$ for all j. Now, it begs the question: What is the time complexity of constructing the set of all efficient assortments? Despite lacking strong theoretical tractability for a general choice model, empirical results have shown that they often can be quickly determined with the "largest marginal revenue" algorithm (Talluri and Van Ryzin 2004) for many widely adopted choice models, e.g., multinomial logit model, generalized attraction model, and nested logit model (see Gallego and Topaloglu 2019). Indeed, in §7, we extend the "largest marginal revenue" algorithm by adding tie-breaking rules and show it works very well for our numerical study. Therefore, throughout the rest of the paper, we make the following assumption:

Assumption 3. We can construct a set of efficient assortments $\mathcal{E}_j$ for each j associated with a given choice model $\{\pi_k(\cdot) : k \in \mathcal{S}\}$ and a list of rewards $\{r_{k,j} : k \in [K], j \in [N]\}$.

The second implication from Proposition 2 is as follows: If we interpret $-\theta_{k,j}$ as the marginal value of satisfying a customer j with fulfillment option k, the second step in the proof of Proposition 2 then suggests that all the resources consumed by the customer j choosing fulfillment options $k \in [K]$ are balanced at the optimal solution across all k. Considering a fulfillment option as a product and a warehouse as a resource in the context of a choice-based network revenue management with multiple resources, this result extends the discussions of efficient sets to the case of flexible products.

4.1 Efficiency of optimal assortment to the original problem

Even though we restrict our attention to efficient assortments throughout the rest of the paper and show that it results in an effective algorithm with provable worst-case bounds, we would like to briefly comment on whether the optimal assortment for the original problem **OPT** satisfies the efficiency condition stated in Definition 1. To check this, we need to analyze the DP formulation of **OPT**. Specifically, let $V_\tau(x_\tau(t), t)$ denote the maximum total expected profit that can be obtained since time t during fulfillment cycle τ with pre-assortment state $x_\tau(t)$. We omit the dependence of x on τ and t and simply write it as $V_\tau(x,t)$. Let δt satisfies $\lambda_{j,\tau}(t)\delta t \ll 1$ for all j. We can formulate the original problem as follows:

$$\begin{aligned} V_\tau(x,t) &= \max_{S_j} \sum_j \lambda_{j,\tau}(t)\delta t \left[\sum_k \pi_k(S_j)\left(r_{k,j} + V_\tau(x+e_{k,j}, t+\delta t)\right)\right] \\ &\quad + \left(\sum_j \pi_\varnothing(S_j) + 1 - \sum_j \lambda_{j,\tau}(t)\delta t\right) V_\tau(x, t+\delta t) + o(\delta t) \\ &= \sum_j \lambda_{j,\tau}(t)\delta t \left[\max_{S_j} \sum_k \pi_k(S)\left(r_{k,j} + \Delta V(x, t+\delta t)\right)\right] + V_\tau(x,t+\delta t) + o(\delta t) \end{aligned}$$

where the boundary condition is

$$V_\tau(y_\tau, 1) = \max_u - \left(\sum_i C_i(w^\top u_i) + \sum_j b_j \left(u_{1,j} - \sum_i u_{1,j,i}\right)\right) + V_{\tau+1}(x_{\tau+1})$$

where $x_{\tau+1}$ is the pre-assortment state for the fulfillment cycle $\tau+1$, and it is obtained by the system dynamics. $V_{\tau+1}(\cdot)$ is the maximized expected profit since cycle $\tau+1$. In the DP formulation, $e_{k,j} \in \mathbb{R}^{KN}$ is a unit vector where the k,j-th element is one and the rest are zero, and

$$\Delta V(x, t+\delta t) = V_\tau(x+e_{k,j}, t+\delta t) - V_\tau(x, t+\delta t)$$

represents the difference in expected profit with or without this additional order $e_{k,j}$. It can be interpreted as a measure of the marginal fulfillment cost of that order. Note that with one more order to fulfill, $V_\tau(x + e_{k,j}, t + dt) \leq V_\tau(x, t + dt)$, and hence $\Delta V(x, t + dt) \leq 0$. Taking limit as $\delta t \downarrow 0$, we know the expected profit $V_\tau(x, t)$ needs to satisfy the Hamilton-Jacobi-Bellman (HJB) equation:

$$\frac{\partial V_\tau(x,t)}{\partial t} = \sum_j \lambda_{j,\tau}(t) \left[\max_{S_j} \sum_k \pi_k(S) \left(r_{k,j} + \Delta V(x,t) \right) \right]$$

and that the optimal assortment to show to a customer from j should solve the following problem

$$\max_S \; R(S) + \sum_k \Delta V(x,t) \pi_k(S)$$

Note that $\Delta V(x, t)$ depends on k via vector $e_{k,j}$. Hence, for an optimal assortment to the original stochastic problem to be efficient, we must show that $\Delta V(x, t)$ is identical for all k. Unfortunately, this is not true in general. Consider the following case: Suppose $t = 1$, and x_τ satisfies $x_{1,j,\tau} > 0$ for all j, $x_{k,j,\tau} = 0$ for all k, j, and all $x_{1,j,\tau}$ will be fulfilled completely but there is no additional capacity to take another order with $k = 1$. In that case, we must have $\Delta V(x, t) = -b_j$ for $k = 1$. In contrast, we can take an order with $k > 1$ and postpone its fulfillment to a later fulfillment cycle, resulting in $\Delta V(x, t) \geq -b_j$ for all $k > 1$ (Zhou et al. 2023, Proposition 2). Therefore, there is no guarantee that equality will always be maintained. According to Definition 1, efficient assortments may cease to be efficient for the original stochastic problem **OPT** depending on the state realization.

5 Proposed algorithm and performance analysis

The analysis in the previous section suggests the following two-step computationally tractable policy: In the first step, we construct all efficient option lists, denoted by $\mathcal{E}_j$, for all j. In the second step, we solve **DET** with $\mathcal{E}_j$ to compute probability $\gamma^*_{j,\tau}(S)$ and fulfillment decisions $u^*_{k,j,i,\tau}$. We then offer the customer from location j an assortment S with probability $\gamma^*_{j,\tau}(S)$ and use rounded version . This approach, referred to as "Certainty Equivalent" policy in the literature, is a classic and widely used heuristic known for its computational tractability and typically near-optimal performance, both theoretically and in practice.

However, this policy *may not* work well in our original setting. As implied by Proposition 2, when an efficient assortment S is offered to a customer j, the marginal fulfillment cost $\theta_{k,j}$ are balanced in **DET** for all $k \in S$. However, the stochasticity in demand realization violates this balance since we have shown that $\Delta V(x, t)$ may not be identical in §4.1, which results in S not being optimal for the original problem. We consider these options that violate the balance of the marginal fulfillment cost *economically nonviable*, which we will define formally later. We then remove them from the set of efficient assortments to ensure that the remaining fulfillment options remain economically viable.

An approach for identifying economically nonviable assortments involves utilizing dynamic marginal fulfillment cost, denoted as $\Delta V(x, t)$. However, computing these costs is intractable due to the curse of dimensionality. We approximate it with the *lowest fulfillment cost increment if we accept an order*, which involves solving a multi-cycle fulfillment problem. Before getting into the details, let us elucidate this concept with an example. For simplicity, we assume that the DM operates a single warehouse and has access to a limited fulfillment capacity Ξ at a unit rate c. The DM has *no* access to expedited fulfillment if the total quantity exceeds the capacity. We can model this using a linear fulfillment cost function on bounded support and setting the penalty cost to $b_j = \infty$. Thus, the marginal fulfillment cost is c if the total fulfillment does not exceed Ξ; otherwise, it is ∞ and economically nonviable.

Let us consider the company has already accepted some orders, denoted as $x = (x_{k,j} : k \in [K], j \in [N])$, at the start of time t in fulfillment cycle τ. When deciding to offer the fulfillment assortment S to a focal customer. If option $k = 1$ is included in S, the lowest fulfillment cost increment is c if $\sum_j x_{1,j} + 1 \leq \Xi$, otherwise, the lowest fulfillment cost increment will be ∞. Furthermore, the cost increment also depends on the subsequent fulfillment cycles as accepting a new order with $k = 1$

triggers a chain reaction: The DM might need to defer another order with a longer fulfillment window to accommodate this immediate order. For instance, the focal customer selects the fulfillment option $k = 1$, resulting in an order with $k = 2$ being postponed. In this case, the cost increment associated with option $k = 1$ will still be ∞ if condition $\sum_j x_{2,j} + 1 \leq \Xi$ does not hold. Consequently, it is imperative to ensure that the current and subsequent fulfillment cycles have adequate capacity. This chain reaction can potentially impact the subsequent $K - 1$ fulfillment cycles. Generally, offering any fulfillment option k necessitates evaluating economic viability for the subsequent fulfillment cycles beginning with the k-th cycle. In other words, to ascertain whether there is sufficient capacity to accept an order with a given fulfillment option k, it becomes essential to address a $(K + 1 - k)$-cycle fulfillment problem.

Formally, when customer j arrives, the pre-assortment state is $x_\tau(t) = (x_{k,j,\tau}(t) : k \in [K], j \in [N])$, $t \in [0, 1]$. The focal customer is presented with an assortment S and selects a fulfillment option $k' \in S$ with a probability $\pi_{k'}(S)$. Consequently, the post-assortment state becomes $y_\tau(t) = x_\tau(t) + e_{k',j}$. We solve the following problem:

$$
\begin{aligned}
\mathcal{F}(x_\tau(t)) = \min_{\tilde{u}} \quad & \sum_{\tau'=\tau}^{\tau+K-1} \left\{ \sum_i C_i(w^\top \tilde{u}_{i,\tau'}) + \sum_j b_j \left(\tilde{y}_{1,j,\tau'} - \sum_i \tilde{u}_{1,j,i,\tau'} \right) \right\} \\
\text{s.t.} \quad & \tilde{y}_{k,j,\tau} = x_{k,j,l,\tau}(t) + \mathbf{1}_{(k=k')}, \quad \forall\, k, j \\
& \tilde{y}_{k,j,\tau'+1} = \tilde{y}_{k+1,j,\tau'} - \sum_i \tilde{u}_{k+1,j,i,\tau'}, \quad \forall\, k < K, j, \tau' < \tau + K - 1 \\
& \sum_i \tilde{u}_{k,j,i,\tau'} \leq \tilde{y}_{k,j,\tau'}, \quad \forall\, k, j, \tau' \\
& \tilde{u}_{k,j,i,\tau'} \geq 0, \quad \forall\, k, j, i, \tau'
\end{aligned}
\tag{3}
$$

The above is a pure fulfillment problem, assuming the existing orders are x, with *no* anticipation of additional demand. The chain reaction described earlier is implicitly managed by the system dynamics. Note that $\mathcal{F}(x)$ can be solved efficiently as we do not require fulfillment action to be integral. The function $\mathcal{F}(x)$ can be considered as the minimum cost to fulfill x orders optimally. Therefore, $\mathcal{F}(x_\tau(t) + e_{k',j}) - \mathcal{F}(x_\tau(t))$ can be interpreted as an approximation to the marginal fulfillment cost $\Delta V_\tau^{k',j}(x, t)$ incurred by accepting order with fulfillment option k' at time t during cycle τ. An option k' is deemed economically nonviable if its associated revenue is smaller than the minimum marginal fulfillment cost. With that, we can formally define an economically viable fulfillment option.

Definition 2 (Economically nonviable assortments)**.** We say the fulfillment option $k' \in S$ is *economically nonviable* for customer from location j if $r_{k',j} < \mathcal{F}(x_\tau(t) + e_{k',j}) - \mathcal{F}(x_\tau(t))$; Otherwise, we say k' is *economically viable.*

For all $k' \in S$, we determine its viability by solving the above fulfillment problem. If k' is nonviable, we remove it from the assortment and only show the economically viable options. If there is no viable option, we offer the customer $S = \varnothing$. The heuristic policy is described in Algorithm 1, and we name it the Economic Assortment Threshold Fulfillment Policy (**EATF**). Note that by design **EATF** is a dynamic policy and depends on the pre-assortment state $x_{k,j,\tau}(t)$. Steps 1 and 2 in Algorithm 1 represent the offline phase in which the set of efficient assortments are constructed and **DET** is optimized. For each fulfillment cycle τ, steps 4 to 12 dynamically determine the personalized fulfillment options by eliminating economically nonviable options from the efficient assortment, and steps 14 to 23 provide the threshold-based fulfillment assignments, which facilitates the structural property of an optimal fulfillment decision with arbitrary demand (see Lemma 3).

The determination of the personalized fulfillment options is also illustrated in Figure 3, where we assume to show an efficient assortment $S = \{(k_1, r_{k_1}), (k_2, r_{k_2}), (k_3, r_{k_3})\}$ to a customer j during fulfillment cycle τ. The figure depicts the marginal fulfillment cost for each option as a rectangle. At elapsed time t_1, since options are viable by Definition 2, we show the entire efficient assortment S. The customer chooses to purchase, leading to an increase in the state. Therefore, when another customer j arrives at time t_2, the marginal fulfillment cost increases such that option k_3 is nonviable. We remove it from S and only show the customer options k_1 and k_2. Later, at time t_3, a new customer j arrives. The calculated marginal fulfillment cost of k_1 at time t_3 exceeds its revenue, so k_1 becomes nonviable and removed. Only k_2 is offered to the customer.

Algorithm 1: Economic Assortment Threshold Fulfillment Policy (EATF)

Input: Customer arrival rate vector $\Lambda_{j,\tau}$, revenue parameters $r_{j,k}$, the fulfillment cost function $C_i(\cdot)$, the penalty vector b_j, $\forall j \in [N]$, and the choice model $\pi_k(S)$.

Output: $S^{\text{ERT}} = \{S_{j,\tau}(t)\}$ and $u^{\text{ERT}} = \{u_{k,j,i,t}\}$

Construct the set of all efficient option lists $\mathcal{E}_j$ for all $j \in [N]$.

Solve **DET** to obtain $\gamma^*_{j,\tau}(S)$ for all $S \in \mathcal{E}_j$, $\bar{u}^*_{k,j,i,\tau}$, and $\bar{y}_{k,j,\tau}$.

for *each cycle* $\tau \in [T]$ **do**

  for *a particular time* t *within each cycle* τ **do**

    if *a customer arrives from location* j **then**

      First, choose one option list $S \in \mathcal{E}$ with probability $\gamma^*_{j,\tau}(S)$.

      for *each option* $k' \in S$ **do**

        Remove that option if it is *economically nonviable* by using Definition 2, i.e., $S_{j,\tau}(t) = S_{j,\tau}(t) \setminus \{k'\}$.

      end

      Finally, offer a personalized delivery assortment $S_{j,\tau}(t)$ to the customer.

    end

  end

  At the end each cycle τ, aggregate all orders and generate fulfillment decisions as follows.

  Set $U_{i,\tau} = 0$ for all i, and $\mathcal{I}_\tau = [M]$.

  for *each order received from location* $j \in [N]$ *in descending order by* b_j **do**

    for *each delivery option* $k = 1, 2, \cdots, K$ **do**

      for *each warehouse* $i \in \mathcal{I}_\tau$ **do**

        Calculate a scale factor $\rho_{k,j,i,\tau} = \frac{\bar{u}^*_{k,j,i,\tau}}{\bar{y}_{k,j,\tau}}$

        Determine fulfillment quantity by

$$u_{k,j,i,\tau} = \begin{cases} \rho_{k,j,i,\tau} y_{k,j,\tau} & \text{if } \mathcal{C}'_i\left(U_{i,\tau} + w_j \rho_{k,j,i,\tau} y_{k,j,\tau}\right) \le b_j \\ z \text{ that solves } \mathcal{C}'_i\left(U_{i,\tau} + w_j z\right) = b_j & \text{if } \mathcal{C}'_i\left(U_{i,\tau}\right) < b_j \le \mathcal{C}'_i\left(U_{i,\tau} + w_j \rho_{k,j,i,\tau} y_{k,j,\tau}\right) \\ 0 & \text{otherwise} \end{cases}$$

        Update $U_{i,\tau} = U_{i,\tau} + w_j u_{k,j,i,\tau}$.

        Remove warehouse i from from $\mathcal{I}_\tau$ if $\mathcal{C}'_i\left(U_{i,\tau}\right) \ge b_j$.

      end

    end

  end

end

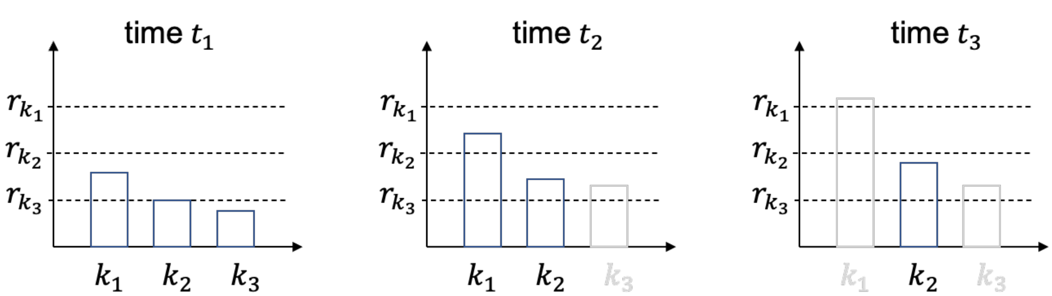

Figure 3: Illustration of dynamically eliminating economically nonviable fulfillment options for some customer j

Remark 1. During any fulfillment cycle τ, we have the following observations. (i) If some fulfillment option $k' \in S$ is nonviable for demand location j at time t, then it must remain nonviable to the end of this fulfillment cycle. This is because the post-assortment state is monotonically increasing within the fulfillment cycle. (ii) If some fulfillment option $k' \in S$ is nonviable at time t, any fulfillment option $k \in S$ such that $r_{k,j} < r_{k',j}$ and $k < k'$ must also be nonviable. This is because the fulfillment option

k cannot have a lower marginal fulfillment cost than option k' since otherwise, we can always fulfill an order with option k' earlier. (iii) For a similar reason, if some fulfillment option $k' \in S$ is viable at time t, then any fulfillment option $k \in S$ such that $r_{k,j} \geq r_{k',j}$ and $k > k'$ must also be viable. Such observations are crucial for implementing **EATF** as it reduces the number of times we solve problem (3).

Next, we upper bound the performance of our algorithm **EATF** in the following theorem:

Theorem 1. Let $\Lambda = \max_{j,\tau} \Lambda_{j,\tau}$, $\overline{w} = \max_j w_j$. Then, the optimality gap of **EATF** is bounded by

$$V^*(x) - V^{\mathrm{EATF}}(x) \leq \overline{w}\overline{C}MTK\sqrt{\frac{1}{2}\Lambda KN}$$

Delegating the detailed proof of Theorem 1 to the electronic companion, here, we outline the main steps of our proof: First, we define a hypothetical policy (denoted by **HYP**) that utilizes the same efficient assortments constructed by step 1 of **EATF** algorithm. However, **HYP** does two things differently. First, it does not eliminate economically nonviable assortments dynamically. Second, it disregards the implicit capacity constraint on fulfillment and incurs higher costs for fulfilling excessive orders. Let $V^{\mathrm{HYP}}(x)$ be the expected total profit of applying this policy. Using sample path arguments, we show that **EATF** cannot yield worse performance than **HYP**, i.e., $V^{\mathrm{HYP}}(x) \leq V^{\mathrm{EATF}}(x)$. Then, in the second step, we characterized the optimality gap of the hypothetical policy and established that it is upper bounded by $\overline{w}\overline{C}MTK\sqrt{\frac{1}{2}\Lambda KN}$.

The performance bound characterized in Theorem 1 shows that the optimality gap of **EATF** is growing linearly with respect to the number of warehouses M and sub-linearly with respect to the number of locations N and the maximum cycle arrival rate Λ. It has the following implications. The total amount of realized demand increases linearly as either the number of locations N or the cycle arrival rate $\Lambda_{j,\tau}$ increases. Since the fulfillment cost is increasing and convex in the quantity of fulfillment, the total cost grows at least linearly in both N and $\Lambda_{j,\tau}$. But the optimality gap grows sub-linearly, implying that the relative optimality gap converges to 0 as either N or $\Lambda_{j,\tau}$ (hence Λ) increases. Consequently, our policy becomes asymptotically optimal in large N or Λ. This implication is particularly crucial in online retailing, where there are typically a vast number of demand locations, and the arrival rate is increasing steadily.

6 Experiment 1: The performance evaluation of EATF

In this section, we conduct a sensitivity analysis to evaluate the performance of **EATF** and generate additional insights. First, we introduce three benchmark policies. Subsequently, we calibrate a choice model that satisfies Assumption 1 using the real data from our industrial partner. This model is extended to encompass a more complex and realistic representation of customer behavior in §7. Finally, we examine the value of jointly optimizing assortments and fulfillment decisions by comparing **EATF** to the benchmark policies across different scenarios derived from the estimated choice model.

We use an upper bound of the percentage profit loss of a policy as the performance measure. It is calculated as $(V^{\mathbf{DET}} - V^{\mathbf{policy}})/V^{\mathbf{DET}} \times 100\%$, where $V^{\mathbf{policy}}$ is the expected profit obtained from some policy, and V^{DET} is the expected profit from **DET**.

6.1 Benchmark policies

E-commerce practitioners often have distinct teams responsible for designing the fulfillment options at checkout and making fulfillment decisions. Consequently, each team may act based on their perspectives without considering the entire system. To emulate these operations, we introduce three benchmark policies. These policies are characterized by whether or not assortment and fulfillment decisions are optimized. The policies are summarized in Table 1:

Table 1: Summary of the Policies.

		Fulfillment	
		Optimized (**TF**)	Myopic (**MF**)
Assortment	Optimized (**EA**)	**EATF**	**EAMF**
	Static (**FA**)	**FATF**	**FAMF**

The first benchmark consists of a fixed-assortment strategy combined with a myopic fulfillment strategy, referred to as the **FAMF** policy. This policy presents a pre-determined and fixed assortment to customers throughout the horizon. Fulfillment decisions are made by solving the following optimization problem:

$$\begin{aligned} \min_u \; & \sum_i C_i(u_i) + \sum_j b_j \left(y_{k,j} - \sum_i u_{k,j,i} \right) \\ \text{s.t.} \; & \sum_i u_{k,j,i} \le y_{k,j}, \quad u \ge 0 \end{aligned}$$

The fixed assortment strategy is selected as follows: We evaluate all possible assortments for each simulation case and choose the one with the best performance as the fixed assortment for that particular case. This policy mimics a DM's operations when they know customer preferences but does not optimize the assortment or fulfillment decisions.

The second benchmark policy only optimizes its assortment decisions but not the fulfillment ones. It solves the following assortment optimization problem under capacity constraints (specified in §6.2.2) to determine the assortment decisions, while fulfillment decisions are made in the same manner as in **FAMF**.

$$\begin{aligned} \max_\gamma \quad & \sum_\tau \sum_j \sum_k \sum_S r_{k,j} \Lambda_{j,\tau} \gamma_{j,\tau}(S) \pi_k(S) & \\ \text{s.t.} \quad & \sum_j \sum_k \sum_S \Lambda_{j,\tau} \gamma_{j,\tau}(S) \pi_k(S) \le \Xi, & \forall\, \tau \\ & \sum_S \gamma_{j,\tau}(S) \le 1, & \forall\, j, \tau \\ & \gamma_{j,\tau}(S) \ge 0, & \forall\, j, \tau, S \in \mathcal{N} \end{aligned}$$

Ξ specifies the fulfillment capacity for each cycle. We refer to this as the Efficient Assortment Myopic Fulfillment policy (**EAMF**). It is important to note that, for each fulfillment cycle, **EATF** *dynamically* adjusts assortment decisions based on the remaining fulfillment capacity, whereas **EAMF** remains *stationary.*

In contrast to the second policy, the last one only optimizes the fulfillment decisions. It fixes an assortment and solves **DET**, then determines the actual fulfillment in the same manner as **EATF**. It may adopt a different, but best, fixed assortment for different cases. We refer to it as the Fixed Assortment Threshold Fulfillment policy (**FATF**).

6.2 Calibration

Our industrial partner is a growing online grocery retailer that began business within the province and is gradually expanding nationwide. Customers visit its website to shop for groceries. At the checkout, they have to choose either home-delivery or self-pickup and specify a date. Due to the retailer's business scale and the operational limit, orders are fulfilled once every week. Thus, customers select the week they wish to receive their orders. A shipping surcharge is displayed based on the selected fulfillment option, allowing customers to explore the cost differences among various options and choose their preferred ones. The transaction is complete once the fulfillment option is finalized and the payment is processed. The retailer fulfills orders from a single warehouse within the customer-selected fulfillment window, using either a local trucking company or a national carrier; both incur flat-rate costs.

The dataset comprises 6,511 transactions (including abandoned transactions) from 2,961 customers from 476 cities across Canada from January 1, 2023, to December 31, 2023. Our analysis focuses on 2,833 transactions from 21 cities, as the local trucking company and self-pickup service are available only to customers in these cities. Due to practical constraints, customers from other locations can only choose home delivery with the national carrier. The dataset contains detailed information on each order, including the number of products, total basket value, the selected fulfillment option, customer location, and whether the transaction was completed.

6.2.1 Basic attraction choice model.

The dataset allows us to calibrate a basic attraction choice model for customer behavior, adhering to Assumption 1. To achieve that, we construct three fulfillment options based on the length of the fulfillment window: fulfill by this week with a mean surcharge of 6.46 Canadian dollars (Option 1), fulfill by next week with a mean surcharge of 4.37 Canadian dollars (Option 2), and fulfill after two weeks with a mean surcharge of 1.69 Canadian dollars (Option 3). We then estimate the attraction of each option (denoted by a) based on the proportion of customers who selected the corresponding time frame. Normalizing the attraction of abandonment to one ($a_\varnothing = 1.00$), we find $a_1 = 3.47$, $a_2 = 10.48$, $a_3 = 0.12$. These results indicate that Option 2, which balances waiting time and fulfillment surcharge, is the most attractive to customers.

6.2.2 Baseline setup.

We introduce a baseline simulation setup calibrated using real data with certain simplifications. We assume a homogeneous demand generation process: customers from different locations arrive following the same stationary Poisson process and choose a fulfillment option following the same choice model. In addition, each order has an identical basket value and generates the same gross profit. These assumptions will be generalized in §7. The billable weight is set to one due to the lack of specific information.

Regarding the fulfillment network, we set $M = 1$ and $N = 21$, as the retailer fulfills orders from 21 cities using a single warehouse. Customers have three options regarding the fulfillment window, i.e., $K = 3$. We consider a planning horizon of one year, which includes 52 fulfillment cycles ($T = 52$). Each week, customers arrive at the system following a stationary Poisson process with an intensity of 19.5 ($\lambda_{j,\tau} = 19.5$ for all j and τ). Each order is valued 140.00 Canadian dollars ($p_j = 140.00$ for all j). Based on the discussion with our industrial partner, we set the percentage applied to each basket size as 5% ($\kappa = 0.05$). Together with the fulfillment surcharge, the revenue generated by each option is $r_{1,j} = 13.46$, $r_{2,j} = 11.37$, and $r_{3,j} = 8.69$ Canadian dollars, respectively.

The retailer employs two trucking companies to fulfill the orders. It can fulfill up to 220 orders per week using the local trucking company for 8.00 Canadian dollars per order. The retailer turns to the national carrier for orders exceeding this number, which can provide sufficient capacity for 16.00 Canadian dollars per order. Thus, we define a two-segment piecewise linear convex function $C(z) = \max\{8z, 16z - 1760\}$ to model the retailer's fulfillment cost. We assign a large value to the penalty cost (e.g., $b_j = 20$ for all j) to avoid using expedited fulfillment service. Clearly, using the national carrier results in a profit loss due to the higher cost than the revenue. We refer to the *fulfillment capacity* as the maximum fulfillment quantity by the local trucking company and the *relative capacity* as the ratio of fulfillment capacity to the expected number of arriving customers ($\sum_j \lambda_{j,\tau} = 21 \times 19.5 = 409.5$). In the baseline case, the capacity is $\Xi = 220$, and the relative capacity is 53.7%.

In what follows, we investigate the value of jointly optimizing the assortments and fulfillment decisions by comparing **EATF** to the benchmark policies under various scenarios.

6.3 Varying customers preferences

We modify the parameters of the choice model to investigate these policies under various customer behaviors. Our estimated choice model represents a market where most customers prefer a balance between fulfillment cost and waiting time. We simulate different customer preferences by swapping the attraction values among the choices. For instance, swapping the values a_1 and a_2 imitates a scenario where customers prefer faster fulfillment (Case 2). We also increase $a_\varnothing$ to 3.0 to simulate the scenarios when more customers abandon carts if they are dissatisfied with offered fulfillment options.

To summarize, we create 6 cases: Case 1 represents the actual market, Case 2 represents customers who prefer faster fulfillment, Case 3 represents customers who prefer cheaper fulfillment, and Cases 4 to 6 replicate Cases 1 to 3 with a higher abandonment probability. We report the percentage profit loss in Table 2. Note that a smaller profit loss indicates a better policy performance. Unsurprisingly, we observe that **EATF** outperforms all benchmark policies in all scenarios.

Table 2: Comparison of EATF with respect to various benchmark policies under different customer preferences

Case	(a_0, a_1, a_2, a_3)	ERT	OL	OF	SM
1	(1.00, 3.47, 10.48, 0.12)	2.7%	4.0%	21.5%	24.7%
2	(1.00, 10.48, 3.47, 0.12)	2.8%	3.8%	33.1%	35.9%
3	(1.00, 0.12, 3.47, 10.48)	1.6%	22.2%	68.0%	68.3%
4	(3.00, 3.47, 10.48, 0.12)	2.8%	4.2%	3.7%	3.8%
5	(3.00, 10.48, 3.47, 0.12)	2.8%	4.1%	20.7%	24.1%
6	(3.00, 0.12, 3.47, 10.48)	1.4%	25.7%	3.6%	28.5%

Upon closer examination, it is evident that demand management (via optimizing assortment decisions) is more crucial than capacity management (via optimizing fulfillment decisions) when fewer customers abandon their carts. This can be seen by comparing **EAMF** to fixed-assortment policies in Cases 1 to 3. Moreover, when customers prefer fast fulfillment in Cases 1 and 2, many will choose Option 1. This necessitates fulfilling most orders by the end of this week, leaving little room for optimized fulfillment decisions to improve performance. As a result, all the additional benefits stem from dynamic assortment management, reflected in the smaller percentage profit losses of **EATF** compared to **EAMF**. In contrast, when customers prefer cheaper fulfillment in Case 3, more orders will have longer fulfillment windows. This scenario requires careful consideration of fulfillment capacity and active postponement of some fulfillments to later cycles. Thus, joint optimization leads to significant improvements.

Analyzing Cases 4 to 6, where the abandonment probability is higher, provides further insights. Customer dissatisfaction with fulfillment options leads to increased cart abandonment in these cases. This emphasizes the importance of offering attractive fulfillment options but simplifies the assortment decisions. Specifically, in Cases 4 and 6, where customers less prefer fast fulfillment, fixed-assortment policies achieve comparable, sometimes even better, performance than **EAMF**. In contrast, **EATF** still performs better than the benchmarks, highlighting its consistency and robustness.

6.4 Varying number of locations

This section aims to investigate the benefits of adopting **EATF** when the retailer expands the business to a larger area. We vary the value of N from 5 to 50. In the meantime, we increase the local fulfillment quantity so that the relative fulfillment capacity remains unchanged. The capacity is set to $\frac{220}{21}N$. We report the results in Table 3, columns 2 to 5.

Table 3: Comparison of EATF with various benchmark policies under different number of customer locations

	Actual Market				Counterfactual Market			
N	ERT	OL	OF	SM	ERT	OL	OF	SM
5	6.0%	7.9%	21.4%	28.3%	2.6%	45.8%	45.1%	68.4%
10	4.1%	5.6%	21.5%	26.1%	2.2%	33.2%	31.5%	68.1%
20	2.9%	4.0%	21.6%	24.8%	1.3%	24.9%	22.0%	68.2%
30	2.3%	3.3%	21.5%	24.3%	1.3%	19.7%	16.5%	68.3%
40	2.0%	2.8%	21.5%	23.9%	1.2%	19.0%	14.8%	68.3%
50	1.9%	2.6%	21.4%	23.7%	1.0%	15.7%	11.6%	68.1%

Consistent with previous results, **EATF** performs the best in all cases. Moreover, the profit loss decreases as N increases. This supports our theoretical performance guarantee in Theorem 1, implying that the retailer will benefit from expanding the business to a larger area, provided it can increase the fulfillment capacity accordingly. In contrast, the profit losses for **FATF** and **FAMF** do not decrease as N increases.

Regarding **EAMF**, the profit loss also decreases similarly to **EATF**. To further understand whether the asymptotic optimality arises from joint optimization or assortment optimization, we repeat the test in a counterfactual market where customers are more price-sensitive and willing to wait (i.e., Case 3 in §6.3). The results, shown in Table 3, columns 6 to 9, reveal that better fulfillment decisions are crucial in this scenario. Therefore, we conclude that asymptotic optimality is a unique benefit of joint optimization.

6.5 Varying fulfillment capacity

This section examines the benefit of **EATF** when the retailer negotiates different capacities with the local trucking company. We adjust the maximum fulfillment quantity from the local trucking company so that the relative capacity varies from 30% to 80%. We report the simulation results in Table 4.

Table 4: Performance comparison of EATF with benchmark policies under different fulfillment capacity

Relative capacity	ERT	OL	OF	SM
80%	1.5%	6.5%	2.1%	2.1%
70%	2.5%	3.5%	5.8%	9.1%
60%	2.6%	3.8%	14.5%	17.7%
50%	2.9%	4.1%	26.5%	29.9%
40%	3.1%	4.5%	45.0%	48.3%
30%	3.7%	5.3%	75.4%	78.7%

As observed, **EATF** outperforms the benchmarks in all the cases. Interestingly, when capacity is relatively insufficient (less than or equal to 70%), meaning the availability of cheaper logistics is limited, the retailer must control the quantity of realized demand with optimized assortment decisions. If too many orders are realized (as with **FATF** and **FAMF**), the expected profit will be lower due to excessive usage of the more expensive carrier. However, the results change drastically when the capacity is relatively sufficient (80%). **EAMF** tends to accept more orders with a lower surcharge than the other three policies, which backfires. It has to use the national carrier more to fulfill those orders, resulting in a profit loss. This interesting finding emphasizes the necessity to align demand management with capacity management.

6.6 Additional scenarios

We also investigate the performance of **EATF** in scenarios when more customers will buy groceries from the retailer (i.e., increasing arrival rate) and when the retailer opens additional warehouses (i.e., increasing the number of warehouses M). We observe that **EATF** achieves better performance if more customers from current locations will buy from the retailer, implying it is asymptotically optimal in arrival rate. This result also supports the performance guarantee in Theorem 1. Regarding the number of warehouses, policy performance seems unaffected by M. It implies that **EATF** is still applicable if the retailer is going to open more warehouses as it expands the business. All details are available in B.

7 Extension to complex customer preference

Previous sections are based on Assumption 1, which essentially implies that customer preference for fulfillment options is solely influenced by the speed (i.e., the maximum duration they have to wait before receiving their orders). However, real-world customer behaviors are more complex. Customers have heterogeneous preferences for shipping speeds and costs, and their choices may be influenced by how much product they buy (we refer to it as the basket value). For example, customers who purchase large baskets may be less sensitive to shipping surcharges as the surcharges account for only a tiny proportion, while those with smaller baskets may be more sensitive. Therefore, online retailers have incentives to offer customers different combinations of fulfillment windows and shipping surcharges. For example, our industrial partner provides various options with varying surcharges for the same fulfillment window.

This section extends the formulation and the **EATF** policy to scenarios with such preferences. Specifically, we denote p_j^q as the basket size of type q from demand location j, where $q \in [Q]$ and $j \in [N]$. A customer from demand location j is considered to be of type q with probability v_j^q, and $\sum_{q\in[Q]} v_j^q = 1$ for all j. Thus, if a customer of type q visits the store, she results in an order with a fulfillment reserve of κp_j^q with probability v_j^q. In addition to the customer types, we extend the main model to adjust fulfillment surcharge based on the customer type dynamically. To operationalize this, we define a discrete set of surcharge prices $\mathcal{R}$. The set of fulfillment options is therefore represented as $\mathcal{S} = \{(k,r) : k \in [K], r \in \mathcal{R}\}$, which implies that the retailer may assign different surcharges to a delivery option with k fulfillment window. We denote $\mathcal{N}$ the powerset of $\mathcal{S}$. Each customer is going to choose one option from an assortment $S \in \mathcal{N}$ with probability $\pi_{(k,r)}(S \mid q)$, reflecting her decision depends on the basket value.

The sequence of events is as follows. First, a customer from location j arrives at the system. Then, she reveals her type and forms an online shopping cart with basket value p_j^q. After that, the online retailer forms a list of fulfillment options S and shows it to the customer at checkout. Finally, the customer chooses a delivery option and completes the purchase or leaves with an abandoned cart. Denote $D_{k,j,\tau}^{q,r}$ the number of customers who are type q and come from location j and choose a delivery option (k,r) during fulfillment cycle τ. Hence, we can calculate the expected revenue as follows:

$$\mathbb{E}\{Revenue\} = \sum_{k,j,\tau} \sum_{r\in\mathcal{R},q\in[Q]} (r + \kappa p_j^q)\mathbb{E}\{D_{k,j,\tau}^{q,r}\}$$

To formulate the DP, we only need to replace the revenue part in (1) with the above term and update the system dynamics according to $y_{k,j,\tau} = x_{k,j,\tau} + \sum_{q,r} D_{k,j,\tau}^{q,r}$. Moreover, to get a modified **DET** formulation, we offer assortment S to $\gamma_{j,\tau}^q(S)$ proportion of q-typed customers from location j during cycle τ. Hence,

$$\mathbb{E}\{D_{k,j,\tau}^{q,r}\} = \Lambda_{j,\tau} v_j^q \sum_S \gamma_{j,\tau}^q(S)\pi_{(k,r)}(S \mid q)$$

and the modified **DET** is as follows.

$$
\begin{aligned}
V^D(x) = \max_{\gamma,\bar{u}} & \sum_{\tau}\sum_{j}\sum_{q}\sum_{S}\sum_{(k,r)\in S}(r+\kappa p_j^q)\Lambda_{j,\tau}v_j^q\gamma_{j,\tau}^q(S)\pi_{(k,r)}(S\mid q) \\
& -\sum_{\tau}\left(\sum_i C_i(w^\top \bar{u}_{i,\tau})+\sum_j b_j\left(\bar{y}_{1,j,\tau}-\sum_i \bar{u}_{1,j,i,\tau}\right)\right)-\sum_{k,j} b_j\bar{x}_{k,j,T+1} \\
\text{s.t.}\quad & \bar{x}_{k,j,1} = x_{k,j}, && \forall\, k,j \\
& \bar{y}_{k,j,\tau} = \bar{x}_{k,j,\tau} + \sum_{q,S}\sum_{r:\,(k,r)\in S}\Lambda_{j,\tau}v_j^q\gamma_{j,\tau}^q(S)\pi_{(k,r)}(S\mid q), && \forall\, k,j,\tau \\
& \bar{x}_{k,j,\tau+1} = \bar{y}_{k+1,j,\tau} - \sum_i \bar{u}_{k+1,j,i,\tau}, && \forall\, k<K,j,\tau \\
& \bar{x}_{K,j,\tau+1} = 0, && \forall\, j,\tau \\
& \sum_i \bar{u}_{k,j,i,\tau} \le \bar{y}_{k,j,\tau}, && \forall\, k,j,\tau \\
& \sum_S \gamma_{j,\tau}^q(S) = 1, && \forall\, q,j,\tau \\
& \bar{u}_{k,j,i,\tau} \ge 0 && \forall\, k,j,i,\tau \\
& \gamma_{j,\tau}^q(S) \ge 0, && \forall\, q,j,\tau,S
\end{aligned}
$$

The above formulation suffers from the exponentially large number of decision variables of $\gamma_{j,\tau}^q(S)$, which has the dimension of $NTQ\cdot 2^{K|\mathcal{R}|}$.

Despite that, our main results still hold with this formulation. Specifically, $V^D(x)$ provides an upper bound for $V^*(x)$ (Proposition 1). In addition, the expected revenue and the purchase probability of type q customer, when the assortment S is offered, is

$$
\begin{aligned}
R_j^q(S) &= \sum_{(k,r)\in S}(r+\kappa p_j^q)\pi_{(k,r)}(S\mid q) \\
Q^q(S) &= \sum_{(k,r)\in S}\pi_{(k,r)}(S\mid q)
\end{aligned}
$$

After that, we modify the definition of efficient assortment with respect to customer location j and type q. As a result, we can modify Proposition 2 to the following

Proposition 3. If $\gamma_{j,\tau}^{q,*}(S)>0$ for some q, j and τ is an optimal solution to **DET**, then S is an efficient assortment.

The proof also consists of two steps: First, we show that if $\gamma_{j,\tau}^{q,*}(S)>0$, then S solves the problem

$$
\max_{S'} R_j^q(S') + \sum_{k,r}\theta_{k,j}\pi_{(k,r)}(S'\mid q)
$$

for some $\theta_j=(\theta_{1,j},\cdots,\theta_{K,j})$. Second, we show that $\theta_{k,j}$ are identical and non-positive for all k. Finally, by Lemma 1, S must be efficient. The proof is almost identical to that of Proposition 2 and thus is omitted.

Proposition 3 implies that we can solve the problem with the **EATF** policy in Algorithm 1 with minor modification. Specifically, we have to find the efficient set $\mathcal{E}_j^q$ for each customer location j and type q. This is because different typed customers have different preferences, and customers from various locations contribute different fulfillment reserves.

However, we cannot directly apply the largest marginal revenue algorithm in Talluri and Van Ryzin (2004) to this case since the assumption that revenues associated with each fulfillment option are monotonic does not hold, resulting in potentially multiple delivery options satisfying the maximum marginal revenue ratio. We modify that algorithm by adding a tie-breaking rule: for any customer type q, we pick the option that yields the smallest purchase probability $Q^q(S)$. By iteratively adding a new option at each step, we can construct all efficient assortments as characterized in Proposition 4. We use $\Omega(S_m)$ to represent the assortment set under consideration during the m-th recursion.

$$\Omega(S_m) = \{S \in \mathcal{N} : \; R_j^q(S) \geq R_j^q(S_m),\, Q^q(S) \geq Q^q(S_m),\, S \notin \mathcal{E}_j^q\} \tag{4}$$

Proposition 4. Algorithm 2 finds all efficient assortments for q typed customer from location j.

Finally, after replacing the set of efficient assortments with those generated by Algorithm 2, we can then use the same **EATF** approach described in Algorithm 1 to solve the joint optimization problem.

Algorithm 2: Modified Largest Marginal Revenue Algorithm

Input: Set of fulfillment window $[K]$, set of surcharges $\mathcal{R}$, reward parameter κp_j^q, and choice model $\pi_{(k,r)}(S \mid q)$.

Output: The efficient assortment set $\mathcal{E}_j^q$

Start from $S_0 = \varnothing$, $m = 0$, and $\mathcal{E}_j^q = \{S_0\}$.

while $\Omega(S_m) \neq \varnothing$ **do**

  $\Theta =_{S' \in \Omega(S_m)} \frac{R_j^q(S') - R_j^q(S_m)}{Q^q(S') - Q^q(S_m)}$

  Set $S_{m+1} =_{S' \in \Theta} Q^q(S')$, breaking ties arbitrarily if there are multiple assortments with the same minimum $Q^q(S')$.

  Update $\mathcal{E}_j^q = \mathcal{E}_j^q \cup \{S_{m+1}\}$ and $m \leftarrow m + 1$.

end

7.1 Experiment 2: Complex customers' behaviors

In this section, we apply **EATF** to a more realistic scenario that adheres to our partner's business model and operations. According to §6.2, most orders are fulfilled within two weeks. Thus, we identify fulfillment options with two windows: a fast option (fulfilled by this week) and a slow one (fulfilled by next week). In addition, our partner offers two types of fulfillment services (home delivery or a self-pickup) and determines the surcharge based on the speed, the service, and the basket value. Therefore, we identify four fulfillment options: home delivery by this week (Option 1), home delivery by next week (Option 2), self-pickup by this week (Option 3), and self-pickup by next week (Option 4).

We calibrate three types of customers. High-value customers submit orders with the highest basket values, medium-value customers with medium basket values, and low-value customers with the lowest basket values. The basket values are determined using the average value of orders within specific percentiles: from the minimum to the 50th percentile for low-value customers, from the 50th to 80th percentile for medium-value customers, and from the 80th percentile to the maximum for high-value customers. The basket values for high-, medium-, and low-value customers are 249.70, 132.31, and 64.06 Canadian dollars, respectively. Additionally, each customer reveals their type as low, medium, or high with probabilities of 0.5, 0.3, and 0.2, respectively.

Next, we calibrate the fulfillment surcharge for each option. We aggregate the dataset to customer type-fulfillment option level and use the average surcharge. The displayed surcharge for each fulfillment option for each customer type is shown in Table 5. After that, we calibrate a basic attraction choice model for each customer type similar to §6.2.1 and summarize the attractions in Table 6. The rest of the calibration follows the baseline setup in §6.2.

We modify the distribution of customer types to create two additional cases: Case 2 has a larger proportion of high-value customers, and Case 3 has a larger proportion of medium-value customers. We compare **EATF** to the benchmark policies described in §6.1. The results are shown in Table 7.

The results underscore the consistent advantage of employing **EATF** across various market conditions. This is evident as it adapts well to different distributions of customer types, continually outperforming benchmark policies. This suggests that joint optimization of assortments and fulfillment decisions is crucial for maximizing profits, especially in markets with heterogeneous customer behaviors.

Table 5: Fulfillment surcharges for each option for each customer type

	Option 1	Option 2	Option 3	Option 4
Low	11.49	8.69	0	0
Medium	5.62	3.58	0	0
High	8.16	5.87	0	0

Table 6: Attraction of each option for each customer type

	Option 1	Option 2	Option 3	Option 4
Low	1.77	4.93	0.97	1.13
Medium	5.46	24.32	1.68	3.25
High	3.36	21.27	1.24	3.42

Table 7: Comparing EATF to benchmark policies in various markets

Case	$(\beta^{\text{low}}, \beta^{\text{mid}}, \beta^{\text{high}})$	ERT	OL	OF	SM
1	(0.5, 0.3, 0.2)	2.3%	2.8%	16.4%	18.5%
2	(0.2, 0.3, 0.5)	1.9%	2.0%	14.8%	16.5%
3	(0.3, 0.5, 0.2)	2.1%	2.4%	21.6%	23.8%

The results also indicate that our partner should prioritize demand management, as it significantly mitigates profit loss, as demonstrated by the **EAMF** policy. This suggests that if the retailer must choose between optimizing delivery assortment decisions and delivery fulfillment decisions, the former yields a larger benefit. With closer investigation, we observe that Options 3 and 4 are not included in any efficient assortment for any customer type. This is mainly because they are free of charge but require the same fulfillment capacity compared to Options 1 and 2. Only home delivery options are considered in **EATF** and **EAMF**. Due to the high attraction of the slow home delivery (Option 2) but limited fulfillment capacity, both **EATF** and **EAMF** show fast home delivery (Option 1) to most customers, resulting in little room to postpone fulfillment and balance capacity strategically.

8 Conclusion

In collaboration with one of the online grocery platforms in Canada, this paper investigates the fulfillment optimization problem with personalized delivery choices. A primary challenge for the online retailer is to decide which delivery options to offer, given customer heterogeneity for differentiated delivery options, and how to fulfill these orders with multiple logistic providers. To the best of our knowledge, our paper is among the first in the literature to address this problem.

To address this challenge, we developed an algorithm called **ERT** that runs in two phases. In the first (offline) phase, considering the customer heterogeneity for differentiated delivery options, **ERT** generates a set of efficient option lists that maximizes the expected revenue while achieving the best possible utilization of fulfillment resources. It then calculates the optimal certainty equivalent fulfillment decisions using the those lists. In the second (online) phase, **ERT** iterates between the personalization and fulfillment stages for the planning cycle. At the delivery personalization stage, **ERT** starts with an efficient option list and dynamically removes the economically nonviable options based on the customer's location and the current state of unfulfilled orders. At the fulfillment stage, the algorithm aggregates all the orders and assigns them to the appropriate warehouses according to the certainty equivalent plan as much as possible. **ERT** is proved to have a worst-case optimality gap that grows sub-linearly in the number of demand locations and demand arrival rate. This implies that **ERT** is asymptotically optimal as the customer network expands, which corresponds to a typical case faced by many online retailers, including our industrial partner.

We performed two series of simulations. First, by calibrating a choice model that adheres to Assumption 1, we illustrate the value of joint fulfillment and personalized delivery optimization. Second, we calibrate a more complex choice model that better describes the industrial partner's business operations. We demonstrate the consistency and robustness of our proposed heuristics across various scenarios.

Our numerical experiments reveal several managerial insights. First, while optimizing personalized fulfillment options and assignments are both crucial, their significance varies across different scenarios. Specifically, when customers prioritize faster fulfillment, demand management is more critical, whereas capacity management plays a minor role. Conversely, capacity management gains importance when customers prefer more cost-effective fulfillment options. Second, online retailers who serve larger areas and/or have more customers from each location benefit more from adopting the **ERT** policy as it is asymptotically optimal. Third, demand management takes precedence when fulfillment capacity is limited, but the focus shifts toward capacity management as capacity increases.

Finally, we would like to conclude by exploring potential avenues for future research. Firstly, our paper assumes that efficient option lists are constructed based on a known choice model. However, in reality, the parameters of a choice model need to be probed and learned jointly, along with the optimal lists and fulfillment decisions. This can be conducted within an online learning framework. Secondly, we assume flat-rate fulfillment costs for all logistic providers. Future research could extend this framework by considering a different fulfillment cost structure. For example, the fulfillment cost consists of fixed and variable parts, which incentivizes operations to ship a large number of orders to pickup locations to lower the average per-order cost. Indeed, our industrial partner is actively forming new pickup partnerships with local convenience stores, believing it is crucial to improve service further and reduce operational costs. We firmly believe that our paper spurs further research on the above extensions, offering valuable insights to the managers of online retailers in improving their customer satisfaction through delivery personalization while minimizing the cost of fulfillment.

A Proofs

A.1 Proposition 1

Before proving Proposition 1, we present a technical lemma and its proof.

Lemma 2. Let $f(x)$ be concave in x, and c is a vector of constant. Then, the following parametric optimization is concave in β

$$P(\beta) := \max_x f(x) + c^\top \beta \quad \text{s.t.} \quad Ax \leq \beta$$

Proof. Consider two parameters β_1 and β_2, and x_1^* and x_2^* are maximizers for $P(\beta_1)$ and $P(\beta_2)$, respectively. Let $\lambda \in [0, 1]$. Note that

$$A(\lambda x_1^*) + A((1-\lambda)x_2^*) \leq \lambda\beta_1 + (1-\lambda)\beta_2$$

which implies $x = \lambda x_1^* + (1-\lambda)x_2^*$ is feasible for $P(\lambda\beta_1 + (1-\lambda)\beta_2)$. Therefore

$$\begin{aligned}
P(\lambda\beta_1 + (1-\lambda)\beta_2) &\geq f(\lambda x_1^* + (1-\lambda)x_2^*) + c^\top(\lambda\beta_1 + (1-\lambda)\beta_2) \\
&\geq \lambda f(x_1^*) + (1-\lambda)f(x_2^*) + c^\top(\lambda\beta_1 + (1-\lambda)\beta_2) \\
&= \lambda(f(x_1^*) + c^\top\beta_1) + (1-\lambda)(f(x_2^*) + c^\top\beta_2) \\
&= \lambda P(\beta_1) + (1-\lambda)P(\beta_2)
\end{aligned}$$

The second inequality holds since f is a concave function. □

Proof of Proposition 1. Consider an arbitrary feasible joint strategy $\rho = (S^\rho, u^\rho)$ to **OPT**, where S^ρ is the assortment decisions and u^ρ is fulfillment assignments, and denote $D = (D_{k,j,\tau} \in \mathbb{Z} : k \in [K], j \in [N], \tau \in [T])$ a sample path of realized demand generated under S^ρ. We consider the following deterministic programming conditioning on D,

$$
\begin{aligned}
V(x \mid D) = \max_{\bar{u}} & \sum_\tau \sum_j \sum_k r_{k,j} D_{k,j,\tau} \\
& - \sum_\tau \left(\sum_i C_i(w^\top \bar{u}_{i,\tau}) + \sum_j b_j \left(\bar{y}_{1,j,\tau} - \sum_i \bar{u}_{1,j,i,\tau} \right) \right) \\
& - \sum_{k,j} b_j \bar{x}_{k,j,T+1} \\
\text{s.t.} \quad & \bar{x}_{k,j,1} = x_{k,j}, && \forall\, k, j \\
& \bar{y}_{k,j,\tau} = \bar{x}_{k,j,\tau} + D_{k,j,\tau}, && \forall\, k, j, \tau \\
& \bar{x}_{k,j,\tau+1} = \bar{y}_{k+1,j,\tau} - \sum_i \bar{u}_{k+1,j,i,\tau}, && \forall\, k < K, j, \tau \\
& \bar{x}_{K,j,\tau+1} = 0, && \forall\, j, \tau \\
& \sum_i \bar{u}_{k,j,i,\tau} \le \bar{y}_{k,j,\tau}, && \forall\, k, j, \tau \\
& \bar{u}_{k,j,i,\tau} \ge 0 && \forall\, k, j, i, \tau
\end{aligned}
$$

Note that $V(x \mid D)$ is a pure fulfillment optimization problem. We wish to show

$$
V^*(x) \le \mathbb{E}[V(x \mid D)] \le V(x \mid \mathbb{E}[D]) \le V^D(x)
$$

The first inequality. Since ρ is feasible to **OPT**, u^ρ must be a feasible fulfillment assignment when the realized demand is D. Hence, u^ρ is feasible to $V(x \mid D)$. Taking expectations yields that the expected profit under joint strategy ρ is upper bounded by $\mathbb{E}[V(x \mid D)]$. Since ρ is arbitrary, the relation must also hold for the optimal joint policy ρ^*.

The second inequality. we claim that $V(x \mid D)$ is concave in D. To show that, first note that the objective function in $V(x \mid D)$ is concave in x and that the constraints are linear. The concavity follows the technical Lemma 2. Then, by Jensen's inequality, we have $\mathbb{E}[V(x \mid D)] \le V(x \mid \mathbb{E}[D])$.

The last inequality. It is sufficient to show that we can construct a feasible assortment policy $\gamma_{j,\tau}(S)$ for $V^D(x)$ such that $\sum_S \Lambda_{j,\tau} \gamma_{j,\tau}(S) \pi_k(S) = \mathbb{E}[D_{k,j,\tau}]$ for all k, j, τ. Let $\bar{u}^*$ an optimal fulfillment assignment to $V(x \mid \mathbb{E}[D])$. Then $(\gamma_{j,\tau}(S), \bar{u}^*)$ is a feasible joint strategy for $V^D(x)$, implying $V^D(x)$ serves as an upper bound. The constructed assortment policy $\gamma_{j,\tau}(S)$ should solve the following problem for all j and τ:

$$
\begin{aligned}
\sum_S \gamma_{j,\tau}(S) \Lambda_{j,\tau} \pi_k(S) = \mathbb{E}[D_{k,j,\tau}], && \forall\, k \\
\sum_S \gamma_{j,\tau}(S) = 1 && \\
\gamma_{j,\tau}(S) \ge 0, && \forall\, S
\end{aligned}
$$

The above system of equations can be compressed as

$$A\gamma = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \Lambda_{j,\tau}\pi_1(S_1) & \Lambda_{j,\tau}\pi_1(S_2) & \cdots & \Lambda_{j,\tau}\pi_1(S_{|\mathcal{N}|}) \\ \vdots & \vdots & \vdots & \vdots \\ \Lambda_{j,\tau}\pi_K(S_1) & \Lambda_{j,\tau}\pi_K(S_2) & \cdots & \Lambda_{j,\tau}\pi_K(S_{|\mathcal{N}|}) \end{bmatrix} \begin{bmatrix} \gamma_{j,\tau}(S_1) \\ \gamma_{j,\tau}(S_2) \\ \vdots \\ \gamma_{j,\tau}(S_{|\mathcal{N}|}) \end{bmatrix} = \begin{bmatrix} 1 \\ \mathbb{E}[D_{1,j,\tau}] \\ \vdots \\ \mathbb{E}[D_{K,j,\tau}] \end{bmatrix}$$

It is easy to show that the matrix A has a full-row rank, which guarantees a solution of γ. For example, let $\mathcal{N} = (S_1, S_2, \cdots, S_{|\mathcal{N}|})$. Moreover, $S_1 = \varnothing, S_2 = \{1\}, S_3 = \{2\}, \cdots, S_{K+1} = \{K\}$. We can construct such a solution as follows: $\gamma_{j,\tau}(S_1) = 1 - \sum_{k=1}^{K} \frac{\mathbb{E}[D_{k,j,\tau}]}{\Lambda_{j,\tau}\pi_k(S_{k+1})}$, $\gamma_{j,\tau}(S_{k+1}) = \frac{\mathbb{E}[D_{k,j,\tau}]}{\Lambda_{j,\tau}\pi_k(S_{k+1})}$ for $k \in [K]$, and $\gamma_{j,\tau}(S) = 0$ otherwise. We still need to verify that $\gamma_{j,\tau}(S_1) \geq 0$. Indeed, it holds by the definition of $D_{k,j,\tau}$. □

A.2 Proposition 2

We present the necessary and sufficient condition for an option list to be efficient from Talluri and Van Ryzin (2004) as a technical lemma before proving Proposition 2. Note that for any given assortment strategy $S_{j,\tau}(t)$, the joint optimization $V^*(x)$ becomes a pure fulfillment optimization problem. Its optimal fulfillment decisions has the following structure (see Zhou et al. 2023, Proposition 2):

Lemma 3 (Optimal threshold fulfillment)**.** For any assortment strategy, if a warehouse i is ever used to fulfill any orders from location j during cycle τ, then the optimal fulfillment from that warehouse must satisfy $C_i'\left(w^\top u_{i,\tau}^*\right) w_j \leq b_j$.

Proof of Proposition 2. The proof consists of two steps: First, we show that if $\gamma_{j,\tau}^*(S) > 0$, then S solves the following problem:

$$\max_{S'} R_j(S') + \sum_k \theta_{k,j}\pi_k(S')$$

for some $\theta_j = (\theta_{1,j}, \cdots, \theta_{K,j})$. Next, we show that $\theta_{k,j}$ are identical and non-positive for all k. Then, by Lemma 1, S must be efficient.

Note that **DET** is feasible and bounded, implying it has a finite optimum with some $\gamma_{j,\tau}^*$ and $\overline{u}_{k,j,i,\tau}^*$. Moreover, strong duality holds since **DET** is a convex optimization with linear constraints.

Let $\theta = (\theta_{k,j,\tau} : k \in [K], j \in [N], \tau \in [T])$, $\eta = (\eta_{k,j,\tau} : k \in [K-1], j \in [N], \tau \in [T])$, $\xi = (\xi_{k,j,\tau} \geq 0 : k \in [K], j \in [N], \tau \in [T])$, $\sigma = (\sigma_{j,\tau} : j \in [N], \tau \in [T])$, $\phi = (\phi_{k,j,i,\tau} \geq 0 : k \in [K], j \in [N], i \in [M], \tau \in [T])$, and $\mu = (\mu_{j,\tau,S} \geq 0 : j \in [N], \tau \in [T], S \in \mathcal{N})$ be the dual variables for constraints (2b), (2c), (2e), (2f), (2g), and (2h), respectively. We indicate the optimal dual variables with superscript *.

The Lagrangian function can be formulated as

$$\begin{aligned}
&\mathcal{L}(\gamma, \overline{u}, \overline{x}, \overline{y}, \theta, \eta, \xi, \sigma, \phi, \mu) \\
=&\sum_{\tau,j} \sigma_{j,\tau} + \sum_{\tau,j,S} \left(\Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_k \theta_{k,j,\tau}\pi_k(S) - \sigma_{j,\tau} + \mu_{j,\tau,S} \right) \gamma_{j,\tau}(S) \\
&- \sum_{\tau,i} \left(C_i(w^\top \overline{u}_{i,\tau}) - \sum_j b_j \overline{u}_{1,j,i,\tau} + \sum_{k<K,j} \eta_{k,j,\tau}\overline{u}_{k+1,j,i,\tau} + \sum_{k,j} \xi_{k,j,\tau}\overline{u}_{k,j,i,\tau} - \sum_{k,j} \phi_{k,j,i,\tau}\overline{u}_{k,j,i,\tau} \right) \\
&- \sum_{\tau,j} \left(b_j \overline{y}_{1,j,\tau} + \sum_k \theta_{k,j,\tau}\overline{y}_{k,j,\tau} - \sum_{k<K} \eta_{k,j,\tau}\overline{y}_{k+1,j,\tau} - \sum_k \xi_{k,j,\tau}\overline{y}_{k,j,\tau} \right) \\
&+ \sum_{k,j,\tau} \theta_{k,j,\tau}\overline{x}_{k,j,\tau} - \sum_{k<K,j,\tau} \eta_{k,j,\tau}\overline{x}_{k,j,\tau+1} - \sum_{k,j} b_j \overline{x}_{k,j,T+1}
\end{aligned}$$

Treating $\bar{x}$ and $\bar{y}$ as primal decision variables, the dual problem can be formulated as follows.

$$\begin{aligned}
\mathcal{D} &= \min_{\theta,\eta,\xi,\sigma,\phi,\mu} \max_{\gamma,\bar{u},\bar{x},\bar{y}} \mathcal{L}(\gamma,\bar{u},\bar{x},\bar{y},\theta,\eta,\xi,\sigma,\phi,\mu) \\
&= \min_{\theta,\eta,\xi,\sigma,\phi,\mu} \Bigg\{ \sum_{\tau,j} \sigma_{j,\tau} + \underbrace{\max_{\gamma} \sum_{\tau,j,S} \left(\Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_k \theta_{k,j,\tau} \pi_k(S) - \sigma_{j,\tau} + \mu_{j,\tau,S} \right) \gamma_{j,\tau}(S)}_{\text{(I)}} \\
&\quad + \underbrace{\max_{\bar{u}} - \sum_{\tau,i} \left(C_i(w^\top \bar{u}_{i,\tau}) - \sum_j b_j \bar{u}_{1,j,i,\tau} + \sum_{k<K,j} \eta_{k,j,\tau} \bar{u}_{k+1,j,i,\tau} + \sum_{k,j} \xi_{k,j,\tau} \bar{u}_{k,j,i,\tau} - \sum_{k,j} \phi_{k,j,i,\tau} \bar{u}_{k,j,i,\tau} \right)}_{\text{(II)}} \\
&\quad + \underbrace{\max_{\bar{y}} - \sum_{\tau,j} \left(b_j \bar{y}_{1,j,\tau} + \sum_k \theta_{k,j,\tau} \bar{y}_{k,j,\tau} - \sum_{k<K} \eta_{k,j,\tau} \bar{y}_{k+1,j,\tau} - \sum_k \xi_{k,j,\tau} \bar{y}_{k,j,\tau} \right)}_{\text{(III)}} \\
&\quad + \underbrace{\max_{\bar{x}} \sum_{k,j,\tau} \theta_{k,j,\tau} \bar{x}_{k,j,\tau} - \sum_{k<K,j,\tau} \eta_{k,j,\tau} \bar{x}_{k,j,\tau+1} - \sum_{k,j} b_j \bar{x}_{k,j,T+1}}_{\text{(IV)}} \Bigg\}
\end{aligned}$$

We take the dual variables at their optimal values for the subsequent analysis.

Part One: For (I), for all τ, j, and S', according to first-order optimality condition, we must have

$$\Lambda_{j,\tau} R_j(S') + \Lambda_{j,\tau} \sum_k \theta^*_{k,j,\tau} \pi_k(S') - \sigma^*_{j,\tau} + \mu^*_{j,\tau,S} = 0$$

implying

$$\Lambda_{j,\tau} R_j(S') + \Lambda_{j,\tau} \sum_k \theta^*_{k,j,\tau} \pi_k(S') - \sigma^*_{j,\tau} = -\mu^*_{j,\tau,S} \leq 0 \tag{5}$$

By the complementary slackness condition $\mu^*_{j,\tau,S}\gamma^*_{j,\tau}(S) = 0$, if $\gamma^*_{j,\tau}(S) > 0$, we must have $\mu^*_{j,\tau,S} = 0$, implying

$$\Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_k \theta^*_{k,j,\tau} \pi_k(S) - \sigma^*_{j,\tau} = 0$$

Subtracting the above equation to (5) yields

$$R_j(S) + \sum_k \theta^*_{k,j,\tau} \pi_k(S) \geq R_j(S') + \sum_k \theta^*_{k,j,\tau} \pi_k(S')$$

which indicates that if $\gamma^*_{j,\tau}(S) > 0$ for some j and τ, then S is an maximizer for $\max_{S'} R_j(S') + \sum_k \theta^*_{k,j,\tau} \pi_k(S')$. This completes the first step.

Part two: we show $\theta^*_{k,j,\tau}$ are identical for all k and that $\theta^*_{k,j,\tau} \leq 0$. To do so, we first identify the relationships among dual variables. We start with (III).

$$\text{(III)} = \max_{\bar{y}} \sum_{\tau,j} (-b_j - \theta^*_{1,j,\tau} + \xi^*_{1,j,\tau}) \bar{y}_{1,j,\tau} + \sum_{\tau,j,k>1} (-\theta^*_{k,j,\tau} + \eta^*_{k-1,j,\tau} + \xi^*_{k,j,\tau}) \bar{y}_{k,j,\tau}$$

The first-order optimality condition implies

$$-b_j - \theta^*_{1,j,\tau} + \xi^*_{1,j,\tau} = 0, \qquad \forall\, j, \tau \tag{6}$$

$$-\theta^*_{k,j,\tau} + \eta^*_{k-1,j,\tau} + \xi^*_{k,j,\tau} = 0, \qquad \forall\, j, \tau, k > 1 \tag{7}$$

Next, we examine $\textcircled{\text{IV}}$. Since $\bar{x}_1$ is given and $\bar{x}_{K,j,\tau+1} = 0$ for all τ, we have the following

$$\textcircled{\text{IV}} = \sum_{k,j} \theta_{k,j,1}\bar{x}_{k,j,1} + \max_{\bar{x}_2,\cdots,\bar{x}_{T+1}} \left\{ \sum_{\tau>1} \sum_{k<K,j} (\theta^*_{k,j,\tau} - \eta^*_{k,j,\tau-1})\bar{x}_{k,j,\tau} + \sum_{k<K,j} (-b_j - \eta^*_{k,j,T})\bar{x}_{k,j,T+1} \right\}$$

which yields the following relations:

$$\eta^*_{k,j,T} = -b_j \qquad \forall\, k < K, j \tag{8}$$

$$\theta^*_{k,j,\tau} = \eta^*_{k,j,\tau-1} \qquad \forall\, k < K, j, \tau > 1 \tag{9}$$

(7) and (9) implies

$$\theta^*_{k,j,\tau} = \theta^*_{k-1,j,\tau+1} + \xi_{k,j,\tau} \tag{10}$$

(6), (7), and (8) imply

$$\theta^*_{k,j,T} = -b_j + \xi^*_{k,j,T} \geq -b_j$$

for all k and j. The inequality holds since $\xi^*_{k,j,\tau} \geq 0$. By backward induction with (10) and the relation in (6), we can further show that, for all k, j, and τ

$$\theta^*_{k,j,\tau} \geq -b_j \tag{11}$$

Lastly, we apply the first-order optimality condition to $\textcircled{\text{II}}$. For all τ and i,

$$C'_i(w^\top \bar{u}^*_{i,\tau})w_j - b_j + \xi^*_{1,j,\tau} - \phi^*_{1,j,i,\tau} = 0$$

$$C'_i(w^\top \bar{u}^*_{i,\tau})w_j + \eta^*_{k-1,j,\tau} + \xi^*_{k,j,\tau} - \phi^*_{k,j,i,\tau} = 0, \qquad \forall\, k > 1$$

Using (6) and (7), above are equivalent to, for all k, j, i, τ,

$$C'_i(w^\top \bar{u}^*_{i,\tau})w_j + \theta^*_{k,j,\tau} - \phi^*_{k,j,i,\tau} = 0 \tag{12}$$

Using (11) and $\phi^*_{k,j,i,\tau} \geq 0$, we conclude that $\theta^*_{k,j,\tau}$ satisfies the follow constraint.

$$\theta^*_{k,j,\tau} \geq -\min\{b_j, \min_i C'_i(w^\top \bar{u}^*_{i,\tau})w_j\} \tag{13}$$

In addition, the complementary slackness condition $\bar{u}^*_{k,j,i,\tau}\phi^*_{k,j,i,\tau} = 0$ implies

$$\bar{u}^*_{k,j,i,\tau}\left(C'_i(w^\top \bar{u}^*_{i,\tau})w_j + \theta^*_{k,j,\tau}\right) = 0 \tag{14}$$

We may interpret $-\theta^*_{k,j,\tau}$ as the marginal value for fulfilling an order index by k, j, and τ. When $C'_i(w^\top \bar{u}^*_{i,\tau})w_j - (-\theta^*_{k,j,\tau}) > 0$, implying the weighted marginal fulfillment cost of warehouse i is higher than the marginal benefit of fulfillment, we must not fulfill such orders from warehouse i, i.e., $\bar{u}^*_{k,j,i,\tau} = 0$. On the other hand, when $C'_i(w^\top \bar{u}^*_{i,\tau})w_j - (-\theta^*_{k,j,\tau}) = 0$, we may fulfill those orders from warehouse i, i.e., $\bar{u}^*_{k,j,i,\tau} \geq 0$.

Next, we will show that $\theta^*_{k,j,\tau}$ are identical and negative for all k. We arbitrarily fix the index j and τ. We need to consider two cases.

Case 1: when some orders from j are fulfilled. We define an index set $\mathcal{K}_{j,i,\tau} = \{k \in [K] : \bar{u}^*_{k,j,i,\tau} > 0\}$. Case 1 implies $\cup_i \mathcal{K}_{j,i,\tau} \neq \varnothing$. We consider the following three sub-cases.

Case 1.1: Orders with different fulfillment window are fulfilled by the *same* warehouse. That is, if there exists an index i such that for all $k, k' \in \mathcal{K}_{j,i,\tau}$, $k \neq k'$,

$$\theta^*_{k,j,\tau} = -C'_i(w^\top \bar{u}^*_{i,\tau})w_j = \theta^*_{k',j,\tau}$$

The relationship holds according to (14). The above relationship indicates that if orders with different windows are fulfilled from the same warehouse, they have the same marginal value, which equals the negative weighted marginal fulfillment cost of that fulfilled warehouse.

Case 1.2: Orders with different fulfillment window are fulfilled by *different* warehouses. That is, if there exists i, i', k, and k' such that $i \neq i'$, $k \in \mathcal{K}_{j,i,\tau}$, $k' \in \mathcal{K}_{j,i',\tau}$, $k \neq k'$,

$$\theta^*_{k,j,\tau} \geq -C'_{i'}(w^\top \overline{u}^*_{i',\tau})w_j = \theta^*_{k',j,\tau}$$
$$\theta^*_{k',j,\tau} \geq -C'_{i}(w^\top \overline{u}^*_{i,\tau})w_j = \theta^*_{k,j,\tau}$$

The inequality follows (13), and the equality follows (14). This implies

$$\theta^*_{k,j,\tau} = -C'_{i'}(w^\top \overline{u}^*_{i',\tau})w_j = -C'_{i}(w^\top \overline{u}^*_{i,\tau})w_j = \theta^*_{k',j,\tau}$$

Cases 1 and 2 imply that if orders with fulfillment window k are ever fulfilled, they have the same marginal value. In other words, $\theta^*_{k,j,\tau}$ takes the same value for all $k \in \cup_i \mathcal{K}_{j,i,\tau}$. Moreover, any warehouse that fulfills to location j during cycle τ has the same weighted marginal fulfillment cost. To simplify the notation for the following discussion, we denote those warehouses by an index set $\mathcal{I}_{j,\tau} = \{i \in [M] : \sum_{k'} \overline{u}^*_{k',j,i,\tau} > 0\}$. We denote this weighted marginal fulfillment cost as C'. That is, $C' = C'_i(w^\top \overline{u}^*_{i,\tau})w_j$ for all $i \in \mathcal{I}_{j,\tau}$ and $\theta^*_{k,j,\tau} = -C'$ for all $k \in \cup_i \mathcal{K}_{j,i,\tau}$.

Case 1.3: Orders with fulfillment window k are not fulfilled (i.e., exists some k such that $k \notin \cup_i \mathcal{K}_{j,i,\tau}$). We assume there exists at least one $k \notin \cup_i \mathcal{K}_{j,i,\tau}$ such that $\theta^*_{k,j,\tau} > -C'$ since otherwise all $\theta^*_{k',j,\tau}$ are identical for all k'. Our goal is to construct another set of dual variables $\hat{\theta}$, $\hat{\eta}$, $\hat{\xi}$, $\hat{\sigma}$, and $\hat{\phi}$, and show that $(\gamma^*, \overline{u}^*, \overline{x}^*, \overline{y}^*, \hat{\theta}, \hat{\eta}, \hat{\xi}, \hat{\sigma}, \hat{\phi}, \mu^*)$ satisfy Karush–Kuhn–Tucker (KKT) conditions.

Denote $\bar{\mathcal{K}}_{j,\tau} = \{k \notin \cup_i \mathcal{K}_{j,i,\tau} : \theta^*_{k,j,\tau} > -C'\}$. Define $\delta_k := \theta^*_{k,j,\tau} + C'$ for all $k \in \bar{\mathcal{K}}_{j,\tau}$. Let $\hat{\theta}$ equals to θ^* everywhere except $\hat{\theta}_{k,j,\tau} = -C'$ for all $k \in \bar{\mathcal{K}}_{j,\tau}$; $\hat{\eta}$ equals η^* everywhere except $\hat{\eta}_{k-1,j,\tau} = \eta^*_{k-1,j,\tau} - \delta_k$ for all $k \in \bar{\mathcal{K}}_{j,\tau}$ and $k \neq 1$, $\hat{\eta}_{k,j,\tau-1} = \eta^*_{k,j,\tau-1} - \delta$ for all $k \in \bar{\mathcal{K}}_{j,\tau}$ and $\tau > 1$; $\hat{\xi}$ equals to ξ^* everywhere except $\hat{\xi}_{1,j,\tau} = b_j - C'$ if $\{1\} \subset \bar{\mathcal{K}}_{j,\tau}$; $\hat{\sigma}$ equals σ^* everywhere except $\hat{\sigma}_{j,\tau} = \sigma^*_{j,\tau} - \Lambda_{j,\tau} \sum_{k \in \bar{\mathcal{K}}_{j,\tau}} \pi_k(S)\delta_k$; $\hat{\phi}$ equals to ϕ^* everywhere except $\hat{\phi}_{k,j,i,\tau} = \phi^*_{k,j,i,\tau} - \delta_k$ for all i and $k \in \bar{\mathcal{K}}_{j,\tau}$.

First, we verify the dual feasibility conditions. It is sufficient only to show that $\hat{\phi}_{k,j,i,\tau} \geq 0$ for all i and $k \in \bar{\mathcal{K}}_{j,\tau}$ and $\hat{\xi}_{1,j,\tau} \geq 0$ if $\{1\} \subset \bar{\mathcal{K}}_{j,\tau}$. The latter must hold according to Lemma 3. For the formal, according to (12), for all $k \in \bar{\mathcal{K}}_{j,\tau}$,

$$\begin{aligned}\hat{\phi}_{k,j,i,\tau} &= \phi^*_{k,j,i,\tau} - \delta_k \\ &= C'_i(w^\top \overline{u}^*_{i,\tau})w_j + \theta^*_{k,j,\tau} - \theta^*_{k,j,\tau} - C' \\ &= \begin{cases} 0 & \textit{if} \quad i \in \mathcal{I}_{j,\tau} \\ C'_i(w^\top \overline{u}^*_{i,\tau})w_j - C' & \textit{otherwise} \end{cases}\end{aligned}$$

when $i \in \mathcal{I}_{j,\tau}$, the last equality holds according to the definition of C'. For $i \notin \mathcal{I}_{j,\tau}$, we can show $C'_i(w^\top \overline{u}^*_{i,\tau})w_j \geq C'$ by contradiction. Suppose the relationship does not hold. For any $i' \in \mathcal{I}_{j,\tau}$, there must exist some k' such that $\overline{u}^*_{k',j,i',\tau} > 0$. Then, we can increase the fulfillment $\overline{u}^*_{k',j,i,\tau}$ and decrease $\overline{u}^*_{k',j,i',\tau}$ by a small amount. This means we re-assign some fulfillment with a fulfillment window of k' from warehouse i' to i, leading to a reduction in the fulfillment cost since, according to the assumption, $C'_i(w^\top \overline{u}^*_{i,\tau})w_j < C' = C'_{i'}(w^\top \overline{u}^*_{i',\tau})w_j$. This contradicts the fact that $\overline{u}^*$ is optimal. Therefore, we conclude $\hat{\phi}_{k,j,i,\tau} \geq 0$ for all i and $k \in \bar{\mathcal{K}}_{j,\tau}$.

Next, we verify the stationarity conditions $\nabla \mathcal{L} = 0$ with respect to the primal variables. Specifically, $\gamma_{j,\tau}(S)$,

$$\frac{\partial \mathcal{L}}{\partial \gamma_{j,\tau}(S)} = \Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_{k'} \hat{\theta}_{k',j,\tau} \pi_{k'}(S) - \hat{\sigma}_{j,\tau} + \hat{\mu}_{j,\tau,S}$$

$$
\begin{aligned}
&= \Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_{k' \notin \bar{\mathcal{K}}_{j,\tau}} \theta^*_{k',j,\tau} \pi_{k'}(S) + \Lambda_{j,\tau} \sum_{k \in \bar{\mathcal{K}}_{j,\tau}} \pi_k(S)(\theta^*_{k,j,\tau} - \delta_k) \\
&\quad - \sigma^*_{j,\tau} + \Lambda_{j,\tau} \sum_{k \in \bar{\mathcal{K}}_{j,\tau}} \pi_k(S)\delta_k + \mu^*_{j,\tau,S} \\
&= \Lambda_{j,\tau} R_j(S) + \Lambda_{j,\tau} \sum_{k'} \theta^*_{k',j,\tau} \pi_{k'}(S) - \sigma^*_{j,\tau} + \mu^*_{j,\tau,S} \\
&= 0
\end{aligned}
$$

For $\bar{u}_{k,j,i,\tau}$ such that $k \in \bar{\mathcal{K}}_{j,\tau}$ and $k > 1$,

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \bar{u}_{k,j,i,\tau}} &= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j - \hat{\eta}_{k-1,j,\tau} - \xi^*_{k,j,\tau} + \hat{\phi}_{k,j,i,\tau} \\
&= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j - \eta^*_{k-1,j,\tau} + \delta - \xi^*_{k,j,\tau} + \phi^*_{k,j,i,\tau} - \delta \\
&= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j - \eta^*_{k-1,j,\tau} - \xi^*_{k,j,\tau} + \phi^*_{k,j,i,\tau} \\
&= 0
\end{aligned}
$$

when $k \in \bar{\mathcal{K}}_{j,\tau}$ and $k = 1$,

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \bar{u}_{1,j,i,\tau}} &= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j + b_j - \hat{\xi}_{1,j,\tau} + \hat{\phi}_{1,j,i,\tau} \\
&= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j + b_j - b_j + C' + \phi^*_{1,j,i,\tau} - \theta^*_{1,j,\tau} - C' \\
&= -C_i'(w^\top \bar{u}^*_{i,\tau}) w_j + \phi^*_{1,j,i,\tau} - \theta^*_{1,j,\tau} \\
&= 0
\end{aligned}
$$

where the last equality follows (12). For $\bar{y}_{k,j,\tau}$ and $\bar{x}_{k,j,\tau}$, $k \in \bar{\mathcal{K}}_{j,\tau}$,

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \bar{y}_{k,j,\tau}} &= \hat{\theta}_{k,j,\tau} - \hat{\eta}_{k-1,j,\tau} - \xi^*_{k,j,\tau} \\
&= -C' - (\eta^*_{k-1,j,\tau} - \theta^*_{k,j,\tau} - C') - \xi^*_{k,j,\tau} \\
&= \theta^*_{k,j,\tau} - \eta^*_{k-1,j,\tau} - \xi^*_{k,j,\tau} \\
&= 0 \\
\frac{\partial \mathcal{L}}{\partial \bar{x}_{k,j,\tau}} &= \hat{\theta}_{k,j,\tau} - \hat{\eta}_{k,j,\tau-1} \\
&= -C' - (\eta^*_{k,j,\tau-1} - \theta^*_{k,j,\tau} - C') \\
&= \theta^*_{k,j,\tau} - \eta^*_{k,j,\tau-1} \\
&= 0
\end{aligned}
$$

Note that $\frac{\partial \mathcal{L}}{\partial \bar{x}_{k,j,\tau}} = 0$ is required to hold only for $\tau > 1$ since $\bar{x}_{k,j,1} = x_{k,j}$ is not a decision variable.

Lastly, regarding the complementary slackness conditions, since $\hat{\theta}$, $\hat{\eta}$, and $\hat{\sigma}$ correspond to equality constraints, we only have to verify $\hat{\phi}_{k,j,i,\tau} \bar{u}^*_{k,j,i,\tau} = 0$ for all i and $k \in \bar{\mathcal{K}}_{j,\tau}$ and $\hat{\xi}_{1,j,\tau}(\sum_i \bar{u}^*_{1,j,i,\tau} - \bar{y}^*_{1,j,\tau}) = 0$ if $\{1\} \subset \bar{\mathcal{K}}_{j,\tau}$. The former holds since $\bar{u}^*_{k,j,i,\tau} = 0$ according to the assumption. The latter holds since the only scenario when $\{1\} \subset \bar{\mathcal{K}}_{j,\tau}$ is when $\bar{y}^*_{1,j,\tau} = 0$. To see that, suppose $\bar{y}^*_{1,j,\tau} > 0$, which indicates, according to the assumption, it is optimal to fulfill some orders with $k' > 1$ but not with $k = 1$. However, we can always find some i such that $\bar{u}^*_{k',j,i,\tau} > 0$. Then, we increase $\bar{u}^*_{1,j,i,\tau}$ by a small amount and reduce $\bar{u}^*_{k',j,i,\tau}$ by the same amount. The resulting fulfillment decision achieves a lower fulfillment cost according to Lemma 3, which contradicts that $\bar{u}^*$ is optimal.

Therefore, $(\gamma^*, \bar{u}^*, \bar{x}^*, \bar{y}^*)$ and $(\hat{\theta}, \hat{\eta}, \xi^*, \hat{\sigma}, \hat{\phi}, \mu^*)$ jointly satisfy KKT conditions. To show the constructed dual variables are optimal, we still need Slater's condition to hold, which is clearly the case.

For instance, we can find a relative interior and feasible primal point by letting $\gamma_{j,\tau}(S) = 1/|\mathcal{N}|$ for all $S \in \mathcal{N}$ and equally assigning $(1-\varepsilon)$ proportion of unfulfilled orders to each warehouse, for any $\varepsilon \in (0,1)$. Thus, we conclude $\hat{\theta}$ is also an optimal dual and $\hat{\theta}_{k,j,\tau} = -C'$ for all k.

Case 2: no order from j is fulfilled. Once again, our goal is to construct dual variables $(\hat{\theta}, \hat{\eta}, \hat{\xi}, \hat{\sigma}, \hat{\phi}, \mu^*)$, and show that, together with primal optimal variables $(\gamma^*, \overline{u}^*, \overline{x}^*, \overline{y}^*)$, they satisfy KKT conditions.

Denote $C' := \min\{b_j, \min_i C_i'(w^\top \overline{u}^*_{i,\tau})w_j\}$. We set $\hat{\theta}_{k,j,\tau} = -C'$ for all k; $\hat{\eta}$ equals to η^* everywhere except $\hat{\eta}_{k-1,j,\tau} = \eta^*_{k-1,j,\tau} - \theta^*_{k,j,\tau} - C'$ for all $k > 1$ and $\hat{\eta}_{k,j,\tau-1} = \eta^*_{k,j,\tau-1} - \theta^*_{k,j,\tau} - C'$ for $\tau > 1$; $\hat{\xi}$ equals to ξ^* everywhere except $\hat{\xi}_{1,j,\tau} = b_j - C'$; $\hat{\sigma}$ equals to σ^* everywhere except $\hat{\sigma}_{j,\tau} = \sigma^*_{j,\tau} - \Lambda_{j,\tau}\sum_{k'}(\theta^*_{k,j,\tau} + C')\pi_{k'}(S)$; $\hat{\phi}$ equals to ϕ^* everywhere except $\hat{\phi}_{k,j,i,\tau} = \phi^*_{k,j,i,\tau} - \theta^*_{k,j,\tau} - C'$ for all i.

For feasibility conditions, we need to show $\hat{\xi}_{1,j,\tau} \geq 0$ and $\hat{\phi}_{k,j,i,\tau} \geq 0$ for all i. The first holds since $\hat{\xi}_{1,j,\tau} = \max\{0, b_j - \min_i C_i'(w^\top \overline{u}^*_{i,\tau})w_j\} \geq 0$. The second holds since, according to (12),

$$
\begin{aligned}
\hat{\phi}_{k,j,i,\tau} &= \phi^*_{k,j,i,\tau} - \theta^*_{k,j,\tau} - C' \\
&= C_i'(w^\top \overline{u}^*_{i,\tau})w_j + \theta^*_{k,j,\tau} - \theta^*_{k,j,\tau} - C' \\
&= C_i'(w^\top \overline{u}^*_{i,\tau})w_j - \min\{b_j, \min_{i'} C_{i'}'(w^\top \overline{u}^*_{i',\tau})w_j\} \\
&= \begin{cases} C_i'(w^\top \overline{u}^*_{i,\tau})w_j - b_j & \textit{if } b_j \leq \min_{i'} C_{i'}'(w^\top \overline{u}^*_{i',\tau})w_j \\ C_i'(w^\top \overline{u}^*_{i,\tau})w_j - \min_{i'} C_{i'}'(w^\top \overline{u}^*_{i',\tau})w_j & \textit{otherwise} \end{cases} \\
&\geq 0
\end{aligned}
$$

For the stationarity conditions $\nabla\mathcal{L} = 0$, similar to the previous discussion in Case 1.3, we can show our constructed dual variables satisfy $\frac{\partial\mathcal{L}}{\partial\gamma_{j,\tau}(S)} = 0$, $\frac{\partial\mathcal{L}}{\partial\overline{u}_{k,j,i,\tau}} = 0$ when $k > 1$, $\frac{\partial\mathcal{L}}{\partial\overline{y}_{k,j,\tau}} = 0$ when $k > 1$, and $\frac{\partial\mathcal{L}}{\partial\overline{x}_{k,j,\tau}} = 0$ for $\tau > 1$. For $k = 1$,

$$
\begin{aligned}
\frac{\partial\mathcal{L}}{\partial\overline{u}_{1,j,i,\tau}} &= -C_i'(w^\top \overline{u}_{i,\tau})w_j + b_j - \hat{\xi}_{1,j,\tau} + \hat{\phi}_{1,j,i,\tau} \\
&= -C_i'(w^\top \overline{u}_{i,\tau})w_j + b_j - (b_j - C') + (\phi^*_{1,j,i,\tau} - \theta^*_{1,j,\tau} - C') \\
&= -C_i'(w^\top \overline{u}_{i,\tau})w_j - \theta^*_{1,j,\tau} + \phi^*_{1,j,i,\tau} \\
&= 0
\end{aligned}
$$

where the last equality follows (12), and

$$
\begin{aligned}
\frac{\partial\mathcal{L}}{\partial\overline{y}_{1,j,\tau}} &= -b_j - \hat{\theta}_{1,j,\tau} + \hat{\xi}_{1,j,\tau} \\
&= -b_j + C' + b - C' \\
&= 0
\end{aligned}
$$

For complementary slackness conditions, it is sufficient only to verify $\hat{\xi}_{1,j,\tau}\left(\sum_i \overline{u}^*_{1,j,i,\tau} - \overline{y}^*_{1,j,\tau}\right) = 0$. If $\overline{y}^*_{1,j,\tau} = 0$, the condition holds since $\sum_i \overline{u}^*_{1,j,i,\tau} - \overline{y}^*_{1,j,\tau} = 0$. Otherwise, if $\overline{y}^*_{1,j,\tau} > 0$, we need to show $\hat{\xi}_{1,j,\tau} = 0$, which is equivalent to show $b_j \leq C_i'(w^\top \overline{u}^*_{i,\tau})w_j$ for all i. This is indeed the case: If there exists a warehouse i that does not satisfy that requirement, we are better off by fulfilling some orders with $k = 1$ from that warehouse to location j, i.e., $\overline{u}^*_{1,j,i,\tau} > 0$, which contradicts that it is optimal to not fulfill all orders from j.

Therefore, we conclude that $(\gamma^*, \overline{u}^*, \overline{x}^*, \overline{y}^*)$ and $(\hat{\theta}, \hat{\eta}, \hat{\xi}, \hat{\sigma})$ satisfy KKT conditions. As discussed previously, Slater's condition also holds. Thus, $\hat{\theta}_{k,j,\tau} = -\min\{b_j, \min_i C_i'(w^\top \overline{u}^*_{i,\tau})w_j\} < 0$ for all k is also an optimal dual.

Since the index j and τ are arbitrarily fixed, the above completes the second part. □

A.3 Theorem 1

Before providing the proof for Theorem 1, we first give the formal definition of the **HYP** policy in Algorithm 3 and two technical lemmas and their proof. The proof of the theorem directly comes from the lemmas.

Algorithm 3: Hypothetical Policy (HYP)

Input: Customer arrival rate vector $\Lambda_{j,\tau}$, $\forall j \in [N], \tau \in [T]$, Fulfillment cost vector C_i, $\forall i \in [M]$, penalty vector b_j, $\forall j \in [N]$
Output: Fulfillment option display decision $S^{\text{HYP}} = \{S^{\text{HYP}}_{j,\tau}\}$ and fulfillment decision $u^{\text{HYP}} = \{u^{\text{HYP}}_{k,j,i,t}\}$
1 Construct the set of all *efficient assortments* $\mathcal{E}_j$ for all $j \in [N]$.
2 Solve **DET** to obtain $\gamma^*_{j,\tau}(S)$ for all $S \in \mathcal{E}_j$ and $\bar{u}^*_{k,j,i,\tau}$.
3 **for** $\tau \in [T]$ **do**
4 At any time $t \in [0, 1]$, if there is a customer from location j, offer an assortment $S \in \mathcal{E}$ with probability $\gamma^*_{j,\tau}(S)$.
5 At the end of the cycle, fulfill orders according to a scaling rule $u^{\text{HYP}}_{k,j,i,t} = \frac{y_{k,j,\tau}}{\bar{y}^*_{k,j,\tau}} \cdot \bar{u}^*_{k,j,i,\tau}$.
6 **end**

Lemma 4. $V^{\text{HYP}}(x) \le V^{\text{EATF}}(x)$

Proof. Consider a sequence of customer arrivals, denoted as $\tilde{D} = (\tilde{D}_{j,\tau}(t) : j \in [N], \tau \in [T], t \in [0, 1])$. This process can also be represented by the specific times of arrival. For each fulfillment cycle τ, let the total number of customers be $L_\tau := \sum_j \tilde{D}_{j,\tau}$. These arrivals are indexed by $l \in [1, L_\tau]$ with corresponding arrival times $t_{l,\tau}$. Each order is identified by its arrival time.

Applying both **EATF** and **HYP** to $\tilde{D}$, the total profit from each policy can be obtained by the cumulative profit from each customer, termed as marginal profit, which can be calculated as the marginal revenue subtracting the marginal fulfillment cost. That is,

$$V^\rho(x \mid \tilde{D}) = \sum_{\tau=1}^{T} \sum_{l=1}^{L_\tau} r^\rho_{l,\tau} - c^\rho_{l,\tau}$$

where $\rho \in \{\text{EATF}, \text{HYP}\}$, $r^\rho_{l,\tau}$ and $c^\rho_{l,\tau}$ are the marginal revenue and cost for the orders arriving at time $t_{l,\tau}$ during cycle τ under policy ρ. The expected profit $V^{\text{EATF}}(x)$ and $V^{\text{HYP}}(x)$ can be obtained by taking the expectation of the profit with respect to the sample path.

$$V^\rho(x) = \mathbb{E}_{\tilde{D}}\left[V^\rho(x \mid \tilde{D})\right]$$

Therefore, to prove this lemma, it is sufficient to show that

$$r^{\text{EATF}}_{l,\tau} - c^{\text{EATF}}_{l,\tau} \ge r^{\text{HYP}}_{l,\tau} - c^{\text{HYP}}_{l,\tau}$$

for all l and τ.

To control the randomness due to customer choice, we assume that if the same assortment is offered under both **EATF** and **HYP** to the customer at time $t_{l,\tau}$, she will choose the same fulfillment option under both policies.

If both policies offer identical assortments for all customer $t_{l,\tau}$ across all τ, the lemma holds trivially. Therefore, we assume, for the first time, **EATF** and **HYP** offer different assortments during cycle τ_0 to the customer at t_{l_0,τ_0} from location j_0. Denote the offered assortments by S^{HYP} and S^{EATF}, respectively, with $S^{\text{EATF}} \subsetneq S^{\text{HYP}}$.

Let us focus on this customer at t_{l_0,τ_0}. If she chooses to leave without the purchase under both policies, $r^\rho_{l_0,\tau_0} = c^\rho_{l_0,\tau_0} = 0$. Otherwise, if she chooses the same fulfillment option $k \in S^{\text{EATF}}$ under both policies, then the marginal revenue will be identical, i.e., $r^{\rho}_{l_0,\tau_0} = r_{k,j}$. However, the marginal fulfillment cost associated with this order under **EATF** will not be larger than that under **HYP**. To show that, we first notice that this order will not be fulfilled in a later cycle under **HYP** than under **EATF** according to Algorithm 1. If this order is fulfilled in the same cycle under both policies, then the marginal fulfillment cost can be considered equal under both policies. However, if this order is fulfilled in a later cycle under **EATF** than under **HYP**, then, according to Algorithm 1, $c^{\text{HYP}}_{l_0,\tau_0} > b_{j_0} \geq c^{\text{EATF}}_{l_0,\tau_0}$.

Conversely, if the customer chooses $k' \in S^{\text{HYP}} \setminus S^{\text{EATF}}$ under **HYP** but $k \in S^{\text{EATF}} \cup \{\varnothing\}$ under **EATF**. Since k' is eliminated in **EATF**, it implies k' is economically nonviable while k is viable. According to Definition 2, the marginal profit with k' is strictly negative, whereas, with k, it is non-negative. This rationale applies to any customers from location j_0 after t_{l_0,τ_0} during cycle τ_0, as k' remains nonviable. Therefore, we can conclude that **EATF** will not have worse performance than **HYP** *by the end of* cycle τ.

For the subsequent cycle $\tau_1 := \tau_0+1$, while the pre-assortment state $x^{\rho}_{\tau_1}$ may differ under **EATF** and **HYP**, this does not impact the assortment strategy S^{HYP} for cycle τ_1. Nor does it affect S^{EATF} until the time instance t_{l_1,τ_1}, when some fulfillment option $k' \in S^{\text{HYP}}$ becomes economically nonviable. Before this time, both policies offer the same assortment. After t_{l_1,τ_1}, k' is removed from S^{EATF}. Using the same logic as before, we can demonstrate that the profit collected under **EATF** during the fulfillment cycle τ_1 will not be less than that under **HYP**. □

Lemma 5. Let $\Lambda = \max_{j,\tau} \Lambda_{j,\tau}$, $\overline{w} = \max_j w_j$. The optimality gap of HYP is bounded by

$$V^*(x) - V^{\text{HYP}}(x) \leq \overline{w}\overline{C}MTK\sqrt{\frac{1}{2}\Lambda KN}$$

Proof. Let $(D^{\text{HYP}}_{k,j,\tau} : k \in [K], j \in [N], \tau \in [T])$ denote the realized demand under HYP. Note that $D^{\text{HYP}}_{k,j,\tau}$ is a counting process with the parameter $\Lambda_{j,\tau} \sum_S \gamma^*_{j,\tau}(S)\pi_k(S)$ over a unit time span, hence,

$$\mathbb{E}\left[D^{\text{HYP}}_{k,j,\tau}\right] = \Lambda_{j,\tau} \sum_S \gamma^*_{j,\tau}(S)\pi_k(S)$$

The expected revenue collected during some fulfillment cycle τ can be calculated by

$$\mathbb{E}\left[\sum_j \sum_k r_{k,j} D^{\text{HYP}}_{k,j,\tau}\right] = \sum_{k,j,S} r_{k,j}\Lambda_{j,\tau}\gamma^*_{j,\tau}(S)\pi_k(S)$$

which equals the revenue in **DET** with an optimal solution. Hence, the difference between the expected profit under **HYP** and **OPT** can be bounded as follows:

$$\begin{aligned} V^*(x) - V^{\text{HYP}}(x) &\leq V^D(x) - V^{\text{HYP}}(x) \\ &= \mathbb{E}\left\{\sum_{\tau=1}^{T}(g(y_\tau, u^{\text{HYP}}_\tau) - g(\overline{y}^*_\tau, \overline{u}^*_\tau)) + \sum_{k,j} b_j(x_{k,j,T+1} - \overline{x}^*_{k,j,T+1})\right\} \end{aligned}$$

where g is the immediate cost function defined as

$$g(y,u) = \sum_i C_i\left(\sum_{k,j} w_j u_{k,j,i}\right) + \sum_j b_j\left(y_{1,j} - \sum_i u_{1,j,i}\right)$$

x_τ and y_τ are the pre- and post-assortment state under **HYP**, respectively. $\overline{x}^*_\tau$ and $\overline{y}^*_\tau$ are the pre- and post-assortment state along the "optimal fluid path".

We first show that the system's evolution under **HYP** follows the "optimal fluid path" in expectation. Specifically, we want to show that

$$\mathbb{E}[x_\tau] = \bar{x}^*_\tau, \quad \mathbb{E}[y_\tau] = \bar{y}^*_\tau, \quad \mathbb{E}[u^{\text{HYP}}_\tau] = \bar{u}^*_\tau$$

for all τ. We can show this by induction. When $\tau = 1$

$$\begin{aligned}\mathbb{E}[y_{k,j,1}] &= \mathbb{E}\left[x_{k,j,1} + D^{\text{HYP}}_{k,j,1}\right] \\ &= x_{k,j,1} + \mathbb{E}[D^{\text{HYP}}_{k,j,1}] \\ &= x_{k,j,1} + \Lambda_{j,1}\sum_S \gamma^*_{j,1}(S)\pi_k(S) \\ &= \bar{y}^*_{k,j,1}\end{aligned}$$

and $\mathbb{E}[u^{\text{HYP}}_{k,j,i,1}] = \bar{u}^*_{k,j,i,1}$ follows the definition in Algorithm 3. Since $x_{k,j,2}$ is linear in both $y_{k,j,1}$ and $u^{\text{HYP}}_{k,j,1}$, then we must have $\mathbb{E}[x_{k,j,2}] = \bar{x}^*_{k,j,2}$.

Suppose for any $\tau > 1$ we have $\mathbb{E}[x_\tau] = \bar{x}^*_\tau$, $\mathbb{E}[y_{\tau-1}] = \bar{y}^*_{\tau-1}$, and $\mathbb{E}[u^{\text{HYP}}_{\tau-1}] = \bar{u}^*_{\tau-1}$ by the inductive assumption. We can easily show the following holds:

$$\mathbb{E}[y_\tau] = \bar{y}^*_\tau, \quad \mathbb{E}[u^{\text{HYP}}_\tau] = \bar{u}^*_\tau, \quad \mathbb{E}[x_{\tau+1}] = \bar{x}^*_{\tau+1}$$

which completes the induction. Therefore,

$$\begin{aligned}V^*(x) - V^{\text{HYP}}(x) &\le \mathbb{E}\left\{\sum_{\tau=1}^{T}(g(y_\tau, u^{\text{HYP}}_\tau) - g(\bar{y}^*_\tau, \bar{u}^*_\tau)) + \sum_{k,j} b_j(x_{k,j,T+1} - \bar{x}^*_{k,j,T+1})\right\} \\ &= \mathbb{E}\left\{\sum_{\tau=1}^{T}\sum_{i=1}^{M}\left(C_i(w^\top u^{\text{HYP}}_{i,\tau}) - C_i(w^\top \bar{u}^*_{i,\tau})\right)\right\} \\ &\le \sum_{\tau=1}^{T}\sum_{i=1}^{M}\bar{C}\mathbb{E}\left\{\left|w^\top u^{\text{HYP}}_{i,\tau} - w^\top \bar{u}^*_{i,\tau}\right|\right\} \\ &\le \sum_{\tau=1}^{T}\sum_{i=1}^{M}\bar{C}\sqrt{Var\left(\sum_{k,j} w_j u^{\text{HYP}}_{k,j,i,\tau}\right)}\end{aligned}$$

The second inequality holds due to Mean Value Theorem and Assumption 2. The last inequality is held by Jensen's inequality and the definition of variance. To further derive the performance guarantee regarding the problem parameters, we must carefully examine the relation of the fulfillment from the warehouse i to different demand indexes k and j. Note that customers' arrival process and the assortment display decision are independent in j. Hence, the post-assortment state $y_{k,j,\tau}$ is independent in j. Since $u^{\text{HYP}}_{k,j,i,t} = \frac{\bar{u}^*_{k,j,i,\tau}}{\bar{y}^*_{k,j,\tau}} y_{k,j,\tau}$, where $\frac{\bar{u}^*_{k,j,i,\tau}}{\bar{y}^*_{k,j,\tau}}$ is deterministic, the randomness of $u^{\text{HYP}}_{k,j,i,\tau}$ completely comes from $y_{k,j,\tau}$. Therefore, $u^{\text{HYP}}_{k,j,i,\tau}$ is independent in j. Therefore

$$\begin{aligned}Var\left(\sum_{k,j} w_j u^{\text{HYP}}_{k,j,i,\tau}\right) &= \sum_j w_j^2 Var\left(\sum_k u^{\text{HYP}}_{k,j,i,\tau}\right) \\ &\le \sum_j \bar{w}^2 K \sum_k Var\left(u^{\text{HYP}}_{k,j,i,\tau}\right)\end{aligned} \tag{15}$$

Moreover, the deterministic part $\frac{\bar{u}^*_{k,j,i,\tau}}{\bar{y}^*_{k,j,\tau}} \le 1$ since $\sum_i \bar{u}^*_{k,j,i,\tau} \le \bar{y}^*_{k,j,\tau}$ and $\bar{u}^*_{k,j,i,\tau} \ge 0$. Therefore, $Var\left(u^{\text{HYP}}_{k,j,i,\tau}\right) \le Var\left(y_{k,j,\tau}\right)$. Then, we bound the variance of $y_{k,j,\tau}$. By the system dynamics, we have

$$y_{k,j,\tau} \le D^{\text{HYP}}_{k,j,\tau} + D^{\text{HYP}}_{k+1,j,\tau-1} + \cdots + D^{\text{HYP}}_{K,j,\tau-K+k}$$

where each term $D^{\text{HYP}}_{k,j,\tau} = \sum_S \gamma^*_{j,\tau}(S)\pi_k(S)\tilde{D}_{j,\tau}$ on the right-hand side can be considered as a splitting Poisson process. This implies $Var(D^{\text{HYP}}_{k,j,\tau}) \le Var(\tilde{D}_{j,\tau}) = \Lambda_{j,\tau}$. Moreover, all terms are independent of each other. Therefore,

$$Var\left(y_{k,j,\tau}\right) \le \Lambda_{j,\tau} + \Lambda_{j,\tau-1} + \cdots + \Lambda_{\tau-K+k} \le (K-k+1)\Lambda_j \le K\Lambda_j$$

where $\Lambda_j = \max_\tau \Lambda_{j,\tau}$. Substituting the above result to (15) yields

$$Var\left(\sum_{k,j} u^{\text{sLR}}_{k,j,i,\tau}\right) \le \sum_j \overline{w}^2 K \sum_k K\Lambda_j \le \frac{1}{2}\overline{w}^2 K^3 N\Lambda$$

where $\Lambda := \max_{j,\tau} \Lambda_{j,\tau}$. Therefore,

$$V^*(x) - V^{\text{HYP}}(x) \le \overline{w}\,\overline{C}MTK\sqrt{\frac{1}{2}\Lambda KN} \qquad \square$$

A.4 Proposition 4

Proof. We omit the customer type index q and location index j in this proof. At each recursion, if there is a single maximizer for the marginal revenue ratio, then Algorithm 2 is identical to the original one proposed by Talluri and Van Ryzin (2004), and it returns the complete sequence of efficient assortments. Therefore, it is sufficient to show that Algorithm 2 can find all assortments that are all maximizers for the marginal revenue ratio at some recursion.

Suppose that Θ contains two assortments, S^1 and S^2, $S^1 \neq S^2$, at m-th recursion. We first consider the case $Q(S^1) \neq Q(S^2)$. Without loss of generality, we assume $Q(S^1) < Q(S^2)$. Clearly, $R(S^1) < R(S^2)$; otherwise, both cannot be maximizers. Denote the maximum marginal revenue ratio at m-th recursion by $\theta = \max_{S' \in \Omega(S_m)} \frac{R(S)-R(S_m)}{Q(S)-Q(S_m)}$, θ can be infinite.

According to the algorithm, we select $S_{m+1} = S^1$. In the $(m+1)$-st recursion, the marginal revenue ratio for assortment S^2 is

$$\begin{aligned}
\frac{R(S^2)-R(S_{m+1})}{Q(S^2)-Q(S_{m+1})} &= \frac{R(S^2)-R(S_m)+R(S_m)-R(S_{m+1})}{Q(S^2)-Q(S_{m+1})} \\
&= \frac{\theta(Q(S^2)-Q(S_m))+\theta(Q(S_m)-Q(S_{m+1}))}{Q(S^2)-Q(S_{m+1})} \\
&= \frac{\theta(Q(S^2)-Q(S_{m+1}))}{Q(S^2)-Q(S_{m+1})} \\
&= \theta
\end{aligned}$$

The second equality holds since $\theta = \frac{R(S^1)-R(S_m)}{Q(S^1)-Q(S_m)} = \frac{R(S^2)-R(S_m)}{Q(S^2)-Q(S_m)}$ according to the assumption and the fact that $S_{m+1} = S^1$.

According to (4), $\Omega(S_{m+1}) \subset \Omega(S_m)$. Therefore, for any other assortment $S \in \Omega(S_{m+1})$, S must also in $\Omega(S_m)$. The marginal revenue ratio for S is

$$\begin{aligned}
\frac{R(S)-R(S_{m+1})}{Q(S)-Q(S_{m+1})} &= \frac{R(S)-R(S_m)+R(S_m)-R(S_{m+1})}{Q(S)-Q(S_{m+1})} \\
&< \frac{\theta(Q(S)-Q(S_m))+\theta(Q(S_m)-Q(S_{m+1}))}{Q(S)-Q(S_{m+1})} \\
&= \frac{\theta(Q(S)-Q(S_{m+1}))}{Q(S^2)-Q(S_{m+1})} \\
&= \theta
\end{aligned}$$

The inequality holds since $S \in \Omega(S_m)$ but not a maximizer, which implies $\frac{R(S)-R(S_m)}{Q(S)-Q(S_m)} < \theta$. We selected S^1 at m-th recursion, and S^2 will be the sole maximizer in the $(m+1)$-st recursion. Therefore, both will be added to $\mathcal{E}$.

When $Q(S^1) = Q(S^2)$, we must also have $R(S^1) = R(S^2)$. In this case, we arbitrarily select an assortment (e.g., S^1) at the m-th recursion. Define $0/0 = \infty$. Then, for the next recursion, the marginal revenue ratio will be infinite for S^2 but strictly smaller than θ for the rest $S \in \Omega(S_{m+1})$. Hence, S^2 will also be added to $\mathcal{E}$.

Suppose there are multiple maximizers $\{S^1, S^2, \cdots, S^n\}$ at m-th recursion and S^1 is selected according to Algorithm 2. In that case, by the same argument as before, we can show that $\{S^2, \cdots, S^n\}$ will be maximizers for $(m+1)$-st recursion. We can continue the procedure until all of them are added to $\mathcal{E}$. □

B Additional numerical studies

In what follows, we present additional numerical studies to explore the performance of EATF and benchmark policies by varying arrival rates, number of warehouses, and the number of delivery options.

B.1 Varying arrival rates

Our theoretical performance bound implies both policies are asymptotically optimal when the arrival rate of each location is large. To test that, we change the arrival rate of all locations from 10 to 60 and adjust the in-contract fulfillment quantity so that the fulfillment capacity remains unchanged. Similar to the study of the number of locations, we consider two types of markets: customers prefer the balanced fulfillment option in the actual markets but the cost-efficient option in the virtual markets. We report the simulation results in Table 8. It can be seen that the relative gaps for **EATF** approach to zero and $\lambda_{j,\tau}$ increases, which support our theoretical performance guarantee. However, this is not seen in the benchmark policies, implying that asymptotical optimality is a unique benefit of joint optimization. In addition, it is clear that **EATF** outperforms all benchmark policies in all cases, similar to what we observed when we changed N.

Table 8: Performance of EATF and benchmark policies under various arrival rates.

	Actual Market				Virtual Market			
$\lambda_{j,\tau}$	EATF	EAMF	FATF	FAMF	EATF	EAMF	FATF	FAMF
10	3.7%	5.6%	21.5%	26.1%	1.9%	31.8%	68.0%	68.3%
20	2.7%	3.8%	21.7%	24.8%	1.2%	18.1%	67.6%	68.3%
30	2.2%	3.2%	21.5%	24.1%	1.2%	15.2%	67.1%	68.4%
40	2.0%	2.7%	21.5%	23.8%	1.1%	13.3%	66.8%	68.2%
50	1.8%	2.5%	21.6%	23.6%	1.0%	9.8%	66.5%	68.3%
60	1.5%	2.2%	21.6%	23.3%	0.7%	9.0%	66.4%	68.3%

B.2 Varying number of warehouses

We change the number of warehouses from 1 to 3 but keep the system-wide fulfillment capacity constant. To do so, we set the cheaper fulfillment quantity for each warehouse by $220/M$. The results are reported in Table 9. Unlike what is suggested in Theorem 1, the simulation results show that the optimality gap does not increase as M increases. This result indicates that **EATF** applies to retailers with multiple warehouses.

Table 9: Performance of EATF and benchmark policies with different number of warehouses.

M	EATF	EAMF	FATF	FAMF
1	2.6%	3.9%	21.7%	24.9%
2	2.7%	4.0%	21.4%	24.7%
3	2.7%	4.0%	21.7%	24.9%

References

Acimovic J, Farias VF (2019) The fulfillment-optimization problem. Operations Research & Management Science in the age of analytics, 218–237, publisher: INFORMS.

Acimovic J, Graves SC (2015) Making better fulfillment decisions on the fly in an online retail environment. Manufacturing & Service Operations Management 17(1):34–51, publisher: INFORMS.

Andrews JM, Farias VF, Khojandi AI, Yan CM (2019) Primal–Dual Algorithms for Order Fulfillment at Urban Outfitters, Inc. INFORMS Journal on Applied Analytics 49(5):355–370, publisher: INFORMS.

Cao Y, Kleywegt AJ, Wang H (2022) Network revenue management under a spiked multinomial logit choice model. Operations Research 70(4):2237–2253, publisher: INFORMS.

Cheung WC, Simchi-Levi D (2016) Efficiency and performance guarantees for choice-based network revenue management problems with flexible products.

Digital Commerce 360 (2023) The shopper speaks: Shipping and delivery satisfaction numbers stand strong. URL https://www.digitalcommerce360.com/2023/08/30/shipping-delivery-satisfaction/.

eMarketer (2023) Retail ecommerce sales worldwide, 2021-2027. URL https://www.emarketer.com/chart/263530/retail-ecommerce-sales-worldwide-2021-2027.

Gallego G, Iyengar G, Phillips R, Dubey A (2004) Managing flexible products on a network.

Gallego G, Li A (2017) Attention, consideration then selection choice model. Consideration then Selection Choice Model (September 24, 2017) .

Gallego G, Phillips R (2004) Revenue management of flexible products. Manufacturing & Service Operations Management 6(4):321–337, publisher: INFORMS.

Gallego G, Topaloglu H (2019) Revenue management and pricing analytics, volume 209 (Springer).

Grand View Research (2022) Global online food delivery market size & share report, 2030. URL https://www.grandviewresearch.com/industry-analysis/online-food-delivery-market-report.

Gönsch J, Koch S, Steinhardt C (2014) Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches. International Journal of Production Economics 153:280–294, publisher: Elsevier.

Harsha P, Subramanian S, Uichanco J (2019) Dynamic pricing of omnichannel inventories: honorable mention—2017 M&SOM Practice-Based Research Competition. Manufacturing & Service Operations Management 21(1):47–65.

Jasin S, Sinha A (2015) An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. Operations Research 63(6):1336–1351, publisher: INFORMS.

Kök AG, Fisher ML, Vaidyanathan R (2009) Assortment planning: Review of literature and industry practice. Retail supply chain management: Quantitative models and empirical studies 99–153, publisher: Springer.

Lei Y, Jasin S, Sinha A (2018) Joint dynamic pricing and order fulfillment for e-commerce retailers. Manufacturing & Service Operations Management 20(2):269–284.

Lei Y, Jasin S, Uichanco J, Vakhutinsky A (2022) Joint product framing (display, ranking, pricing) and order fulfillment under the multinomial logit model for e-commerce retailers. Manufacturing & Service Operations Management 24(3):1529–1546.

Liu Q, Van Ryzin G (2008) On the choice-based linear programming model for network revenue management. Manufacturing & Service Operations Management 10(2):288–310, publisher: INFORMS.

Ma Y, Rusmevichientong P, Sumida M, Topaloglu H (2020) An approximation algorithm for network revenue management under nonstationary arrivals. Operations Research 68(3):834–855, publisher: INFORMS.

Mahar S, Wright PD (2009) The value of postponing online fulfillment decisions in multi-channel retail/e-tail organizations. Computers & operations research 36(11):3061–3072.

Qi M, Mak HY, Shen ZJM (2020) Data-driven research in retail operations—A review. Naval Research Logistics (NRL) 67(8):595–616, publisher: Wiley Online Library.

Talluri K, Van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. Management Science 50(1):15–33, publisher: INFORMS.

Wei L, Kapuscinski R, Jasin S (2021) Shipping consolidation across two warehouses with delivery deadline and expedited options for e-commerce and omni-channel retailers. Manufacturing & Service Operations Management 23(6):1634–1650.

Wolfe A (2024) Top retail fulfillment challenges and strategies. URL https://cart.com/blog/retail-fulfillment-challenges-strategies.

Xie Y, Ma W, Xin L (2023) The benefits of delay to online decision-making.

Xu PJ, Allgor R, Graves SC (2009) Benefits of reevaluating real-time order fulfillment decisions. Manufacturing & Service Operations Management 11(2):340–355.

Zhou Q, Gumus M, Miao S (2023) E-commerce order fulfillment problem with limited time window.

Zhu W, Topaloglu H (2023) Performance guarantees for network revenue management with flexible products. Manufacturing & Service Operations Management Publisher: INFORMS.