Essays on E-Commerce and Omnichannel Retail Operations

by

Aravind Govindarajan

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in the University of Michigan
2019

Doctoral Committee:

Assistant Professor Joline Uichanco, Chair
Professor Hyun-soo Ahn
Assistant Professor Eric. M. Schwartz
Assistant Professor Cong Shi
Principal Scientist Amitabh Sinha

Aravind Govindarajan

arav@umich.edu

ORCID iD: 0000-0003-4658-3523

Acknowledgments

I am extremely grateful to my advisors for their guidance and patience during my PhD – to Joline Uichanco for her constant support and encouragement during difficult times, and helping me grow as a researcher; to Amitabh Sinha for always being there to offer advice and mentorship, and for opening new doors for my career; and finally to Hyun-Soo Ahn for graciously taking me under his wing, and inspiring me with his bottomless excitement for research. I also thank my thesis committee members – Eric Schwartz and Cong Shi, for their comments that significantly improved this dissertation. I am extremely grateful to Professor C. Rajendran for his inspiration and encouragement that made me apply for a PhD. I am also thankful to the faulty members at the Ross School of Business, in particular, Stefanus Jasin and Roman Kapuscinski, for their endless support throughout my PhD.

I gratefully acknowledge that my dissertation was partially supported by the National Science Foundation under Grant Number 1561791.

I cannot begin to express my thanks to Zhaohui (Zoey) Jiang, my cohort, best friend, and confidante, for her unrelenting belief in me during the toughest times, and for helping me grow as a person. I'm extremely grateful to my seniors, especially Yixin (Iris) Wang, who I'm fortunate to call a close friend that I can count on, and Guihua Wang, whose help and advice I value a lot. Burdens at work felt lighter with Eun Woo Kim, Frank Li and Rebecca Chae, discussing topics ranging from time travel and alternate universes, to winning lotteries and becoming YouTube celebrities. My roommates, Madan Ravi Ganesh and Sakthi Manokar Dhandapani, were always there by my side during my happiest and saddest times, and helped create a home far away from home. My friends from IIT, Shankar Sridhar and Vijayakumar Ganapathy, have made me realize that real friendship defies distance and time. Last but not the least, I would like to thank my family for their unconditional love and support, without which none of this would have been possible. My circle of love has grown bigger with my newborn nephew Adhav, who is going to have the best uncle in the world.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

ABSTRACT

The advent of e-commerce has impacted the retail industry, as retail firms have innovated in response to customers increasingly preferring to purchase products online. This dissertation studies operational problems that accompany such retail innovations, and provides tractable heuristic solutions developed using stochastic and robust optimization methods. In particular, the first two chapters focus on the value of fulfillment flexibility – online orders can be fulfilled from any node in the firm's fulfillment network. The first chapter is devoted to omnichannel retailing, where e-commerce demand is integrated with the physical network of stores through ship-from-store fulfillment. For a retailer with a network of physical stores and fulfillment centers facing two demands (online and in-store), we consider the following interlinked decisions – how much inventory to keep at each location and where to fulfill each online order from. We show that the value of considering fulfillment flexibility in inventory planning is highest when there is a moderate mix of online and in-store demands, and develop computationally fast heuristics with promising asymptotic performance for large scale networks, which are shown to improve upon traditional strategies.

The second chapter considers a pure play e-commerce fulfillment network, and studies the inventory placement decision. As e-commerce demands are volatile due to a variety of factors (price-matching, recommendation engines, etc.), we consider a distributionally robust setting, where the objective is to minimize the worst-case expected cost under given mean and covariance matrices of the underlying demand distribution. For this NP-hard problem, we develop computationally tractable heuristic in the form of a semi-definite program, with dimension quadratic in the size of the

network. In the face of distribution uncertainty, we show that the robust heuristic outperforms inventory solutions that assume incorrect distributions.

The final chapter offers a new take on a classic problem in retail – customer returns, which has grown to be an important issue in recent times with firms competing to provide lenient and convenient return policies to boost their e-commerce sales. However, several customers take advantage of such policies, which can lead to loss in revenue and increase in inventory costs. We study different return policies that a firm can employ depending on the information about customers' return behavior that is available to the firm. We derive the structure of the optimal return policies and show that personalizing return policies based on customers' historical data can significantly improve the firm's profits, but allows the firm to extract all customer surplus.

CHAPTER 1

Introduction

Internet has changed how humans interact and transact with businesses. Customers are increasingly preferring to conduct their shopping online, which provides new challenges to retail firms to reinvent supply chain strategies in the digital era.

There are several operational problems that need to be addressed by firms which face online demands – inventory placement (where to position inventory in the network), fulfillment decisions (where to fulfill an incoming online order), assortment (online and in-store, online recommendations), pricing (dynamic pricing, price-matching with competitors, differentiation between online and in-store), returns (free returns or partial returns, leniency in return window), etc. While most of these problems have been studied in literature in the context of brick-and-mortar retail, e-commerce demand introduces additional flexibility and challenges. In particular, from the point of view of the online shopper, policies (pricing, assortment, return policy, etc.) can now be personalized based on a customer's historical data collected by the firm; from the point of view of the firm, online orders can be fulfilled from any node (stores, fulfillment centers, etc.) in the network, which can be used as a strategic lever to manage operations across the network.

In this thesis, we look at three such important problems faced by modern retail firms. The first two chapters deal with inventory and fulfillment decisions that are brought about by the *fulfillment flexibility* in dealing with online demand – an online order can be fulfilled from any node in the fulfillment network. The third chapter focuses on the value of personalizing return policies for customers based on their historical return behavior.

In the first chapter, we study the problem of inventory and fulfillment optimization for omnichannel retail firms. Omnichannel refers to the seamless integration of a retailer's sales channels, such as in-store and online. While this integration is motivated by giving flexibility to customers, it leads to pooling of demands within and across locations. Thus, such integration can lead to reduction in cost that can be achieved through efficient inventory management. To this end, we consider a retailer with a network of physical stores and fulfillment centers facing two demands (online and in-store), where online demand can be fulfilled from any location with available inventory. We model the setting as a stochastic optimization problem, by considering order-up-to policies for a general multi-period model with multiple locations and zero lead time, and online orders fulfilled multiple times in each period. We develop a simple, scalable heuristic for the multi-location problem based on analysis from the two-store problem, for the special case where online orders are only fulfilled at the end of a period. For the case where fulfillment is done dynamically, we develop a simple threshold-based policy which reserves inventory at stores for future in-store demand. We then employ a realistic numerical study to analyze the benefits of using the combined inventory and fulfillment heuristic over traditional decentralized and myopic strategies.

In the second chapter, we consider the inventory placement problem in e-commerce fulfillment centers through a distributionally robust approach. In network inventory planning, the joint distribution of the random demands is needed to optimize inventory levels at each node in the network. However, in the case of e-commerce demands, the exact distribution may be inaccessible due to high volatility in online customer behavior arising from factors such as competition, the use of dynamic price-matching strategies and flash promotions, recommendation engines that manipulate click-streams, etc. Assuming that the firm knows only the mean and covariance matrices, we solve a distributionally robust multi-location newsvendor model for network inventory optimization. The objective is to minimize the worst-case expected cost over the set of demand distributions satisfying the known mean and covariance information. For the special case of two homogeneous customer locations with correlated demands, we show that a six-point distribution achieves the worst-case expected cost, and derive a closed-form expression for the optimal inventory decision. The general multi-location problem can be shown to be NP-hard. We develop a computationally tractable upper bound through the solution of a semidefinite program (SDP), which also yields heuristic inventory levels, for a special class of fulfillment cost structures, namely nested fulfillment structures. We also develop an algorithm to convert any general distance-based fulfillment cost structure into a nested fulfillment structure which tightly approximates the expected total fulfillment cost.

In the third chapter, we consider the important problem of managing customer return policies. With lenient return policies growing popular in recent times, several customers take advantage of such policies. Retail firms keep track of customer return behavior through the data they collect themselves or through third-party companies – recently, Amazon has banned several customers who were considered to be fraudulent returners. We study how the firm can use information about customers' return behavior to tailor personalized return policies. For heterogeneous customers who differ in their perceived hassle cost of returns (which can be thought of as a proxy for return rates under lenient return policies), we derive the firm's optimal return policy, which consists of two components: 1) a return window (short or long), and 2) a refund fee. The firm benefits from a shorter return window, as returned items are less likely to be damaged and more likely to be resold during the selling season (modeled by a higher salvage price), whereas customers are inconvenienced by shorter windows (modeled by increase in their return hassle). When the firm offers full refunds to returning customers, consistent with Amazon's practice, we show that low-hassle customers must be banned from returning. However, when the firm is allowed to personalized return fees, we show that the firm benefits from selling to these low-hassle customers under strict return policies (short window and high return fees). Identifying and targeting customers based on their historical return behavior can lead to significant increase in profits, however, we show that customer surplus is wiped out. This provides implications for customers' privacy in retail settings, and the value of consumer behavior data.

CHAPTER 2

Joint Inventory and Fulfillment Decisions for Omnichannel Retail Networks

2.1 Abstract

With e-commerce growing at a rapid pace compared to traditional retail, many brick-and-mortar firms are supporting their online growth through an integrated omnichannel approach. Such integration can lead to reduction in cost that can be achieved through efficient inventory management. A retailer with a network of physical stores and fulfillment centers facing two demands (online and in-store) has to make important, interlinked decisions – how much inventory to keep at each location and where to fulfill each online order from, as online demand can be fulfilled from any location. We consider order-up-to policies for a general multi-period model with multiple locations and zero lead time, and online orders fulfilled multiple times in each period. We first focus on the case where fulfillment decisions are made at the end of each period, which allows separate focus on the inventory decision. We develop a simple, scalable heuristic for the multi-location problem based on analysis from the two-store case, and prove its asymptotic near-optimality for large number of omnichannel stores under certain conditions. We extend this to the case where fulfillment can be done multiple times within a period and combine it with a simple, threshold-based fulfillment policy which reserves inventory at stores for future in-store demand. With the help of a realistic numerical study based on a fictitious retail network embedded in mainland USA, we show that the combined heuristic outperforms a myopic, decentralized planning strategy under a variety of problem parameters, especially when there is an adequate mix of online and in-store demands. Extensions to positive lead times are discussed.

2.2 Introduction

By the end of 2016, e-commerce sales accounted for around 9% of the total retail sales in the United States (U.S. Census Bureau, 2016). Although this is a small portion of the total sales, online sales have been increasing at a rapid growth rate of around 16% each year (Zaroban, 2018), and projected to account for 17% of all retail sales within the next five years (Lindner, 2017). In comparison, the growth in traditional retail has dwindled to around 2% in recent years. With customers increasingly favoring the online channel, traditional brick-and-mortar (B&M) firms are compelled to develop their e-commerce capabilities to remain competitive against pure play e-commerce firms like Amazon (Leiser, 2016), which alone accounted for 53% of the e-commerce sales growth in 2016 (Kim, 2017). In order to improve efficiency and flexibility, retailers resort to an omnichannel approach to integrate the online channel with their physical stores.

Omnichannel refers to the seamless integration of a retailer's sales channels, such as in-store and online. Customers can purchase an item in different ways, including placing an order through the online store (websites), through mobile devices (mobile apps), as well as through the traditional practice of walking into physical stores. In addition, customers placing orders online can also choose how they receive the item, which has led to various omnichannel initiatives: they can pick up their items from a nearby physical store (in-store pickup) or from designated self-service kiosks like Amazon Lockers, or simply have the item shipped directly to their homes (ship-to-customer).

Providing an omnichannel customer experience is regarded as a brand differentiator by many retailers, and integrating the online channel with the physical stores increases revenue, reduces shipping costs and improves customer satisfaction (Forrester, 2014). Hence, there is an industry-wide shift to omnichannel retailing, with onetime B&M firms like Macy's and Walmart leveraging their existing network of retail stores in their integration of the online channel (Nash, 2015). Amazon has also joined these firms through the acquisition of a network of physical stores across the US by means of its purchase of Whole Foods Market. This allows Amazon to not only operate an omnichannel grocery chain, but also absorb the stores into its distribution network to reduce logistic costs.

One of the key aspects of this channel integration is *store fulfillment*, which is the use of physical stores to fulfill online orders. Store fulfillment has now become indispensable for firms like Walmart and Macy's, that rely on a network of physical stores close to population centers to offer same day and next-day delivery options to customers (Giannopoulos, 2014). Dedicated floor space and store staff are required to fulfill online orders from stores.

In spite of potential benefits, many firms have struggled in their implementation of channel integration: from 2010 to 2014, even as retail and online sales increased, inventory turnover decreased (Kurt Salmon, 2016). One possible cause for this inefficiency could be insufficient planning in inventory management. While firms have traditionally managed inventory levels at stores based on demands in the corresponding locations, such a decentralized approach ceases be optimal in an integrated system.

The optimal inventory decisions depends on the fulfillment policy followed, and there does not seem to be a standard approach to online fulfillment across the industry. Some firms primarily fulfill from online fulfillment centers (FCs), and resort to store fulfillment in case the online FC runs out of stock. Some firms fulfill online orders from stores, but are agnostic to store inventory levels, while others do not fulfill from stores running low on inventory.

In this paper, we study the problem of an omnichannel firm with a network of physical stores and online FCs facing online (ship-to-customer) and in-store demands, by means of a general multi-period, multi-location model. We consider a dynamic setting, where we allow online fulfillment decisions to be made multiple times within each period. Online orders can be routed to any store or online FC in the network, and items are picked off the shelves, packed, labeled and shipped to the customers' homes. This has several advantages over the dedicated use of online FCs including reduced shipping costs, quicker deliveries and efficient use of store inventory (UPS Compass, 2014).

Our goal is to *optimize inventory levels and fulfillment decisions* for a single product. The decisions have to be made based on the network as a whole as opposed to a decentralized approach, in order to take into account demand pooling of online demands across the network, in addition to demand pooling of in-store and online demands in each region.

The firms's problem is described as follows. A retail firm owns a network of stores and online FCs, and has integrated the online channel into the physical stores through store fulfillment. Following a periodic review inventory model, each store orders up to a certain level at the beginning of each review period, to fulfill in-store demand (customers walking into physical stores) and online demand (customers ordering online, expecting items to be shipped directly to them) during the course of the period. The in-store demand at a store is fulfilled as it arrives, until that store runs out of inventory.

Unlike in-store demand, online demand can be fulfilled from any location in the network, and there is typically a delay between the time an order is placed and when items are picked off the shelf. Firms may delay fulfillment decisions due to various reasons:

- for strategic reasons, orders from the same customer or region can be consolidated to lower shipping costs (Xu, Allgor, and Graves, 2009; Wei, Jasin, and Kapuscinski, 2017),
- or for practical reasons, as the timing of orders fulfilled from stores is affected by store staffing schedules and pick-up times of third-party carriers like UPS.

To model this dynamic, a review period is further divided into T fulfillment epochs, where in-store demands are fulfilled as they arrive, and online fulfillment decisions (assigning online orders to fulfillment locations) are made at the end of each epoch after observing the demands during the epoch, with unmet demands being lost. The inventory and fulfillment decisions are made centrally by the firm to minimize holding, penalty and shipping costs. For the sake of clarity, the two units of time are described below:

- a *review period* is the amount of time between two consecutive inventory replenishments. For stores that are replenished daily, the review period is a single day.
- a *fulfillment epoch* is the time between two fulfillment decisions. Over the course of an epoch, online orders are aggregated, and fulfillment decisions are made at the end of each epoch. For stores replenished daily, the length of an epoch can range from a whole day (e.g. Macy's stores fulfill online orders once a day through UPS (Lewis, 2013)) to a few minutes (e.g. firms like Amazon make more frequent fulfillment decisions).

As described, the definition of a fulfillment epoch carries flexibility, and by choosing large enough values for T, we can closely approximate the continuous time setting, where firms make fulfillment decisions as online orders arrive.

The online fulfillment decisions are similar to transshipment decisions for online demand, except that instead of items being shipped between stores, they are shipped directly to the customer. The setting can thus be cast as planning of order-up-to levels in a transshipment problem with a replenishment leadtime of $T-1$ periods, with a planning horizon of T periods. This makes the problem hard, as it has been shown that optimal transshipment decisions are intractable, let alone joint optimization of initial inventory levels and transshipment decisions, even for two locations (Tagaras and Cohen, 1992).

The general structure of the problem is also subject to complications from other sources - multiple locations, multiple fulfillment epochs, and two non-identical classes of demands. Our main contribution is a combined inventory and fulfillment heuristic for omnichannel retailing, which we derive from a general multi-location, multi-period model shown to be mathematically intractable due to the various generalizations involved. Specifically, the inventory heuristic calculates the stocking levels at each location based on the demands in the network, rather than individually at that location, and the fulfillment heuristic provides location-specific, time-varying inventory thresholds which dictate the rationing between in-store and online demands.

The strength of our combined heuristic lies in the ease of computation and comprehension, and we show by means of a realistic numerical study that our heuristic creates value by planning for virtual pooling of online demands across locations, and diligently reserving inventory at stores for future demands. Our solutions are generalizable, and offer a framework to build further complexities on, which can yield valuable decision support tools for firms.

The approach we take to address this problem is as follows. We model the general problem in Section 2.4, and describe the complexities involved. To obtain a heuristic solution to this problem, we first decouple the inventory and fulfillment decisions by considering the case with a single fulfillment epoch ($T = 1$) in Section 2.5. When there is no leadtime, a myopic fulfillment policy would be optimal in this case - fulfill online demand as much as possible with the available inventory in each review period. Given this fulfillment policy, we discuss the optimal inventory solution for the two-store case, and develop a simple, asymptotically near-optimal inventory heuristic for the multi-location case.

In Section 2.6, we extend this inventory heuristic to the general problem where online orders are fulfilled multiple times within each review period ($T > 1$), and develop a simple threshold fulfillment policy in each fulfillment epoch, where stores fulfill online orders only when the inventory levels are above a certain threshold.

In Section 2.7, by means of a realistic numerical study on a network of stores and online FCs embedded in mainland USA, we show that our combined inventory and fulfillment heuristic improves greatly upon a benchmark solution which naively sets inventory levels in a decentralized fashion and fulfills online orders myopically. We test the relative performance of our heuristic over a variety of problem parameters such as shipping costs, online market share, network size, etc. Finally, we conclude with

Section 2.8 by discussing further generalizations including non-identical leadtimes and costs, and areas for future research.

2.3 Literature Review

Omnichannel retailing is a relatively new area in operations management literature, and has been gaining traction in recent years. Readers are referred to Rigby (2011) and Brynjolfsson et al. (2013) for comprehensive reviews of the topic. Existing papers in this area focus on the impact of online channel integration: Gao and Su (2017) study the impact of implementing store pickup on store operations, and Gallino et al. (2017) focus on sales dispersion from implementing store pickup. Other papers study the impact from the customers' point of view: Bell et al. (2017); Ansari et al. (2008), and Gallino and Moreno (2014) study customer migration due to product information, and Gao and Su (2016) analyze the effect of information provided to strategic omnichannel customers on store operations.

When there is no in-store demand, the problem is analogous to the pure play e-commerce setting, which has enjoyed recent attention in literature: Acimovic and Graves (2017) study the optimal allocation of replenishment to fulfillment centers to reduce shipping costs and mitigate costly spillovers, Lei et al. (2018) consider the joint pricing and fulfillment strategy to maximize the expected profits (revenue minus shipping costs), and Acimovic and Graves (2014) focus on fulfillment strategies to minimize outbound shipping costs.

There have been some studies which discuss integration of online demand to physical stores by means of a separate online fulfillment center, as this was the primary mode of fulfillment in the e-commerce channel in its nascent stages. Seifert et al. (2006) consider the inventory management of a system where an online warehouse handles online orders, and in case of stockouts, stores can fill these orders. Chen et al. (2011) consider a three location system consisting of two stores and an etailer, with a hierarchy to fulfillment - the etailer can fulfill online orders with the least cost, followed by store 1 and then store 2.

We consider a generalized setting representing the current retailing situation wherein physical stores are the primary ports of online fulfillment. To the best of our knowledge, the study closest to ours in emulating the problem setting, where online demand is integrated with the physical stores through store fulfillment is by Jalilipour Alishah et al. (2015). They consider a single store with online and in-store demands, and analyze decisions at three levels — fulfillment structure, inventory optimization and inventory rationing. They show that the optimal rationing policy between in-store and online demands is threshold-based, but their results do not extend to the multi-store case due to the complexity involved in an additional rationing decision - online orders from other regions. This setting is rather important in the context of e-commerce, and falls under the purview of the transshipment literature, where it has been shown to be an intractable problem to solve.

The key feature that online demands can be fulfilled from any store in the system is analogous to a reactive transshipment setting with zero transshipment lead time, as pointed out by Yang and Qin (2007), who called this 'virtual lateral transshipment'. In addition, our problem has multiple demand classes (online and in-store), where some classes of demand (in-store) cannot be subject to transshipment. For an extensive review of the transshipment literature, the readers are referred to Paterson et al. (2011).

The fact the the problem in question can be related to the transshipment literature offers little solace. Transshipment problems are infamously hard to solve, and analytical approaches can be done only for simplified cases with zero replenishment and transshipment leadtimes and two locations (Tagaras, 1989) or identical shipping costs across locations (Dong and Rudi, 2004). Tagaras and Cohen (1992) show that when there is positive replenishment leadtime, the problem becomes intractable even for two locations, as obtaining the optimal transshipment policy is mathematically complex due to its interdependence on demands during the leadtime, on-hand inventory and in-transit inventory.

Obtaining optimal order-up-to policies are by extension intractable as well, as they need to be calculated based on the optimal transshipment policy. Yao et al. (2016) have recently considered the optimal joint initial stocking and transshipment decisions for the two-store case, where stocking is done once at the beginning of a selling season, and transshipment is done multiple times during the season. Their analysis is limited to two stores, as key mathematical properties like submodularity do not extend to multiple locations.

Characterization of optimal policies in periodic review systems are especially difficult for lost-sales (see Bijvank and Vis 2011 for a review). However, Huh et al. (2009) find that optimal inventory levels assuming backordering provide a reasonable approximation to the lost-sales case when the penalty costs are very high compared to holding costs.

Due to the various complexities involved such as multiple stores, multiple epochs and periods, lost sales and joint optimization of inventory and fulfillment, one cannot hope to obtain a provably tight bound for a problem of this stature, let alone analytically finding the optimal solutions. We will instead develop simple, tractable and scalable heuristics, which perform well compared to naive strategies in most cases, with help from techniques used in literature.

Finally in the zero leadtime case, when online demand is fulfilled only once at the end of each review period, we show that the problem is analogous to a newsvendor network, with virtual lateral transshipment as a 'discretionary policy' (van Mieghem and Rudi, 2002). Newsvendor networks have been analyzed in great detail by van Mieghem and Rudi (2002) and van Mieghem (2003), building up from the multi-dimensional newsvendor models proposed by Harrison and van Mieghem (1999). However, as we shall show later, the canonical approach to optimizing inventory levels is difficult even for two stores due to the number of random demands involved, and is intractable for the multi-store case.

2.4 The General Problem - Model and Assumptions

Consider a system composed of a firm which owns N facilities $R_1, R_2, \ldots, R_N$ in different customer regions, selling a single product. Considering multiple products introduces complex combinatorial features to the fulfillment problem as a multi-item order can be fulfilled in different ways (Jasin and Sinha, 2015); we disregard this in our analysis to better study the interplay between inventory and fulfillment decisions. There are two classes of demand originating in each region i, modeled by non-negative and continuous random variables with well-behaved density functions.

1. the *in-store demand* (D_{is}) consists of customers picking items off the shelves (all the inventory is available on the shelf), with unmet demand lost immediately

2. the *online (ship-to-customer) demand* (D_{io}), consisting of customers ordering through the website or mobile app, with items delivered directly to their homes. For orders fulfilled from stores, the store staff pick up the item from the shelf, followed by packing and labeling in the store backroom, and shipping to the customer. A sale is lost when there is no available inventory for fulfillment at any location.

The demands are exogenous and are temporally independent, but can have any general channel or location correlation structure, while we require that the total demands in each region and across the system have continuous and well-defined density

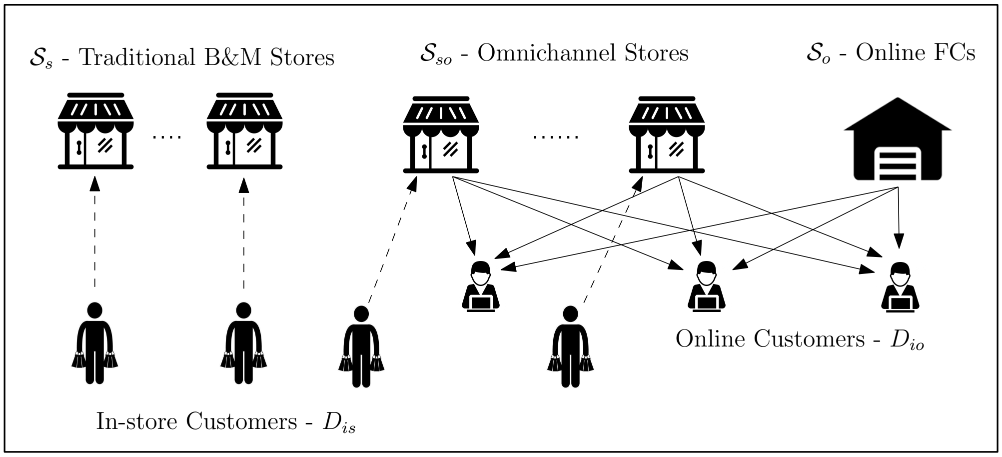

Figure 2.1: Three types of facilities in a fulfillment network a) Traditional brick-and-mortar stores ($\mathcal{S}_s$), b) Omnichannel stores ($\mathcal{S}_{so}$), and c) Online Fulfillment Centers ($\mathcal{S}_o$).

functions.

A typical retail fulfillment network is shown in Figure 2.1, where dashed lines represent customers visiting physical stores and solid lines represent items shipped to customers' homes. We consider three different types of facilities described by the following sets:

- $\mathcal{S}_s$ - physical stores which handle only in-store demand.
- $\mathcal{S}_o$ - online fulfillment centers (OFCs) which handle only online orders.
- $\mathcal{S}_{so}$ - omnichannel physical stores which handle both online and in-store demands.

Since traditional B&M stores plan for inventory independent of other facilities in the network, we exclude them from our analysis. We are hence interested in locations involved in online fulfillment, namely the omnichannel stores and online fulfillment centers, denoted by the set of facilities $\mathcal{S} = \mathcal{S}_o \cup \mathcal{S}_{so}$, and the number of such facilities is $N = |\mathcal{S}|$.

An important feature to be noted in the omnichannel problem is that unfulfilled in-store demand at one region cannot be fulfilled by stores in other regions. Any facility with available inventory can fulfill an online order, and hence there is pooling of online demands across regions in addition to pooling of in-store and online demands within each region.

2.4.1 Periodic Review Setup

We consider a periodic review model, where an order is placed by each facility at the start of each review period, and received with zero replenishment leadtime. The demands are realized during the course of the period based on the facility considered. We are interested in an optimum from the class of order-up-to policies, due to ease of implementation and practical relevance, and the order-up-to levels in each period are $y_1, \ldots, y_N$.

Based on conversations with industry executives, there are certain situations in the context of omnichannel stores where the leadtime is effectively negligible: in major cities like New York, store replenishment can only be done at night-time due to traffic restrictions. Such stores handle high volumes of sales, and are usually replenished daily from warehouses in nearby cities. An order placed in the afternoon can often be replenished before the following day. Positive leadtimes can significantly complicate analyses, and we discuss extending our heuristics to the case of non-identical leadtimes across locations in Section 2.8.

We assume that *online orders are fulfilled in multiple batches in each review period*, which we model by dividing a review period into T fulfillment epochs: in each epoch, in-store demand is fulfilled as it arrives, whereas online fulfillment decisions are made at the end of the epoch after observing demand, and orders are fulfilled with the available inventory.

The assumption reflects practical constraints in store operations: fulfillment activities in stores are usually done by store personnel, who in most cases also share additional store responsibilities. In such situations, it is better to fulfill online orders in batches, as opposed to having store staff picking items every time an online order is received.

2.4.2 Cost Parameters

We consider a per-unit service cost s_{ij} for online demand from region j fulfilled by R_i, which encapsulates the cost of picking the item off the shelf, packing and labelling, as well as the shipping cost for delivery. We have $s_{ij} > s_{ii}, \forall j \neq i$, as it is costlier to ship an item over longer distances. We will refer to the service costs s_{ii} (within the same region) as *shipping* costs, and s_{ij} (across regions) as *cross-shipping* costs.

In practice, the handling (pick-pack-and-label) component of the service cost is higher for stores fulfilling online demand, as it involves human labor, than for OFCs where the process can be automated and streamlined. The shipping component of the service cost can be higher for the OFCs which are usually located farther away from population centers.

We have identical costs at each location, including shipping costs $s_{ii} = s$, $\forall i$. At the end of a fulfillment epoch, each unit of unused inventory incurs an overage cost h, and each unit of unfulfilled in-store and online demands incur penalty costs p_s and p_o respectively. We assume that $p_s > p_o - s > 0$, as in-store demand is fulfilled first and costlier to lose, and cross-shipping always leads to a myopic reduction in cost: $s_{ij}\ (= s_{ji}) < h + p_o$, $\forall i, j$. We ignore the purchasing cost of inventory, but this can be incorporated through linear terms.

2.4.3 Stochastic Programming Formulation

We are now ready to write the total expected per period cost function for the case where online demand is fulfilled over T fulfillment epochs in each review period. We focus on the single period to obtain order-up-to levels, which we show in Section 2.5.2 to be optimal in a multi-period setting in the case of negligible replenishment leadtimes.

In each fulfillment epoch t, let the starting inventory levels be denoted by $\mathbf{x^t} = (x_i^t)_i$, and $\tilde{\mathbf{D}}^{\mathbf{t}} = (D_{is}^t, D_{io}^t)_i$ denotes the demands. From location R_i, let z_i^t be the amount of inventory used to fulfill the in-store demand, and Z_{ij}^t be the amount of inventory shipped to fulfill online demand from region j, denoted in vector form as $\mathbf{z^t}, \mathbf{Z^t}$ respectively.

We have a T-stage stochastic program, with the cost-to-go function in epoch t, $C_t(\mathbf{x^t}, \tilde{\mathbf{D}}^{\mathbf{t}})$ is given by:

$$C_t(\mathbf{x^t}, \tilde{\mathbf{D}}^{\mathbf{t}}) = \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \left[P(\mathbf{x^t}, \tilde{\mathbf{D}}^{\mathbf{t}}, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E} C_{t+1} \left(\left(x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t \right)_i, \tilde{\mathbf{D}}^{\mathbf{t+1}} \right) \right] \tag{2.1}$$

where $P(\mathbf{x^t}, \tilde{\mathbf{D}}^{\mathbf{t}}, \mathbf{z^t}, \mathbf{Z^t})$ is the total cost in fulfillment epoch t, given by:

$$\begin{aligned} P(\mathbf{x^t}, \tilde{\mathbf{D}}^{\mathbf{t}}, \mathbf{z^t}, \mathbf{Z^t}) = {} & \sum_{i=1}^{N} h \left(x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t \right) + \sum_{i=1}^{N} p_s (D_{is}^t - z_i^t) \\ & + \sum_{j=1}^{N} p_o \left(D_{jo}^t - \sum_{i=1}^{N} Z_{ij}^t \right) + \sum_{i=1}^{N} s Z_{ii}^t + \sum_{i=1}^{N} \sum_{j=1, j \neq i}^{N} s_{ij} Z_{ij}^t \end{aligned} \tag{2.2}$$

and Δ is the set of feasible fulfillment decisions, described by the following set of constraints:

$$\begin{aligned}
z_i^t + \sum_{j=1}^{n} Z_{ij}^t \le x_i^t, & \qquad \forall i \in [N], \forall t \in [T] \\
z_i^t \le D_{is}^t, & \qquad \forall i \in [N], \forall t \in [T] \\
\sum_{i=1}^{n} Z_{ij}^t \le D_{jo}^t, & \qquad \forall j \in [N], \forall t \in [T] \\
\mathbf{z^t}, \mathbf{Z^t} \ge 0, & \qquad \forall t \in [T]
\end{aligned} \tag{2.3}$$

The first inequality in Δ represents the supply constraint, and the second and third inequalities model the fulfillment constraints. Note that the online demand in one region can be fulfilled from any facility in the network, as seen in the third inequality in (2.3).

The goal is to obtain the initial stocking level $\mathbf{y} = (y_i)_i$. The single period, T-epoch problem can thus be stated as follows: $\min_{\mathbf{y} \ge \mathbf{0}} \mathbb{E}[C_1(\mathbf{y}, \tilde{\mathbf{D}})]$. This is a convex minimization problem, as we will later show in Section 2.6, but it is intractable to solve. The fulfillment decisions are similar to optimal transshipment decisions with non-negligible lead time, as decisions in any fulfillment epoch depend on future demands in that review period. As pointed out by Tagaras and Cohen (1992) for the two-store case in traditional transshipment, while the optimal fulfillment policy may be threshold-based, the optimization becomes intractable due to the complexity of the decision space in the dynamic programming formulation.

We cannot hope to solve this problem to optimality, and we resort to heuristic solutions that perform well compared to simple, naive strategies and hindsight optimal lower bounds. Note that a heuristic solution specifies both the initial stocking level and fulfillment policy.

We first develop the inventory heuristic in the following way: treat the T-epoch problem as a single fulfillment epoch. A similar method was also used by Tagaras and Cohen (1992) to set heuristic inventory levels for the two-location transshipment problem with leadtime, based on numerical evidence that most transshipments took place at or near the end of the planning horizon, when stockouts are more likely to happen.

Our problem is different in two aspects: 1) we have in-store demands which are more costly to lose than online demands and do not have pooling flexibility, and 2) demands follow lost sales. However, we adopt this single fulfillment epoch approximation as it provides a tractable alternative by decoupling inventory and fulfillment decisions, because:

1. a myopic fulfillment policy is optimal, where online demands are fulfilled to the maximum possible extent with the available inventory, and as a result,

2. the inventory problem reduces to a single stage stochastic linear program.

With the help of results obtained through this approximation, we formulate inventory and fulfillment heuristic solutions for the multi-period, multi-location problem in Section 2.6, and numerically test their performance in Section 2.7.

2.5 The Single Fulfillment Epoch Case (T=1) - Model and Analysis

In this setting, items are ordered and received at the beginning of the period with zero lead time, and in-store demand is fulfilled as it arrives. Due to the single fulfillment epoch assumption, the fulfillment of online demand is done once at the end of the review period, after in-store demands are fulfilled. There is no benefit to reserving inventory for future demands as replenishments arrive immediately. In such a case, a myopic fulfillment policy is optimal, where online orders are fulfilled to the maximum possible extent in each period.

The case of single fulfillment epoch is quite common in present day omnichannel retailing where stores are replenished daily. Most stores still rely on third party carriers such as UPS and FedEx to ship items to customers. Online orders to be shipped are loaded onto these trucks once a day from the store backroom, usually towards the end of the day. This is especially popular in the context of same-day and next-day deliveries, where stores allow online ordering until a cutoff time, and these orders are ready to be shipped by the end of the day. However with developments in drone technology in the future, one can easily envision stores that fulfill multiple times in a day, which we address through the general case of multiple fulfillment epochs ($T > 1$) in Section 2.6.

We first consider the two-store setting to exhibit the complicated nature of the decoupled inventory problem alone, given the optimal fulfillment policy is myopic. The insights derived in this case inform our analysis of a generalized multi-location case, which includes a network of omnichannel stores and online FCs.

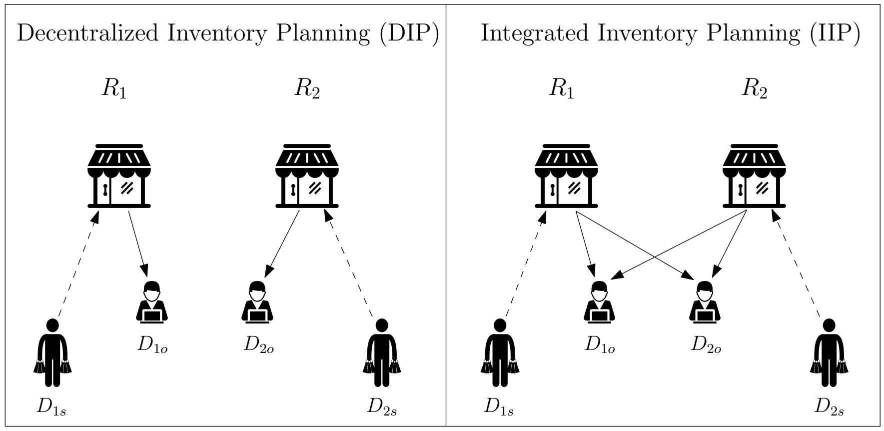

Figure 2.2: Two methods of inventory planning - 1) Decentralized inventory planning (DIP) and 2) Integrated inventory planning (IIP)

2.5.1 The Two-store System

A firm owns two retail stores R_1 and R_2 serving different regions, with two demand streams originating form each region – in-store demand (D_{1s}, D_{2s}), and online demand (D_{1o}, D_{2o}). The objective is to set the initial inventory levels y_1 and y_2 to minimize the total expected cost. We consider two solutions – decentralized inventory planning (DIP) and integrated inventory planning (IIP), which are represented in Figure 2.2. The assumptions on cost parameters are recapitulated in the set Ψ in Equation 2.4.

$$\Psi = \left\{ \; p_s > p_o - s_i > 0, \, \forall i; \quad h + p_o > s_{ij} > s, \, \forall i, j \neq i \; \right\} \tag{2.4}$$

2.5.1.1 The Decentralized Inventory Planning (DIP) Strategy (Pooling within Regions)

We first consider the case where the firm plans for inventory at its stores in a decentralized fashion, without planning in advance for cross-shipping. This serves as a benchmark for any inventory heuristic we may develop for the centralized planning case. The inventory level at store i is set with an objective to minimize the total expected cost incurred in meeting the demands from that region, given by:

$$\begin{aligned} C^{DIP}(y_i) = &\mathbb{E}\Big[h\left((y_i - D_{is})^+ - D_{io}\right)^+ + p_s(D_{is} - y_i)^+ \\ &+ p_o\left(D_{io} - (y_i - D_{is})^+\right)^+ + s\min\left((y_i - D_{is})^+, D_{io}\right)\Big] \end{aligned} \tag{2.5}$$

where $x^+ = \max(x, 0)$. The cost function is convex, which can be seen by expressing Equation 2.5 in terms of the total demands $D_i = D_{is} + D_{io}$ as follows:

$$C^{DIP}(y_i) = s\mu_{io} + \mathbb{E}\Big[h\left(y_i - D_i\right)^+ + \left(p_o - s\right)\left(D_i - y_i\right)^+ + \left(p_s - \left(p_o - s\right)\right)\left(D_{is} - y_i\right)^+\Big] \tag{2.6}$$

where $\mu_{io} = \mathbb{E}[D_{io}]$. The simplification is done using the identities $\min(x, y) = y - (y - x)^+$, and $(D_{is} - y_i)^+ + \left(D_{io} - (y_i - D_{is})^+\right)^+ = (D_i - y_i)^+$, the latter holds when demands are non-negative. The optimal inventory levels (y_1^{DIP}, y_2^{DIP}) can obtained from implicit equations:

$$\left(h + p_o - s\right) F_i\left(y_i^{DIP}\right) + \left(p_s - p_o + s\right) F_{is}\left(y_i^{DIP}\right) = p_s, \qquad \forall i = 1, 2 \tag{2.7}$$

where F_i is the cumulative distribution function of demand D_i. A line search yields unique optimum, as the left hand side is increasing in y_i^{DIP}, and the right hand side is constant.

2.5.1.2 The Integrated Inventory Planning (IIP) Strategy (Pooling within and across Regions).

This is similar to the DIP scenario, except that after R_i has fulfilled its own in-store and online demands, unfulfilled online orders from region j ($\neq i$) can be fulfilled using any available inventory at R_i. In the two-store problem, the cross-shipped quantity from store R_i to region j can be explicitly calculated as the minimum of the inventory available at R_i and the unfulfilled online demand at R_j, after each store has attempted to fulfill its own demands. The total expected one-period cost function is:

$$\begin{aligned} C^{IIP}(y_1, y_2) = \mathbb{E}\Bigg[\sum_i \Bigg(& h\left((y_i - D_{is})^+ - D_{io}\right)^+ + p_s(D_{is} - y_i)^+ \\ & + p_o\left(D_{io} - (y_i - D_{is})^+\right)^+ + s\min\left((y_i - D_{is})^+, D_{io}\right)\Bigg) \\ & + (s_{12} - h - p_o)\min\left(\left((y_1 - D_{1s})^+ - D_{1o}\right)^+, \left(D_{2o} - (y_2 - D_{2s})^+\right)^+\right) \\ & + (s_{21} - h - p_o)\min\left(\left((y_2 - D_{2s})^+ - D_{2o}\right)^+, \left(D_{1o} - (y_1 - D_{1s})^+\right)^+\right)\Bigg] \end{aligned} \tag{2.8}$$

The additional terms in Equation 2.8 that are absent in Equation 2.5 represent the value of cross-shipping: the total savings by cross-shipping a unit from R_i to region j, $h + p_o - s_{ij}$, times the total quantity cross-shipped from R_i to region j. The total cross-shipped quantity can be expressed as

$$\sum_i \left(D_{io} - (y_i - D_{is})^+\right)^+ - \left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+$$

The first term represents the total unfulfilled online demand if there was no cross-shipping allowed, and the second term represents the unfulfilled online demand with cross-shipping. Naturally, the difference yields the cross-shipped quantity. By using this expression, as well as the simplification techniques used in Equation 2.6, we can simplify Equation 2.8 as follows:

$$\begin{aligned} C^{IIP}(y_1, y_2) = s \sum_i \mu_{io} + \sum_i \mathbb{E}\Big[& h\,(y_i - D_i)^+ + (p_s - p_o + s)(D_{is} - y_i)^+ + (p_o - s)\,(D_i - y_i)^+ \Big] \\ & + (s_{12} - h - p_o) \left[\sum_i (D_i - y_i)^+ - \sum_i (D_{is} - y_i)^+ - \left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+ \right] \end{aligned} \tag{2.9}$$

We can rearrange the terms to a convex expression, except $\left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+$, which is non-convex in y_i's. This is seen by keeping y_1 constant and changing y_2.

$$\begin{aligned} & \left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+ \\ & = \begin{cases} \left(D_{1o} + D_{2o} - (y_1 - D_{1s})^+\right)^+, & \text{if} \quad y_2 \leq D_{2s} \\ D_2 + D_{1o} - (y_1 - D_{1s})^+ - y_2, & \text{if} \quad D_{2s} < y_2 < D_2 + D_{1o} - (y_1 - D_{1s})^+ \\ 0, & \text{if} \quad y_2 \geq D_2 + D_{1o} - (y_1 - D_{1s})^+ \end{cases} \end{aligned} \tag{2.10}$$

In the event that $D_{is} = 0, \forall i$ (similar to traditional transshipment considered by Dong and Rudi, 2004), the formulation in Equation 2.9 would directly yield a convex cost function. Convexity is not obvious in our case, as the nested piecewise linear function in Equation 2.10 is neither convex nor concave, and this is purely due to the fact that in-store demand is fulfilled first and cannot be subject to cross-shipment. However, the total cost can be shown to be jointly convex in the inventory levels (Proposition 2.5.1):

Proposition 2.5.1 *Under the conditions on cost parameters in* Ψ*,*

(a) $C^{IIP}(y_1, y_2)$ *is jointly convex in the order-up-to levels.*

(b) There exist regions $\Omega_k(y_1, y_2)$ *in the demand space, such that in each region the dual-price vector* λ^k *corresponding to the variables* y_1*,* y_2 *remains constant, and the gradient of the IIP cost function can be written as*

$$\nabla C^{IIP}(y_1, y_2) = (h, h)^\intercal - \sum_k \lambda^k \mathbb{P}\left(\Omega_k\left(y_1, y_2\right)\right) \tag{2.11}$$

All proofs are relegated to the Appendix. We first observe that under the assumptions in Ψ, C^{IIP} can be expressed as the expectation of a linear program, through which joint convexity in inventory levels is established. By noting structural similarities with a newsvendor network (van Mieghem and Rudi, 2002), we derive an expression for the gradient based on the dual prices $\lambda = (\lambda_1, \lambda_2)^\intercal$, which are simply the shadow prices of the constraints involving y_1 and y_2 in the linear program representation (Equation A.1, Appendix A.1).

The demands are shown to be separable into independent regions Ω_k based on the values of y_1 and y_2, within which the dual prices $\lambda^k = (\lambda_1^k, \lambda_2^k)$ are constant (refer to Appendix A.2 for a detailed discussion), which enables formulating the gradient as shown in Equation 2.11. The optimal solution (y_1^{IIP}, y_2^{IIP}) can thus be obtained by a gradient descent algorithm. Given values of $(y_1,\ y_2)$ in each iterative step, the probability of realization of every demand region has to be calculated. As we extend to N stores, we face the following hurdles:

- exponentially increasing number of demand regions Ω_k (in which the dual prices remain constant), whose identification is non-trivial, and
- repeated probability calculations of a $2N$-dimensional multivariate distribution for these demand regions.

The non-triviality in identification of these demand regions arises from the fact that cross-shipment quantities are now set by a transportation linear program, as compared to explicit expressions in the two-store case. Hence we develop a tractable lower bound which yields a heuristic solution for the two-store case, which we later extend to multiple locations.

2.5.1.3 Lower Bound and Heuristic for the Two-Location Problem

An important feature which complicates the IIP cost function is that the in-store demands are not pooled across regions, which in turn leads to complex and non-convex coupled terms in the cost function. We relax this by treating unfulfilled in-store demand as online demand which can be fulfilled by cross-shipping. Specifically, we replace the total unfulfilled demand $\sum_i (D_{is} - y_i)^+ + \left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+$ by its lower bound, which is the total unfulfilled demand when all demands are pooled $\left(\sum_i D_i - \sum_i y_i\right)^+$. Substituting this in Equation 2.9 and simplifying, we get the following cost function:

$$\begin{aligned}C^{LB}(y_1, y_2) = s(\mu_{1o} + \mu_{2o}) + \mathbb{E}\Big[&h\,(y_1 + y_2 - D)^+ + (p_o - s_{12})\,(D - y_1 - y_2)^+ \\ &+ (p_o - s - (p_o - s_{12}))\,(D_1 - y_1)^+ + (p_o - s - (p_o - s_{12}))\,(D_2 - y_2)^+ \\ &+ (p_s - (p_o - s))\,(D_{1s} - y_1)^+ + (p_s - (p_o - s))\,(D_{2s} - y_2)^+\Big]\end{aligned}$$

where $D = D_1 + D_2$, the total demand. Proposition 2.5.2 establishes C^{LB} as a lower bound:

Proposition 2.5.2 $C^{LB}(y_1, y_2) \leq C^{IIP}(y_1, y_2)$, $\forall y_1, y_2 \geq 0$

By removing the nested piecewise linear terms in C^{IIP} from Equation 2.9, we no longer need the gradient descent approach, as the first order conditions for C^{LB} are greatly simplified:

$$(h + p_o - s_{12})F_D\left(\sum_{j=1,2} y_j\right) + (s_{12} - s)F_{D_i}(y_i) + (p_s - p_o + s)F_{D_{is}}(y_i) = p_s, \quad \forall i \tag{2.12}$$

We have a system of two equations with two variables, which can be solved using numerical methods to yield a heuristic solution $\mathbf{y^{LBH}}$ with expected cost $C^{LBH} = C^{IIP}(\mathbf{y^{LBH}})$. Equation 2.12 is of a similar structure to the first order conditions obtained by Dong and Rudi (2004) for the case of constant transshipment cost, with a key difference: there is an additional term stemming from the presence of in-store demands with a higher underage cost than the online demands. This allows us to fix inventory levels at each location separately, in contrast to Dong and Rudi (2004) where the optimality equation only yields a system-wide inventory level.

We make the following relaxation to formulate the lower bound: $\sum_i (D_{is} - y_i)^+ +$

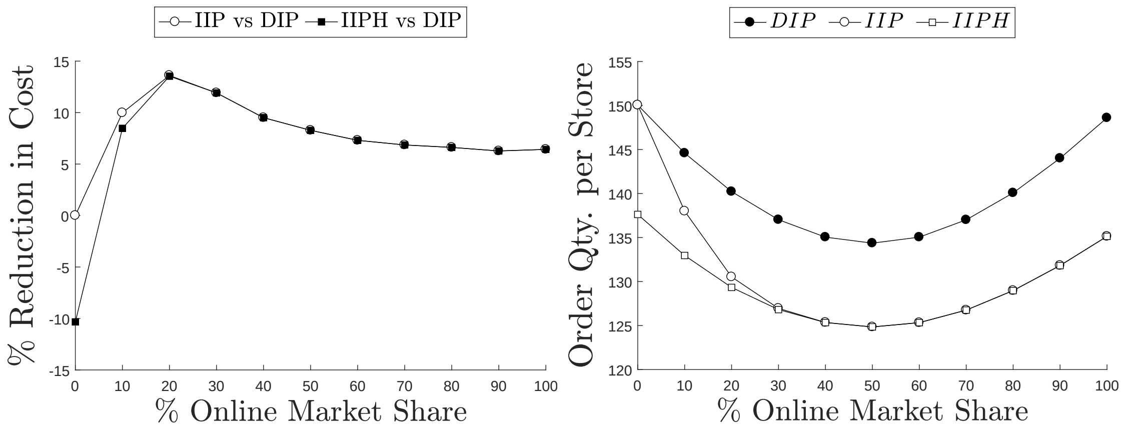

(a) (b)

Figure 2.3: Shows the effect of online market share on C^{IIP}, C^{LBH}, and C^{DIP} (left) and the corresponding optimal order quantities per store (right).

$\left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+$ with $\left(\sum_i D_i - \sum_i y_i\right)^+$; this relaxation will be tight when the in-store demand is very small compared to the online demand, as the optimal inventory levels are set based on the total demands. We test this numerically by changing the mix of in-store and online demands in Figure 2.3. The mean in-store and online demands are calculated as a proportion of a fixed total mean demand ($= 100$) in each region. The demands are normal and identical across regions, with the coefficient of variation fixed at 0.3 for each demand. The cost parameters are: $h = 5$, $p_s = 100$, $p_o = 100$, $s = 10$, $s_{12} = 15$.

From Figure 2.3a, we see that the heuristic provides savings over the DIP strategy for most cases, except for small values of online market share ($< 10\%$). However, we note that for such small values of online market share, the potential savings from centralized planning is also small, as seen from comparing the IIP and DIP costs. In such cases, one can simply resort to planning for each region separately using the DIP strategy.

Centralized inventory planning is most valuable when there is a moderate mix of online and in-store demands. As online demand grows in comparison to in-store demands, the effect of pooling across regions increases, due to two reasons: 1) more demand is pooled across regions which leads to a bigger reduction in variability of the total online demand, and 2) pooled online demands can better absorb the variability in the in-store demands. Thus, the maximum savings is achieved when there is a good mix of online and in-store demands so that the pooling across channels and locations work in synergy.

As the in-store demand becomes smaller, the probability that there will be unfulfilled in-store demand decreases, and the heuristic solution converges to the optimal IIP solution (Figure 2.3b). Thus for high values of online market share, in-store demand can effectively be treated as online demand which explains the stable savings achieved by the IIP solution.

The cost savings directly arise from a change in inventory levels in anticipation of pooling across locations. Proposition 2.5.3 addresses this observation from Figure 2.3b that the IIPH solution consistently stocks less than the DIP solution at each store.

Proposition 2.5.3 *For identical stores and normally distributed demands, $y^{LB} \leq (\geq)y^{DIP}$ whenever $y^{DIP} \geq (\leq)\mu$, where μ is the mean total demand at a store. Under perfect positive correlation across locations, $y^{LB} = y^{DIP} = y^{IIP}$.*

Similar to the intuition in newsvendor settings, $y^{DIP} \geq \mu$ would hold when underage costs are greater than overage costs, but this does not translate into an analytical proof due to the structure of the optimality equations in Equation 2.7, which has a mixture distribution as compared to a simple normal distribution in newsvendor theory. Lastly, positive correlation across locations reduces the pooling benefits achieved by cross-shipping, and under perfect correlation, there is no benefit from pooling as all locations either have too much or too little inventory without any imbalance.

2.5.2 The Multi-Location Problem

We extend the two-store problem discussed so far to a generalized setting with multiple regions, as described earlier in Section 2.4 (Figure 2.1). The cross-shipping costs are taken to be $s_{ij} = s + f(d_{ij})$, where d_{ij} is the distance between location R_i and region j, and f is a non-negative, increasing function such that $f(d) \to 0$ as $d \to 0$. Also, $\sup_{d \in \mathcal{D}} f(d) \leq h + p_o - s$, where $\mathcal{D} = \{d_{ij}, \forall i, j\}$, so that the conditions in Equation 2.4 hold true.

The decentralized solution $\mathbf{y^{DIP}}$ derived from Equation 2.7 readily extends to the multiple locations as the problem is decoupled by region, whereas the optimal IIP solution cannot be obtained due to the computational infeasibility even of the two-store approach. However, we can extend the heuristic and lower bound developed in the two-store case, by lowering all cross-shipping costs to $s_{\min} = \min_{i \neq j} s_{ij}$, yielding the first order conditions:

$$(h + p_o - s_{\min})F_D\big(\sum_{j \in \mathcal{S}} y_j\big) + (s_{\min} - s)F_{D_i}(y_i) + (p_s - p_o + s)F_{D_{is}}(y_i) = p_s, \quad \forall i \in \mathcal{S} \tag{2.13}$$

The corresponding cost function yields a lower bound to the multi-location problem, satisfying Propositions 2.5.2 and 2.5.3 (the proofs are similar to the two-store case, and hence omitted). The optimal solution can be found easily for small number of stores by iterative root-finding algorithms such as the Newton-Raphson method. The computational burden of this solution, although reduced from the newsvendor network approach by van Mieghem and Rudi (2002), is still significant for omnichannel networks in practice with thousands of stores due to the number of variables involved. A small change to the parameters: reducing $s_{\min}$ to s yields a weaker lower bound:

$$\begin{aligned} C^{LBN}(y_1, \ldots, y_N) = s \sum_{i \in \mathcal{S}} \mu_{io} + \Bigg[&\mathbb{E}h \left(\sum_{i \in \mathcal{S}} y_i - D_{\mathcal{S}} \right)^+ + \mathbb{E}\,(p_o - s) \left(D_{\mathcal{S}} - \sum_{i \in \mathcal{S}} y_i \right)^+ \\ &+ \mathbb{E}\,(p_s - p_o + s) \sum_{i \in \mathcal{S}_{so}} (D_{is} - y_i)^+ \Bigg] \end{aligned} \tag{2.14}$$

C^{LBN} is convex in the inventory levels, and can be solved to yield a heuristic solution y^{LBN} characterized by the first order conditions:

$$(h + p_o - s)F_{D_{\mathcal{S}}} \left(\sum_{j \in \mathcal{S}} y_j^{LBN} \right) + (p_s - p_o + s)F_{D_{is}}(y_i^{LBN}) = p_s, \quad \forall i \in \mathcal{S} \tag{2.15}$$

Parallels can be drawn to Equation 2.12 and Dong and Rudi (2004), as the presence of in-store demands enables the characterization of inventory levels at each individual location. As a consequence, the calculation of y^{LBN} is computationally light, established by the following Proposition.

Proposition 2.5.4 *The heuristic solution is unique, and when demands follow a multivariate normal distribution, the heuristic inventory levels at stores are at the same critical fractile of their corresponding in-store demands.*

In contrast with Equation 2.13, we only need to solve for one variable, namely the common critical fractile of the in-store demands. This reduces the computational effort drastically, even for very large networks. However, the optimal solution has zero inventory in the OFCs – this is because when all cross-shipping costs are lowered to s, a unit of inventory at the OFC can lead to a decrease in total cost if it was instead at a store, as it can also serve to fulfill in-store demands.

We modify y^{LBN} to obtain the heuristic solution y^{IIPH} for multiple locations by calculating order quantities for the OFCs separately, and using them in Equation 2.15 to compute order quantities for the omnichannel stores. The order-up-to quantities for OFCs are calculated from the pooled total order quantity for OFCs, which is determined using the newsvendor quantity for the combined online demand $D_{\mathcal{S}_o} = \sum\limits_{i \in \mathcal{S}_o} D_{io}$.

$$\sum_{j \in \mathcal{S}_o} y_j^{IIPH} = F_{D_{\mathcal{S}_o}}^{-1} \left(\frac{p_o - s}{h + p_o - s} \right) \tag{2.16}$$

The actual underage cost for online demands at the OFCs would be less than $p_o - s$ and would depend on inventory information of stores, as stores can fulfill these online orders with available inventory. The calculation of inventory levels at stores and OFCs are dependent on each other, but since we are forced to estimate the inventory at OFCs separately, we inflate the underage cost to $p_o - s$ which yields a higher overall inventory level at the OFCs. This is a limitation that arises out of our heuristic approximation, but it allows us to extend the heuristic to the case where OFCs have a different service cost (s_o) compared to the stores (s), as the inventory calculation for the OFCs is done separately.

To calculate the individual order quantities y_i^{IIPH}, $i \in \mathcal{S}_o$, we use the method of obtaining order-up-to quantities for multiple products with capacity constraints, as described in Chopra and Meindl (2007, p. 367). The total capacity is the total order-up-to quantity calculated from Equation 2.16, and the order-up-to quantity for each product corresponds to the order-up-to quantity for each OFC. Each unit from $\sum\limits_{j \in \mathcal{S}_o} y_j^{IIPH}$ is allocated incrementally to the OFCs based on the individual expected marginal costs. Once the order-up-to quantities for the OFCs are obtained, they are used in Equation 2.17 to determine order-up-to levels for other omnichannel stores.

$$(h + p_o - s) F_{D_{\mathcal{S}}} \left(\sum_{j \in \mathcal{S}} y_j^{IIPH} \right) + (p_s - p_o + s) F_{D_{is}} \left(y_i^{IIPH} \right) = p_s, \quad \forall i \in \mathcal{S}_{so} \tag{2.17}$$

Note that individual store inventory levels are directly obtained from Equation 2.17 due to the presence of in-store demands. Calculating the heuristic solution y^{IIPH} is also computationally fast, as Proposition 2.5.4 still applies to Equation 2.17. The cost of the heuristic solution is given by $C^{IIPH} = C^{IIP}(y^{IIPH})$. We capture the effect of virtual pooling among the facilities in this heuristic, and the systematic approach is shown in Algorithm 1.

Algorithm 1 Procedure to calculate the heuristic solution $\mathbf{y}^{IIPH}$

1: For physical stores in set $\mathcal{S}_s$, set $y_i^{IIPH} = F_{is}^{-1}\left(\frac{p_s}{h+p_s}\right), \forall i \in \mathcal{S}_s$.

2: **for** $i \in \mathcal{S}_o$ (OFCs) **do**

3: Calculate total order quantity: $y^{TOT} = F_{D_{\mathcal{S}_o}}^{-1}\left(\frac{p_o-s}{h+p_o-s}\right)$, where $D_{\mathcal{S}_o} = \sum_{i\in\mathcal{S}_o} D_{io}$.

4: Set $y_i^{IIPH} = 0, \forall i \in \mathcal{S}_o$, and $rem = \lfloor y^{TOT} \rfloor$.

5: Calculate marginal cost $MC_i\left(y_i^{IIPH}\right) = -\left(p_o - s\right)(1 - F_{D_{io}}(y_i^{IIPH})) + hF_{D_{io}}\left(y_i^{IIPH}\right)$

6: Choose $i^* = \min_{i\in\mathcal{S}_o} MC_i(y_i^{IIPH})$. Set $y_{i^*}^{IIPH} \leftarrow y_{i^*}^{IIPH} + 1$

7: Set $rem \leftarrow rem - 1$. If $rem > 0$, go to Step 3.

8: **for** $i \in \mathcal{S}_{so}$ **do**

9: Calculate order quantities implicitly from the optimality equations:
$$\left(h + p_o - s\right) F_{D_\mathcal{S}}\left(\sum_{j\in\mathcal{S}} y_j^{IIPH}\right) + \left(p_s - p_o + s\right) F_{D_{is}}\left(y_i^{IIPH}\right) = p_s,\ \forall i \in \mathcal{S}_{so}.$$

The performance of the heuristic clearly depends on the structure of the network which directly influences the cross-shipping costs, in addition to the mix of in-store and online demands. However in practice, the range of shipping costs is not too large: for a 5lb package, the ratio $\max_{i,j} s_{ij}/s$ is less than 2 for the UPS Ground option, and less than 3 for the UPS Next Day Air option (UPS, 2017) for locations within the mainland US. We test the sensitivity for factors that adversely affect heuristic performance in Section 2.7 (Figure 2.5).

As the problem scale increases, and the number of stores grows large within a given area to accommodate the increase in demand, it is highly likely that a store with unfulfilled online demand can find a close-by store with available inventory, and hence, almost all cross-shipping takes place over short distances, at a cost close to s. Thus, we can expect the heuristic solution to be close to the optimal solution, and as a consequence of this notion, Proposition 2.5.5 shows that the heuristic is near optimal in an asymptotic sense.

Proposition 2.5.5 *As the number of omnichannel stores in a given area increases, with demands bounded and i.i.d. across regions, for sufficiently small $h > 0$, the heuristic is near optimal in an asymptotic sense with a constant approximation factor, i.e.*

$$\frac{C^{IIPH}}{C^{LBN}(y^{IIPH})} \leq \frac{h + p_s}{p_s - p_o + s}, \; as \; N \to \infty$$

The proposition holds when all locations have omnichannel stores, and $y^{LBN} = y^{IIPH}$. We first show that reducing cross-shipping costs to s preserves optimality in the asymptotic setting, by considering a simplified setting where the stores are uniformly distributed in the given region, which is in-turn divided into identical sub-regions. As the number of stores grows large, each sub-region has sufficient supply to fulfill its demands, and hence cross-shipping takes place only within the sub-regions with costs converging to s.

The assumption that in-store demands are pooled can affect the heuristic when online demands are small. However, we bound the heuristic performance by a constant approximation factor dependent only on cost parameters. While this bound is not tight, it shows that the heuristic is not critically affected by its assumptions as the problem scale grows.

2.6 Multi-Period, Multi-Location, and Multiple Fulfillment Epochs

So far, we have discussed the single review period setting where online fulfillment is done once, at the end of the period. We now switch back to the general version of the problem described in Section 2.4, with multiple review periods and online demand fulfilled over T fulfillment epochs in each review period. This is a more realistic representation of practice, as we closely approximate the continuous time case, because the value of T can be flexibly large. We start by proving convexity for the single period problem described in Equation 2.1.

Proposition 2.6.1 *The single-period, T-fulfillment-epoch expected cost function given by $C(\mathbf{y}) = \mathbb{E}C_1(\mathbf{y}, \tilde{D})$ is jointly convex in the inventory levels y_i.*

The proof follows by induction. Let the optimal solution to the single period problem be denoted by $\mathbf{y^{IIP}}$. We extend our analysis to the finite horizon case with multiple periods.

Proposition 2.6.2 *For the finite horizon problem with lost sales and zero replenishment leadtime, a stationary base-stock policy is optimal, with order-up-to levels $\mathbf{y^{IIP}}$.*

For the zero replenishment leadtime case with lost sales, the multi-period problem reduces to solving a single-period problem, and the proof is similar to traditional multi-period inventory problems involving lost-sales. As noted earlier, solving for $\mathbf{y}^{\mathbf{IIP}}$ is difficult, as optimal fulfillment decisions are intractable. Hence, we resort to heuristic solutions developed from our analysis of the single period problem.

2.6.1 Inventory Levels

To obtain a heuristic solution to set order-up-to levels, we use the procedure described in Algorithm 1, by approximating the problem as a single fulfillment epoch problem. Naturally, the demands used to calculate the heuristic solutions are the total review-period demands at each location. For example, the review-period in-store demand at store i is given by $\mathcal{D}_{is} = \sum_{t=1}^{T} D_{is}^{t}$. Also, the holding cost parameter used in the algorithm is the review-period holding cost, which is given by $\bar{h} = h^{*}T$.

We compare this heuristic solution with the naive strategy which plans for inventory in a decentralized fashion. We extend the DIP solution derived in Equation 2.7, by using the total review-period demands for each location and holding cost $\bar{h}$. We will continue to denote the heuristic solution derived in this fashion by IIPH and the decentralized solution as DIP for the numerical studies in the following sections.

2.6.2 Fulfillment Policies

We consider two fulfillment policies, which dictate how online orders are fulfilled:

1. the *myopic fulfillment (MF)* policy, where online demands in the current fulfillment epoch are fulfilled to the maximum possible extent with the available inventory, without consideration for demands in the future, and

2. the *threshold fulfillment (TF)* policy, which reserves inventory at each location for future in-store demands, by halting online fulfillment from a location when the inventory level falls below a certain threshold in each fulfillment epoch.

As future in-store demands are costlier to lose and do not have the additional flexibility of cross-shipping, it is intuitive that the TF policy can lead to reduction in costs compared to the MF policy when implemented well. Rationing inventory between high-priority and low-priority demands has been studied in literature (for a review, refer to Kleijn and Dekker, 1999), and along similar lines, Jalilipour Alishah et al. (2015)

Algorithm 2 Implementation of the Threshold Fulfillment (TF) Policy

1: At the start of the review period, evaluate thresholds w_i^t, $\forall i, t$ using Equation 2.18.
2: In each fulfillment epoch t, each location first fulfills its own in-store demand to the maximum possible extent, and the leftover inventory at location i is $\hat{x_i}^t$.
3: Calculate fulfillment capacities for each location i as $K_i^t = (\hat{x_i}^t - w_i^t)^+$.
4: Online fulfillment decisions Z_{ij}^t are obtained from the transportation linear program: $\Big\{ \min \ \sum_{i,j} (s_{ij} - h - p_o) Z_{ij}^t, \quad \text{subject to: } \sum_k Z_{kj}^t \leq D_{jo}, \ \sum_k Z_{ik}^t \leq K_i^t, \ Z_{ij}^t \geq 0, \ \forall i, j \Big\}$

prove the existence of an optimal threshold rationing policy between in-store and online demands at a single store.

In our case it is not straightforward to estimate the underage cost for the low-priority (online) demand, as it is endogenized by the fulfillment policy followed and depends on where an order is fulfilled from. The optimal thresholds depend on in-store and online demands in a complicated, network-based fashion, as online demands are pooled across locations, and their calculation is akin to obtaining optimal transshipment decisions based on such a threshold structure. We propose simple newsvendor-based thresholds which only take into account future in-store demands. In any fulfillment epoch t, an amount w_i^t is reserved at store i for future in-store demands in that review period, where

$$w_i^t = F_{\mathcal{D}_{is}^t}^{-1} \left(\frac{p_s}{h(T - t + 1) + p_s} \right), \quad \text{where} \quad \mathcal{D}_{is}^t = \sum_{\hat{t}=t+1}^{T} D_{is}^{\hat{t}} \tag{2.18}$$

We have developed a static fulfillment policy, as these thresholds can be evaluated at the start of the review period based on the demand forecasts. We formalize the TF policy in Algorithm 2. The MF policy places no such restriction on fulfillment, and can simply be recovered from Algorithm 2 by setting the thresholds w_i^t to be zero in step 1.

Note that the fulfillment heuristic is agnostic to current inventory levels and online demands. While including such information would be valuable, we show that such a simple policy, when combined with a good inventory heuristic which positions inventory in a calculated fashion, can provide considerable savings compared to naive strategies.

To evaluate the performance of the fulfillment policies, we compare them with the so-called hindsight-optimal policy. The cost of this policy can be evaluated through a linear program which minimizes the total cost in the review period, given that all uncertainty is realized at the beginning of the period. Given inventory levels, the cost of such a policy is a natural lower bound for the cost of any fulfillment policy, and we numerically show that the simple TF policy performs very well compared to this lower bound in Section 2.7.

2.7 Numerical Analysis

We employ a realistic setting to test the performance of the inventory and fulfillment heuristic solutions, based on a fictitious network embedded in mainland US. We shall mainly focus on the case with zero lead time and multiple fulfillment epochs.

We evaluate the total expected costs through a Monte-Carlo simulation with a sample size of 10^4, for two inventory heuristics - IIPH (integrated planning heuristic) and DIP (decentralized planning), and two fulfillment heuristics - MF (myopic) and TF (threshold-based). We mostly focus on comparing our combined heuristic, the $\langle$IIPH,TF$\rangle$ strategy, to the benchmark $\langle$DIP,MF$\rangle$ strategy, which represents a naive solution.

2.7.1 Network Setup

We take the locations of the stores to be at the most populous cities in mainland US (Wikipedia, 2016) and the OFCs are located according to the list of most efficient warehouses in the US, in terms of possible transit lead-times (Chicago Consulting, 2013). The shipping costs are calculated using the cost equation estimated by Jasin and Sinha (2015) based on UPS Ground shipping rates for an item weighing one pound: $s_{ij} = 9.182 + 0.000541 d_{ij}$, where d_{ij} is the distance in miles from region i to region j. We also perform sensitivity analysis for the slope of the shipping cost with respect to distance, to study the effect of shipping costs on the relative performance of our combined heuristic. Other cost parameters used are: $\bar{h} = 5$, $p_s = p_o = 100$, $s = 9.182$.

The review-period demands are taken to be independent and normally distributed with mean and standard deviations calculated based on the population of the cities. To study the effect of online market share (α) on the performance of the heuristic solutions, we take that the sum of the mean in-store and online demands in each region to be a fixed proportion of the cities' populations. This represents the average market size of the region, and the review-period mean in-store and online demands are calculated as $1 - \alpha$ and α proportions respectively of this mean market size in each region. The coefficient of variation of the review-period demands are fixed at 0.2.

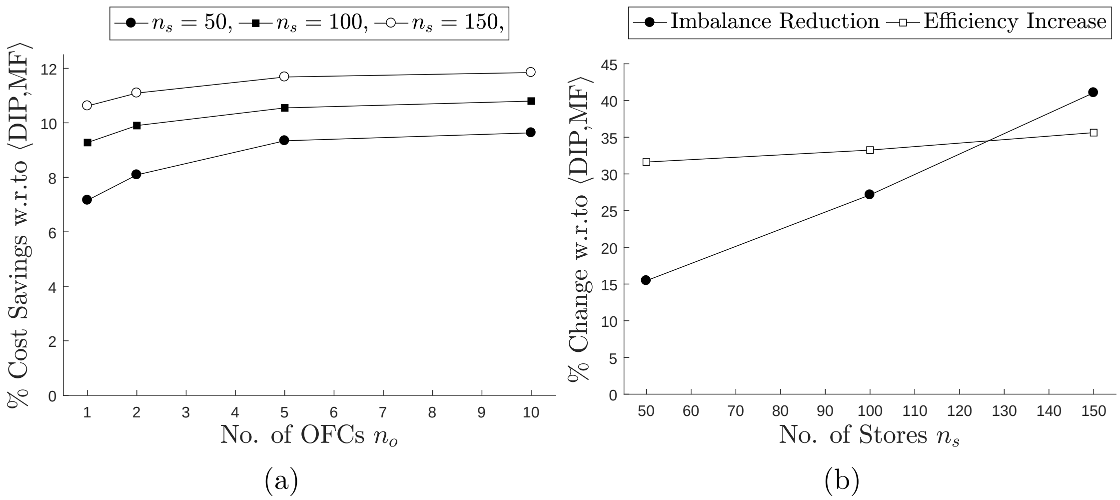

Figure 2.4: Shows the effect of network size on the performance of ⟨IIPH,TF⟩ compared to ⟨DIP,MF⟩, in terms of cost (left), inventory imbalance and inventory efficiency (right)

Demands are identical and independently distributed across fulfillment epochs, with parameters calculated from the review-period demands. In the base case, $\alpha = 0.5$ and $T = 5$, and we perform sensitivity analyses with respect to these parameters. Let n_s be the number of physical stores and n_o be the number of OFCs. An online order can be fulfilled from any physical store or OFC with available inventory. Further details on the numerical setup and a brief overview of the simulation process can be found in Appendix A.3.

2.7.2 Results

We tabulate the results obtained. We mainly focus on comparing the cost of the combined heuristic ⟨IIPH,TF⟩ to that of the naive strategy ⟨DIP,MF⟩. In some cases, to test the severity of assumptions made to derive the inventory heuristic, we compare ⟨IIPH,TF⟩ and ⟨DIP,TF⟩, keeping the fulfillment policy fixed.

2.7.2.1 Network Size.

As the network size increases, centralized inventory planning and strategic fulfillment can be valuable, as there is more flexibility in terms of options available in fulfillment. Figure 2.4a shows that increasing network size have a positive and marginally decreasing effect on the relative performance of the combined heuristic.

We also compare the strategies based on two important metrics, inventory imbalance and inventory efficiency, and the results are shown in Figure 2.4b for $n_o = 2$. Higher imbalance can lead to costly spillovers and local stockouts (Acimovic and

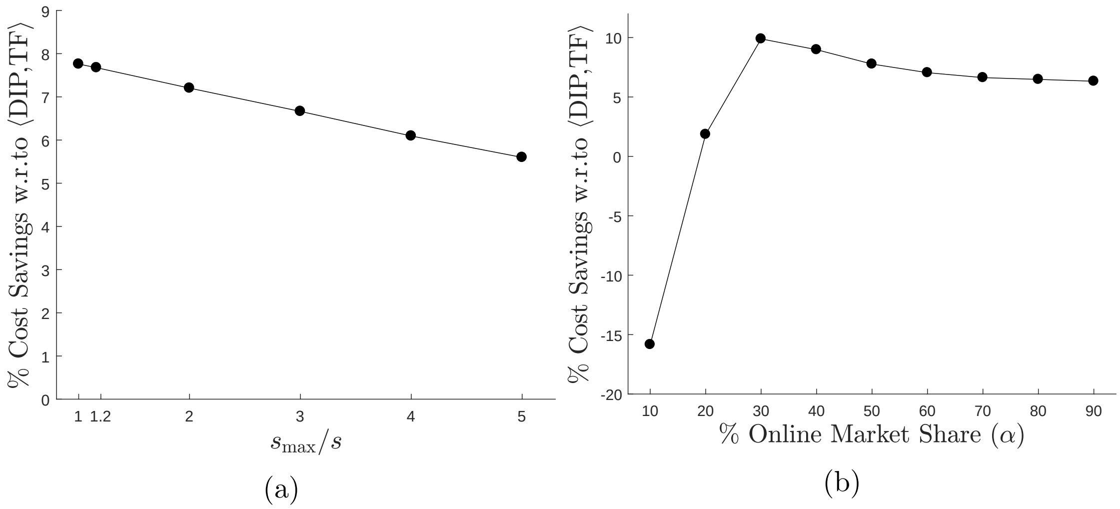

Figure 2.5: Shows the effect of the slope of cross-shipping costs with distance by varying the ratio $s_{\max}/s$ (left) and online market share (right) on the performance of ⟨IIPH,TF⟩ compared to ⟨DIP,TF⟩.

Graves, 2017), which in turn can cause markdowns in stores. We measure imbalance by recording the variance of ending inventory positions across locations at the end of each epoch, and taking the average value over the review period. Although this is different from the metric used by Acimovic and Graves (2017), it captures the essence of imbalance among locations in an omnichannel network. We see that our combined heuristic achieves a lower imbalance across locations as compared to the ⟨DIP,MF⟩ strategy, and this effect is more pronounced for larger networks.

We define another metric, inventory efficiency, as an equivalent measure for inventory turnover, calculated as the ratio of the total fulfilled demand to the average inventory level of the system in a review period (calculated as the mean of the starting inventory level and expected ending inventory at the end of the review period). Higher efficiency achieved by the heuristic stems from a reduction in inventory levels without a considerable decrease in service levels, due to planning in advance for cross-shipping. This offers a potential solution to decreasing trend in turnovers in the retail industry in recent years (Kurt Salmon, 2016).

2.7.2.2 Cross-shipping Costs and Online Market Share.

As discussed in Section 2.5.2, two major factors affect the inventory heuristic performance – shipping cost structure and online market share. For fixed fulfillment policy TF, we compare the ⟨IIPH,TF⟩ and ⟨DIP,TF⟩ strategies to understand the effect of these parameters on the inventory heuristic. We found similar results when comparing to the ⟨DIP,MF⟩ strategy.

We first vary the slope of shipping costs with respect to distance, thereby increasing the ratio $s_{\max}/s$ (value of 1.2 corresponds to the base case setting). As expected, the relative performance of the heuristic decreases as shipping costs become more sensitive to distance (Figure 2.5a). For a perspective, the costliest shipping option, the UPS Next Day Air, has a ratio of $s_{\max}/s$ less than 3 for shipping a 5lb package within mainland US. Hence the heuristic provides significant savings for most existing shipping cost structures.

Figure 2.5b shows the effect of online market share. As expected, we see that the heuristic performs worse than the decentralized solution when the online market share is low ($<$ 20%). This reflects the deficiency noted in the two-store case, as the inventory heuristic assumes that in-store demands are pooled across locations. When the online demand is very low compared to the in-store demand, the value from centralized planning is limited (as previously seen in Figure 2.3a), and the firm can simply resort to decentralized planning.

However, the heuristic provides a valuable alternative to the decentralized solution for products that have adequate online market shares: for example, books, computers and consumer electronics have an online market share of about 50% (FTI Consulting, 2015). Additionally, with rapidly increasing online sales, firms can obtain considerable savings through centralized inventory strategies, and for most cases, our heuristic serves as a viable proxy for inaccessible optimal decisions.

2.7.2.3 Number of Fulfillment Epochs (T).

By increasing the number of times online fulfillment decisions are made, we can closely model the continuous time case. We keep the total review-period demand parameters constant, and keep demands across fulfillment epochs independent and identically distributed. To reduce the computational burden associated with higher values of T, we use a smaller network with $n_s = 10,\ n_o = 2$.

The results are shown in Figure 2.6. In Figure 2.6a, we compare the MF and TF fulfillment strategies with IIPH inventory levels, against the hindsight optimal strategy HF, which makes fulfillment decisions with all uncertainty realized at the start of the review period.

As T increases, the MF policy is punished for failing to reserve inventory for future in-store demands (Figure 2.6a). The TF policy on the other hand proves to be a simple but effective fulfillment strategy, achieving costs within 0.5% of the HF lower bound.

For a fixed fulfillment policy TF, we compare the ⟨IIPH,TF⟩ and ⟨DIP,TF⟩ strate-

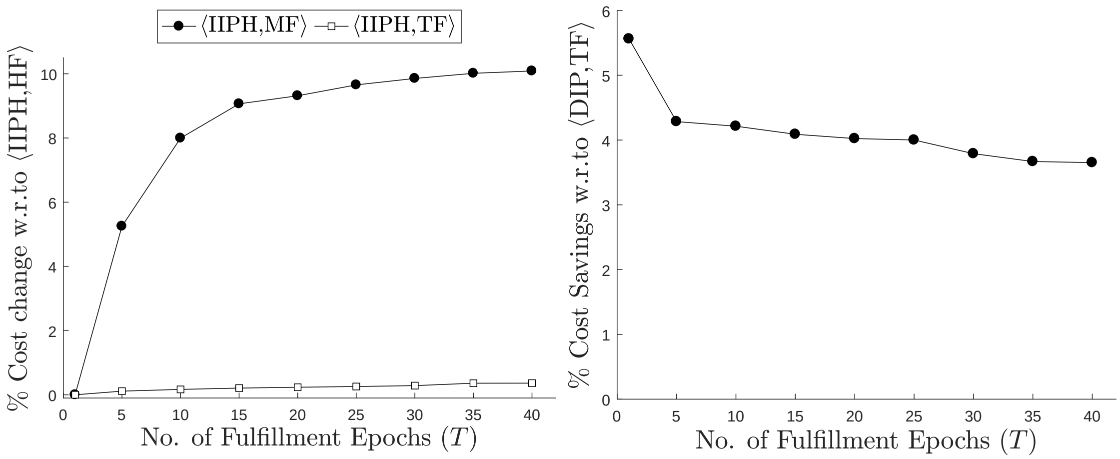

Figure 2.6: Shows the effect of increasing the number of fulfillment epochs in a single review period on the cost of fulfillment policies with respect to the hindsight optimal policy (left) and the performance of ⟨IIPH,TF⟩ compared to ⟨DIP,TF⟩ (right).

gies in Figure 2.6b, and see that the effect of increasing T has a decreasing effect on the relative performance of the inventory heuristic.

Finally, we note that our heuristics are extremely scalable with respect to network size - for a network with $n_s = 150$, $n_o = 10$ and $T = 5$, calculating the inventory levels using the heuristic takes only around 10 seconds, and the calculation of fulfillment thresholds takes around 2 minutes. Real-life retail networks are often much bigger in size – for instance, Target ships online orders from more than 1000 stores (Lindner, 2016), and our heuristic can provide considerable improvements compared to traditional strategies in most cases.

2.8 Conclusion

Despite numerous retailers struggling with the operational problems posed by omnichannel retailing, the area has received comparatively less attention in literature. Our research addresses an important facet of omnichannel retailing — inventory management, by demonstrating the value in utilizing the pooling benefits offered by omnichannel retailing, through a combined inventory and fulfillment policy.

Our heuristic policies, though derived from a complicated multi-location and multi-period model, are quite generalizable. We can extend our analysis to demands originating from abstract regions, by treating them as OFCs that carry zero inventory. Disparity in service costs at OFCs and stores can also be taken into account by using

s_o, the service cost from OFCs, instead of s in Equation 2.16, as inventory planning is done separately for OFCs. We still need to make the assumption that demands from a region with an omnichannel store can be fulfilled from that store with the least cost. Otherwise, the demand at this store will be assigned to be fulfilled from the online FC with the least fulfillment cost, which can lead to different first order conditions in inventory planning for the online FC.

We can also extend the heuristic solutions to the case of positive leadtimes as follows: assuming each location i has a replenishment leadtime of L_i review periods, the total planning horizon for order-up-to policies is (L_i+1) review periods, or equivalently, $(L_i + 1)T$ fulfillment epochs for each location. Using the total demands during the planning period for each location instead of review period demands, we can directly extend our inventory heuristic to set order-up-to levels for each location.

For the fulfillment heuristic, an additional threshold for inventory position needs to be calculated based on future in-store demands in the remainder of the current planning horizon, which can also be computed based on a simple newsvendor formula. Online fulfillment from a location is temporarily stopped in an epoch when either threshold is violated.

An important direction for future research is to include multiple classes of online demand, especially in-store pickups, which is a popular mode of omnichannel fulfillment. A heuristic control for managing multiple products is also an interesting and important extension to be considered. Future research may also focus on further extensions such as capacities and stochastic leadtimes. We believe that our framework provides a platform to build further complexities on, which can yield important decision support tools for the industry.

CHAPTER 3

The Distribution-free Inventory Problem for E-commerce Fulfillment Networks

3.1 Abstract

With rapidly increasing e-commerce sales, firms are leveraging a network of inventory nodes to fulfill online orders by its customers. This results in demand spillovers: demand is first fulfilled using the nearest node, but then demand can spill over to other nodes when there are stockouts. Inventory planning for e-commerce is challenging due to this complex nature of fulfillment. Further, e-commerce demand is difficult to estimate due to flash promotions, recommendation engines, and other strategies typically employed by e-tailers. We address this by solving a distributionally robust inventory problem where the fulfillment network has a *nested* hierarchy and the firm only knows the mean and covariance of the demand. The objective is to minimize the worst-case expected total cost of procurement, fulfillment, and penalties for any unmet demand.

If there are two nodes in the nested network, we derive a tight bound for the expected cost of unmet demand that only requires mean, variance, and correlation of the demands in the two locations. We show that this new bound is significantly tighter than the well-known Scarf bound in the regime when inventory levels are low and demand spillovers are likely to occur. For general nested fulfillment networks, the problem is NP-hard. We develop a heuristic by deriving an upper bound to the expected cost of unmet demand that ensures the nested structure of fulfillment is preserved. The heuristic is computationally tractable since it relies on solving a semidefinite program with dimension quadratic in the number of nodes. We also develop an algorithm to approximate any general distance-based fulfillment cost structure with a nested structure, which we show in numerical experiments to result in tight approximations to the expected total cost.

3.2 Introduction

As e-commerce continues to grow rapidly (Zaroban, 2018), retail firms are equipping themselves with the ability to fulfill online orders from multiple inventory nodes (stores, fulfillment centers, etc.) in their network. In many modern e-commerce fulfillment systems, a customer order may be fulfilled by a shipment originating from any inventory node in the retail network. Shipping rates are proportional to distances, so typically the nearest node to the customer location is chosen first. However, in the event of stockouts, the demand would "spill over" to other nodes (Acimovic and Graves, 2017), guaranteeing that the demand is not lost while there is inventory still remaining in the retail network. This is reminiscent of the flexibility offered in brick-and-mortar retail with periodic store transshipments. However, a key difference is that, since the transaction is conducted through a virtual store, e-commerce always allows for this flexibility without the need for inventories to be first prepositioned in a customer's location.

Allowing demand spillovers essentially pools the geographically separate inventories. Hence e-commerce fulfillment requires less inventory than what would be recommended by traditional decentralized inventory models that do not account for demand spillovers. Therefore, in order to reduce the burden of carrying too much inventory, e-commerce inventory planning must use network-based models that capture fulfillment flexibility. However, there are several challenges in inventory planning for e-commerce retail.

One such challenge is demand estimation, since e-commerce demand often has a higher variance than brick-and-mortar demand. Reasons for this include the ease with which online customers could choose to purchase from any of multiple competing e-tailers, the use of dynamic price-matching strategies and flash promotions, recommendation engines that manipulate click-streams, etc. Hence the empirical distribution of past sales is a less reliable estimate of the distribution of future e-commerce demands.

A common heuristic is to assume that the underlying uncertainty (i.e., the vector of demands in customer locations) has a multivariate normal distribution. Such an assumption can help an inventory planner in two ways. First, describing a multivariate normal demand vector requires information about only the first two moments, namely the mean and covariance matrices, which can be reliably estimated. Second, the normal distribution lends itself to simple analytic solutions (e.g. Dong and Rudi, 2004). However, the normal distribution assumption can lead to solutions that overestimate pooling benefits if the true demand distribution is non-normal. Eppen (1979) showed that for demand distributions that are of 'light-tailed nature' (including the normal

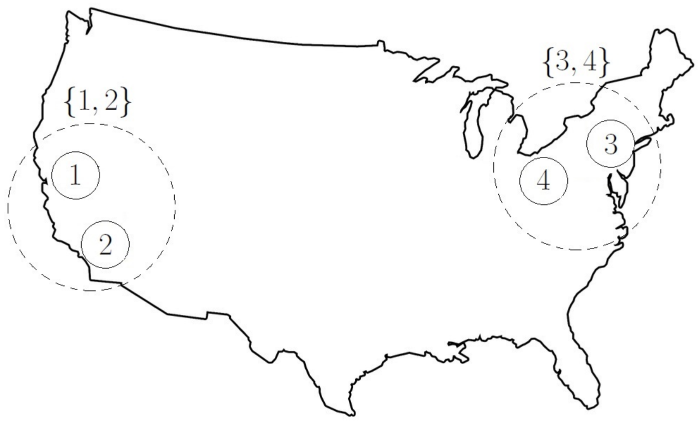

Figure 3.1: Example of a nested network: 1 - San Francisco, 2 - Los Angeles, 3 - New York, 4 - Pittsburgh.

distribution), pooling can lead to savings in expected cost that scale with $\sqrt{n}$, where n is the number of random demands being pooled. However, Bimpikis and Markakis (2015) show that the pooling benefits can scale significantly lower than $\sqrt{n}$ when the demand distribution is heavy-tailed. There have been earlier studies that show evidence of real-life demands exhibiting non-normal distributions: Bimpikis and Markakis (2015) give empirical evidence of heavy-tailed demands for movies at Netflix and shoes at a major retailer, Agrawal and Smith (1996) show that the negative binomial distribution fits the sales data for men's slacks at a major retailer better than Poisson or Normal.

In this paper, we address the challenge in demand estimation by adopting a *distributionally robust* approach. Since this approach assumes an adversary always chooses a distribution resulting in the highest expected cost, it leads to robust decisions for a firm that has access to only partial information of the demand distribution. Specifically, we assume that the only information known about the demand vector is its mean and covariance.

Another challenge in e-commerce inventory planning is the complex nature of fulfillment in the retail network due to demand spillovers. Inspired by e-commerce fulfillment, we assume that the fulfillment network is *nested*, which results in a natural hierarchy in the demand spillovers. Figure 3.1 is an example of a nested fulfillment network with four inventory nodes (San Francisco, Los Angeles, Pittsburgh, and New York).

Consider an online customer located in New York. When the nearest fulfillment center (node 3, New York) stocks out, the next best option is to fulfill from the nearby node in Pittsburgh (node 4). If this node is also stocked out, fulfillment from San

Francisco (node 1) or Los Angeles (node 2) will roughly have the same cost—one that is higher than the cost of fulfillment from Pittsburgh—as these locations are farther away. This is similarly true for a customer in San Francisco—node 1 is preferred, followed by node 2, then by nodes 3 and 4. Thus, to fulfill an unmet demand, nodes or groups of nodes are considered progressively based on their proximity (in terms of distance or cost) to the unmet demand. We refer to cost structures that induce such hierarchical levels of fulfillment as *nested fulfillment structures*, and networks with such cost structures as *nested networks*.

The focus of this paper is the problem of deciding the inventory levels for firm that fulfills sales of a single product through a nested network of multiple inventory nodes (such as warehouses, stores, etc.) when the firm only knows the mean and the covariance matrix of the demand vector. The firm incurs inventory purchasing costs, penalty costs for unmet demand, and fulfillment costs. We consider a single-period model where network fulfillment occurs after the demands across multiple locations are realized.

Main Results and Contributions

1. The distributionally robust inventory problem with a single node can be analytically solved using the well-known result by Scarf (1958) that provides a tight bound for the expected cost of unmet demand which only requires mean and variance. We extend this result to the case of demand spillovers by deriving a tight bound for the expected cost of unmet demand in a nested fulfillment network with two nodes. This bound only requires the mean, variance, and correlation of demands in the two locations. This yields a closed-form expression for the optimal inventory levels in the distributionally robust problem under a nested network with two nodes. We show that the new bound is significantly tighter than the Scarf bound in the regime of low inventory levels where demand spillover is most likely to occur.

2. We introduce the class of nested fulfillment structures that greatly simplify the computation of the expected cost of the firm. Specifically, we derive a closed-form expression for the expected cost under this general class. We develop a simple algorithm (based on hierarchical agglomerative clustering) to approximate any general distance-based fulfillment cost structure as a nested fulfillment structure. We show empirically that this structure tightly approximates the expected total fulfillment cost under a variety of distributions. A nested fulfillment structure not

only yields tractability for our robust problem, but also for stochastic systems.

3. For general nested fulfillment networks with multiple nodes, the distributionally robust problem is NP-hard. We develop a heuristic for the robust inventory problem by deriving an upper bound to the expected cost of unmet demand in a nested network with demand spillover that requires only mean and covariance data. The bound is constructed to ensure that the nested structure of fulfillment is preserved. Moreover, it is computationally tractable since it relies on solving a semidefinite program (SDP) with dimension $\mathcal{O}(n^2)$, where n is the number of inventory nodes. By means of numerical experiments, we show that the distributionally robust heuristic can lead to significant savings in expected cost as compared to stochastic solutions that assume an incorrect distribution.

3.2.1 Literature Review

Our study is related to the literature on inventory pooling, since e-commerce fulfillment virtually pools geographically separate inventories. We mention only works that establish the importance of distributional properties of the stochastic demand in pooling. Eppen (1979) showed that if n normal and uncorrelated demands are pooled, the benefit from inventory pooling is $\sqrt{n}$, and the benefit decreases with increasing positive correlation among demands. Corbett and Rajaram (2006) extended Eppen's result to more general distributions. Yang and Schrage (2009) study various cases of 'inventory anomaly' (a situation where pooling leads to an increase in inventory as opposed to a reduction), one of which is for right-skewed demand distributions with product substitution. Berman et al. (2011) found through numerical simulations that the normal distribution misestimates the benefits of pooling stemming from a reduction in variance. Bimpikis and Markakis (2015) find that the benefit from pooling under heavy-tailed demand distributions can be significantly lower than $\sqrt{n}$. Specifically, they show that the benefit from pooling decreases as the tail of the demand becomes heavier. All these studies indicate that pooling benefits crucially depend on the distribution of the demands being pooled.

The inventory problem for e-commerce network fulfillment is mathematically identical to reactive lateral transshipments in brick-and-mortar retail networks, which have been discussed in great detail in the literature (for a review, refer to Paterson et al., 2011). Two features make the transshipment problem difficult to analyze. First, for more than two locations, analytically optimal solutions become elusive as a linear program recourse is needed to model the network flow problem among multiple locations

(Robinson, 1990). Second, in a multi-period setting with leadtimes, the optimal transshipment decisions are intractable due to complexity in the state space even for the two location problem, as it can be *ex-post* optimal to reserve inventory for future use rather than transshipping to another location (Tagaras and Cohen, 1992). In this paper, we focus on the multi-location problem in a single period setting. Dong and Rudi (2004) consider a similar setting for the special case where the transshipment cost between any two locations is a constant. However, none of these works consider a distributionally robust framework.

Distributionally robust inventory problems have a long history, dating back to Scarf (1958) who considers the classic newsvendor problem with only mean, variance and support information. He shows that a two-point demand distribution results in the smallest expected profit given the inventory level, and derives the optimal inventory level of the robust max-min problem. After 35 years, Gallego and Moon (1993) extended this framework to multiple products with moments (mean and variance) of the marginal distributions. Recourse decisions were also studied in Gallego and Moon (1993) (by allowing an additional order to be placed after demand is realized) and in Mostard et al. (2005) (by allowing returned products to be resold if there is sufficient demand).

The multiple-product setting was further extended by Hanasusanto et al. (2015) to include mean and covariance of the joint distribution, and they show that the resulting distributionally robust newsvendor problem is NP-hard. Natarajan and Teo (2017) develop a tractable heuristic for this problem in the form of a semi-definite program. They achieve this by expressing piecewise linear terms through integer variables, and relaxing the equivalent completely positive program into a semi-definite program. They relax the integrality constraints through a boolean quadric polytope, previously studied by Padberg (1989). Natarajan et al. (2017) use similar techniques to derive tractable heuristics for the multi-item newsvendor with known mean, covariance and semivariance information, which additionally captures asymmetry in the distribution.

Our work is related to distributionally robust inventory problems over a network that allow for recourse network flows after the demand is realized. Chou et al. (2006) addresses such a problem by assuming that the transshipment quantities are linearly dependent on some primitive uncertainties with known support, forward and backward deviations. Linear decision rules are common in approximating multi-stage programs (Ben-Tal et al., 2004), however they are not necessarily optimal. In contrast to their work, our study models the optimal flows by directly approximating the nested fulfillment cost structure, and assuming known mean and covariance matrices of the demands, rather than allowing the demands to depend on primitive uncertainties that may be hard to estimate. Recently, Yan et al. (2018) used techniques from Natarajan et al. (2011) to obtain an exact reformulation of the distributionally robust min-cost network flow problem as a completely positive program, given the mean and covariance information and under restricted 0-1 edge costs. Our paper in contrast allows for general nested fulfillment cost structures under which the problem becomes more difficult, but which are more appropriate for e-commerce network fulfillment.

We note that the nested structure is similar to a tree metric, which has been studied in the Computer Science literature. Bartal (1998) and Fakcharoenphol et al. (2004) consider probabilistic approximations of metric spaces using a tree metric, where the nodes of a graph forms the leaves of a rooted tree. The distance between two nodes is approximated by the sum of edge weights on the shortest path between them, and probabilistic approximations of $\mathcal{O}(\log n)$ are available. A distinction in the quality of approximation is that these papers study how closely the generated tree metric approximates the actual distance metric, whereas we consider the closeness in approximating the expected fulfillment cost.

Finally, our study can also be related to the growing literature in e-commerce inventory and fulfillment optimization. Acimovic and Graves (2017) show that decentralized inventory solutions can lead to costly spillover effects, and perform poorly compared to network-based policies. Govindarajan et al. (2018) consider joint optimization of inventory and fulfillment decisions in an omnichannel setting where in-store demands cannot be flexibly fulfilled from other locations, whereas e-commerce demand can. More generally, we relate to the problem of capacity allocation in networks with flexibilities, which has been recently gaining relevance due to applications in e-commerce; Lyu et al. (2017) study the optimal allocation policy given target fill rates, and find that the required capacity levels in a long-chain network are close to the levels in a fully flexible network. DeValve et al. (2018) study the benefit of adding fulfillment flexibility to a large online retailer's network by combining an allocation policy based on a stochastic program with a fulfillment policy which restricts the spillover demand that is fulfilled. Given the volatility in online customer behavior, we contribute to the above streams of literature by studying the distributionally robust inventory allocation problem where only lower order moments of the demands can be reliably estimated.

3.2.2 Preliminaries

For any integer n, we use notation $[n]$ to denote the set $\{1, 2, \ldots, n\}$. We denote by $2^{[n]}$ the power set of $[n]$, defined as the set of all subsets of $[n]$ including the empty set. We denote by $\mathbf{e}_n$ the n-dimensional column vector of all ones, where we drop the subscript if the size is clear from the context. We denote by $\mathbf{I}_n$ the identity matrix of size $n \times n$, and $\mathbf{0}_{m,n}$ the zero matrix of size $m \times n$.

We denote by $\Re$ the set of real numbers, and by $\Re_{\geq 0}$ the set of nonnegative real numbers. We similarly denote by $\Re^n$ the set of n-dimensional vectors of real numbers, and $\Re^n_{\geq 0} := \{\mathbf{x} \in \Re^n | \mathbf{x} \geq 0\}$ its subset of nonnegative vectors. For a scalar variable $x \in \Re$, we define $x^+ := \max(0, x)$ as the positive part of x. For a column vector $\mathbf{x} = (x_i) \in \Re^n$, we define $\mathbf{x}^+ := (x_i^+)$ as the positive part of each element in $\mathbf{x}$. We write $\mathbf{A} \succeq 0$ if a square matrix $\mathbf{A}$ is symmetric positive semidefinite. We write $\mathbf{B} \geq 0$ if all entries of the matrix $\mathbf{B}$ are nonnegative.

3.2.3 Outline

The rest of the paper is organized as follows. In Section 3.3, we describe the model, and introduce the nested fulfillment cost structure. Section 3.4, presents closed-form bounds on the expected unmet demand in a nested network with two nodes, and contrasts the bound to Scarf. Section 3.5 is devoted to developing computationally tractable heuristics for the multi-location case for nested networks. In Section 3.6, we develop an algorithm to recover the nested structure from a general cost structure, and empirically test the strength of approximation. In Section 3.7, we analyze the multi-location heuristic solutions numerically to understand the effect of additional information and to test the performance of the heuristic solutions. Extensions and future directions follow in Sections 3.8 and 3.9.

3.3 The Model

Consider a firm managing the inventory in a network of nodes (e.g., stores or warehouses) that support sales of a product during a selling horizon. We assume that there is no inventory replenishment during the selling horizon (as is typically the case when the horizon is short compared to the procurement lead time), so the firm only needs to decide the initial inventory levels. We assume that demand for the product originates from n geographic regions. For simplicity of the model, we consider a fulfillment network with n inventory nodes, where one node is located in each customer region. Our framework can easily be extended to general networks by modeling customer regions without a fulfillment node as zero-inventory nodes.

3.3.1 Nested fulfillment networks

We assume that the firm incurs a per-unit "fulfillment" cost (e.g. shipping cost and/or handling cost) for using an inventory node to fulfill demand in a customer region, and this cost depends on both the inventory node and the customer region. Specifically, if a unit of demand in region $j \in [n]$ is met with inventory from the node in the same region j, then the fulfillment cost is s_{jj}, where $s_{jj} > 0$; if met with inventory from a node $i \neq j$, then the fulfillment cost is s_{ij}. We assume that $s_{ij} = s_{ji}$ and $s_{ij} \geq s_{jj}$ for all i, j. We further assume that not all fulfillment costs are the same (otherwise, the location of inventory does not matter). We denote the $n \times n$ matrix of per-unit fulfillment costs as $\mathbf{S} = (s_{ij})$.

Motivated by e-commerce, we focus our attention to fulfillment networks where the cost structure has a nested hierarchy. We refer to these as *nested fulfillment networks* since they result in a hierarchical ordering of nested node sets. Before introducing the general definition, we first provide a simple example of such a network.

Example 1 (A 3-level nested network) *Consider the network discussed earlier in Figure 3.1 with $n = 4$ (nodes) regions. Suppose that in-location fulfillment has a per-unit cost $s_0 > 0$ (using our notation, $s_{jj} = s_0$ for $j = 1, 2, 3, 4$). Fulfillment from region 1 to 2 (and vice versa) incurs cost s_1, and fulfillment between regions 3 and 4 also has cost s_1 (i.e., $s_{12} = s_{34} = s_1$). Fulfillment between any other pairs of regions has cost s_2 (i.e., $s_{13} = s_{14} = s_{23} = s_{24} = s_2$). If $s_0 < s_1 < s_2$, the fulfillment structure induces the nested hierarchy illustrated in Figure 3.2 since a higher level fulfillment is used only if fulfillment in a lower level is not possible due to lack of inventory. Since there are three levels to this hierarchy, we refer to this as a* 3-level nested fulfillment network.

We now describe a general L-level nested fulfillment network. For a given level ℓ, where $\ell = 0, 1, \ldots, L-1$, the regions are partitioned into n_ℓ sets, where $n_\ell \leq n$. (In the previous example, there are four sets (i.e., $\{1\}, \{2\}, \{3\}, \{4\}$) in the level 0 partition, and two sets (i.e., $\{1, 2\}, \{3, 4\}$) in the level 1 partition). We denote the n_ℓ sets in partition ℓ as $\{\mathcal{I}_1^{(\ell)}, \mathcal{I}_2^{(\ell)}, \ldots, \mathcal{I}_{n_\ell}^{(\ell)}\}$. By definition, the sets in a partition must cover all regions, and that the intersection of any two sets is empty.

The fulfillment network is nested because of the following property: any set in level ℓ is the union of sets in the preceding level $\ell - 1$. That is, $\mathcal{I}_k^{(\ell)} = \cup_{m \in \mathcal{K}_k^{(\ell)}} \mathcal{I}_m^{(\ell-1)}$ where

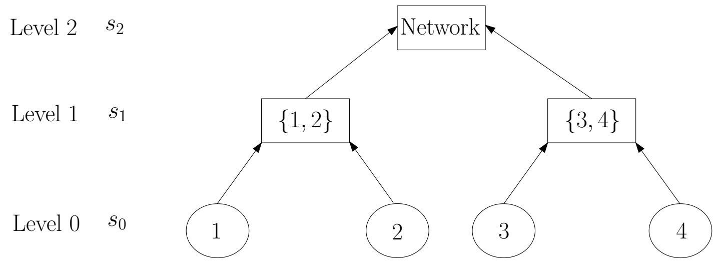

Figure 3.2: Example of a 3-level nested fulfillment cost structure with four warehouses $W_1, .., W_4$.

$\mathcal{K}_k^{(\ell)}$ are the indices of the level $\ell - 1$ children of $\mathcal{I}_k^{(\ell)}$. Note that $n_0 > n_1 > \cdots > n_{L-1}$, since there are fewer sets in higher order levels. We assume that for level 0, $n_0 = n$ and $\mathcal{I}_i^{(0)} = \{i\}$ for all $i \in [n]$. For the final level $L-1$, $n_{L-1} = 1$ and $\mathcal{I}_1^{(L-1)} = [n]$. Any L-level nested fulfillment structure can be represented as an L-level tree similar to Figure 3.2 (another example is Figure 3.7b).

The nested hierarchical structure can also be represented through the assignment matrices $\Xi = \{\mathbf{E}_0, \mathbf{E}_1, \ldots, \mathbf{E}_{L-1}\}$. We define the level ℓ assignment matrix $\mathbf{E}_\ell$ as the binary matrix of size $n_\ell \times n$ where the (k, i) entry is equal to 1 if and only if $i \in \mathcal{I}_k^{(\ell)}$. Note that $\mathbf{E}_0$ is the $n \times n$ identity matrix, and that $\mathbf{E}_{L-1}$ is the row vector of all ones.

To complete the description of the nested fulfillment network, we next discuss the fulfillment costs. If two regions are in set $\mathcal{I}_k^{(\ell)}$, then the per-unit cost of fulfillment between the two regions is $s_{\ell,k}$. (In Example 1, regions 1 and 2 are in the same level 1 set, so the fulfillment cost between them is s_1. Note that in the example, all level 1 costs are equal; however, in general, we allow sets in the same level to have different costs.) To induce the nested hierarchy, we assume that it is less costly to fulfill demand using fulfillment in lower levels. Mathematically, if $k^{(\ell)}(i)$ is the level ℓ set index of region i, then we assume that $s_{0,k^{(0)}(i)} \leq s_{1,k^{(1)}(i)} \leq \cdots \leq s_{L-1}$. We denote by $\mathbf{s} = \{s_{\ell,k}\}$ the set of all fulfillment costs.

The nested hierarchy, Ξ, and the fulfillment costs, $\mathbf{s}$, fully characterize the nested fulfillment network. As we later discuss in Section 3.3.4, this class of fulfillment structures can approximate general fulfillment cost structures.

3.3.2 The cost function

At the start of the selling horizon, the firm decides the vector of initial inventory levels $\mathbf{y} = (y_i)$ to fulfill demands that arrive throughout the selling season. We assume that fulfillment is done at the end of the selling horizon, so that the problem can be approximated as a single period problem.[1] The single-period approximation allows us to study inventory pooling in a distribution-free context by side-stepping the complications that arise from dynamic fulfillment decision-making, and focusing on the inventory decisions. In Section 3.7, we empirically show that the single-period assumption is a tight approximation for the dynamic setting under the common practice of myopic fulfillment.

We let $\tilde{d}_j$ denote the stochastic demand in region $j \in [n]$ (a variable with a tilde placed on top refers to a random variable; the same variable without the tilde is a particular realization). The vector of stochastic demands in the n customer regions is $\tilde{\mathbf{D}} = (\tilde{d}_j)$ (unless otherwise stated, any vector is a column vector). In the single period setting, the vector of customer demands $\mathbf{D} = (d_j)$ is realized at the end of the period, and the firm fulfills the demand with the objective of minimizing the total newsvendor cost (i.e., penalty, overage, and fulfillment costs). Unmet demand in any region incurs a per-unit penalty cost p, while unsold inventory in any node incurs a per-unit overage cost h. Without loss of generality, we assume that any fulfillment cost in the nested network does not exceed $p + h$; that is, the firm prefers to fulfill any unmet demand if there is available inventory in the network. (If $s_{\ell,k} > p + h$, then the children nodes of set k will never use level ℓ fulfillment, and the problem decomposes as each level $\ell - 1$ child of set k can be removed to form new networks.)

Mathematically, if z_{ij} units of inventory from node i is used to satisfy demand in region j, then the fulfillment quantities $\mathbf{Z} = (z_{ij})$ are determined through the following network flow problem:

$$
\begin{aligned}
C(\mathbf{y}, \mathbf{D}) := \underset{\mathbf{Z} \geq 0}{\text{minimize}} \quad & h \cdot \sum_{i \in [n]} \left(y_i - \sum_{j \in [n]} z_{ij} \right) + p \cdot \sum_{j \in [n]} \left(d_j - \sum_{i \in [n]} z_{ij} \right) + \sum_{i \in [n]} \sum_{j \in [n]} s_{ij} z_{ij} \\
\text{subject to} \quad & \sum_{j \in [n]} z_{ij} \leq y_i, \quad \forall i \in [n] \\
& \sum_{i \in [n]} z_{ij} \leq d_j \quad \forall j \in [n]
\end{aligned} \tag{3.1}
$$

where the terms in the objective are the overage cost, the penalty cost, and the fulfillment cost, respectively. The first constraint specifies that the units used to fulfill demand from inventory node i should not exceed the initial stocking level y_i. The

[1] While taking into account dynamic fulfillment decisions is more realistic, the problem becomes complicated due to the well-known curse of dimensionality (see Tagaras and Cohen (1992) for a stochastic system, and Ben-Tal et al. (2004) for the robust system). Multi-location considerations often complicate the problem further by adding complexity to the action space, as a linear program recourse is needed to make fulfillment decisions.

second constraint specifies that the total units used to fulfill demand in region j must not exceed d_j.

Since the fulfillment network has a nested hierarchical structure, we are able to express the total cost (3.1) in closed form.

Lemma 3.3.1 *Under the L-level nested fulfillment cost network,*

$$C(\mathbf{y}, \mathbf{D}) = h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{D}) + \mathbf{s}_0^\top \mathbf{D} + \sum_{\ell=0}^{L-1} \eta_\ell^\top (\mathbf{E}_\ell \mathbf{D} - \mathbf{E}_\ell \mathbf{y})^+ \tag{3.2}$$

where $\eta_{L-1} = p + h - s_{L-1}$ and, for $\ell \leq L-2$, $\eta_\ell = (\eta_{\ell,k})_{k \in [n_\ell]}$ with $\eta_{\ell,k} = s_{\ell+1, m^{(\ell+1)}(k)} - s_{\ell,k}$ where $m^{(\ell+1)}(k)$ is the index of the level $\ell + 1$ parent of set $\mathcal{I}_k^{(\ell)}$.

Note that $\eta_{\ell,k} = s_{\ell+1,m^{(\ell+1)}(k)} - s_{\ell,k} \geq 0$, and can be interpreted as the marginal benefit of fulfillment in level ℓ instead of level $\ell + 1$ of any demand occuring in $\mathcal{I}_k^{(\ell)}$. Similarly, η_{L-1} is the marginal benefit using a unit of inventory for fulfillment with the highest cost instead of holding onto it.

Proof. To obtain the fulfillment cost of a demand realization $\mathbf{D}$ in an L-level nested fulfillment cost structure, we sum the fulfillment costs in each level. The total fulfillment cost in level 0 is $\sum_{i \in [n]} s_{0,i} \cdot \min(d_i, y_i) = \mathbf{s}_0^\top \mathbf{D} - \sum_{i \in [n]} s_{0,i} \cdot (d_i - y_i)^+$, where $\mathbf{s}_0 = (s_{0,i})_{i \in [n]}$. For $\ell = 1, 2, \ldots, L-1$, the total units of demand in regions of $\mathcal{I}_k^{(\ell)}$ fulfilled at level ℓ (at a per-unit fulfillment cost $s_{\ell,k}$) is

$$\underbrace{\sum_{m \in \mathcal{K}_k^{(\ell)}} \left(\sum_{i \in \mathcal{I}_m^{(\ell-1)}} d_i - \sum_{i \in \mathcal{I}_m^{(\ell-1)}} y_i \right)^+}_{\text{unmet demand in } \mathcal{I}_k^\ell \text{ after level } \ell - 1} - \underbrace{\left(\sum_{i \in \mathcal{I}_k^{(\ell)}} d_i - \sum_{i \in \mathcal{I}_k^{(\ell)}} y_i \right)^+}_{\text{unmet demand in } \mathcal{I}_k^{(\ell)} \text{ after level } \ell},$$

where $\mathcal{K}_k^{(\ell)}$ are all level $\ell - 1$ children of set $\mathcal{I}_k^{(\ell)}$. Since $p + h$ is strictly greater than all fulfillment costs, then the penalty cost is $p \cdot \left(\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y}\right)^+$, and the overage cost is $h \cdot \left(\mathbf{e}^\top \mathbf{y} - \mathbf{e}^\top \mathbf{D}\right)^+$. Therefore, the total cost (overage, penalty and fulfillment) is equal to

$$\begin{aligned} C(\mathbf{y}, \mathbf{D}) = h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{D}) + \mathbf{s}_0^\top \mathbf{D} + (p + h - s_{L-1}) \cdot \left(\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y}\right)^+ \\ + \sum_{\ell=0}^{L-2} \sum_{k \in [n_\ell]} (s_{\ell+1,m^{(\ell+1)}(k)} - s_{\ell,k}) \cdot \left(\sum_{i \in \mathcal{I}_k^{(\ell)}} d_i - \sum_{i \in \mathcal{I}_k^{(\ell)}} y_i \right)^+, \end{aligned}$$

where $m^{(\ell+1)}(k) \in [n_{\ell+1}]$ is the level $\ell+1$ parent of set $k \in [n_\ell]$. Using compact notation with the parameters η_ℓ and assignment matrices $\mathbf{E}_\ell$, we obtain the lemma. ■

3.3.3 Distributionally robust framework

If the firm knew that the stochastic multi-location demand $\tilde{\mathbf{D}}$ has a joint distribution $f : \Re^n \mapsto \Re^+$, then the firm will choose the initial inventory vector $\mathbf{y}$ so as to minimize the expected total cost. Mathematically, the firm's problem is equivalent to solving the two-stage stochastic program:

$$\min_{\mathbf{y}\geq 0} \ \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}}) \right], \tag{3.3}$$

where $\mathbb{E}_f$ is the expectation operator under the joint probability distribution f. Note that the objective is to minimize the expected total cost, where $C(\mathbf{y}, \mathbf{D})$ is as described in (3.2). Since inventory in one location can be routed to meet demand in another location, we will refer to (3.3) as the *multi-location newsvendor problem with inventory risk pooling*. Note that the above problem can be numerically solved as a linear program either by sample average approximation using large enough number of samples, or by approximating the joint distribution by a discrete distribution. Under the nested structure, we can also obtain closed-form first order conditions, which can be solved numerically to yield the optimal solution.

In reality, however, firms do not have a complete description of the joint distribution of the multi-location demands. At best, the firm may only have partial information about the distribution. We will assume that the firm only has knowledge of the mean vector $\mathbf{m}$ and the covariance matrix $\mathbf{\Sigma}$. We chose this information set as in practice, e-commerce firms may have a good sense of how demands across locations are correlated. Moreover, it is known from the literature that covariance across locations is crucial for decisions that take pooling benefits into account.

As discussed in Section 3.2, if the firm assumes a particular distribution for the demand vector, say a multivariate normal distribution with parameters $(\mathbf{m}, \mathbf{\Sigma})$, then the optimal decision resulting from the stochastic program (3.3) may be suboptimal with a high expected cost under the true (unknown) demand distribution. To protect against such cases, we adapt a minmax distributionally robust approach (Scarf, 1958; Gallego and Moon, 1993; Hanasusanto et al., 2015; Natarajan et al., 2017) that aims to choose inventory levels $\mathbf{y}$ to minimize the maximal expected cost over all demand distributions consistent with the information known to the firm.

To understand the minmax robust approach, assume that after the firm makes a decision on the inventory levels $\mathbf{y}$, an "adversary" is able to choose a joint distribution f that results in the highest expected cost $\mathbb{E}_f(\mathbf{y}, \tilde{\mathbf{D}})$ for the firm. However, the adversary cannot choose just any f; consistent with the known mean and covariance, it has to belong to the distribution set:

$$\mathcal{F}_{\geq 0} := \left\{ f : \Re^n_{\geq 0} \mapsto \Re_{\geq 0} \mid \mathbb{E}_f(1) = 1,\ \mathbb{E}_f\left(\tilde{\mathbf{D}}\right) = \mathbf{m},\ \mathbb{E}_f\left(\tilde{\mathbf{D}}\tilde{\mathbf{D}}^\top\right) = \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \right\}, \tag{3.4}$$

which is the set of all joint probability distributions of the n-dimensional demand, whose support is nonnegative, where the sum of probabilities equal to 1, the expectation is $\mathbf{m}$, and the covariance is $\mathbf{\Sigma}$. The firm's best strategy against this adversary is to choose $\mathbf{y}$ that minimizes the "worst-case" expected cost (i.e., the maximum expected cost among distributions in $\mathcal{F}_{\geq 0}$). Mathematically, this is done by solving the following minmax robust problem

$$C^* := \min_{\mathbf{y} \geq 0} \sup_{f \in \mathcal{F}_{\geq 0}} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}}) \right]. \tag{3.5}$$

We denote the optimal value of (3.5) as C^* and its optimal solution as $\mathbf{y}^*$. If the firm chooses the initial inventory level $\mathbf{y}^*$, then it can be guaranteed that the expected cost is no larger than C^* under any joint demand distribution with mean-covariance $(\mathbf{m}, \mathbf{\Sigma})$. Since the inventory levels are chosen to be robust to any specific distribution, we also refer to (3.5) as the *distributionally robust* multi-location newsvendor problem with inventory risk pooling.

3.3.4 Discussion of the model

For the special case of $s_{ij} > h + p$ for all $i \neq j$, it is never optimal to allow demand spillover (i.e., $z_{ij} = 0$ for any $i \neq j$ and $z_{ii} = \min(d_i, y_i)$), so the cost reduces to n separable single-location newsvendor costs. Note that in this special case, while the newsvendor cost is separable by location, the minmax robust problem (3.5) is not, due to the joint constraints (3.4) on the joint probability distribution. Hanasusanto et al. (2015) proved that a minmax robust problem of n single-location newsvendor costs with joint mean and covariance information is NP-hard even in the absence of constraints on the support.

The minmax robust problem under a general fulfillment network (where $s_{ij} \leq h+p$) is intractable, due to the recourse network flow linear program. In general, there are two ways that previous studies in the literature deal with such issues. The first method is to assume a linear decision rule (Ben-Tal et al., 2004). Using this technique, Chou et al. (2006) develop a tractable formulation for the robust transshipment problem, where the decision variables are assumed to be linear functions of some underlying primitive uncertainties, and that these primitives have known support, and known forward and backward deviations. Linear decision rules usually have little basis in reality and are used solely to yield tractability, and hence there is no guarantee that they approximate the optimal second-stage flows well.

The second method (which this paper employs), is to approximate the cost structure of the network flow problem. Recently, Yan et al. (2018) studied the distributionally robust network flow problem under restricted 0-1 cost structures. This simplified cost structure gives rise to binary solutions in the dual program, which are then exploited using techniques from Natarajan et al. (2011). In our study, we extend this analysis to a more general class of fulfillment cost structures which preserve tractability. It is worthwhile to note that the simplest case of a 2-level nested fulfillment structure is similar to the problem analyzed by Yan et al. (2018).

The nested fulfillment structure is a good approximation whenever the geographical region inherently contains this hierarchical cluster structure. In particular, this is common in countries like the US that have dispersed population centers, where inter-cluster distances are much higher than intra-cluster distances. Indeed, errors in approximation arise when we ascribe a single fulfillment cost to fulfillment between any two locations in two different clusters. However the number of units being shipped in these higher levels is small when more demands are being pooled within each cluster, and as a result the error in approximation of the total fulfillment cost is small.[2]

The ubiquity of the nested structure can also be understood by noting its similarity to a tree metric, which is an approximation for a general metric on n nodes derived from an edge-weighted rooted tree with the n nodes as leaves. The tree metric has found applications in various problems that exhibit hierarchial characteristics, and the approximation of a general metric by a tree metric has been studied in the Computer Science literature (see Bartal, 1998; Fakcharoenphol et al., 2004).

In Section 3.6, we develop a simple algorithm that employs hierarchical agglomerative clustering to recover an n-level nested structure from a general distance-based

[2]Govindarajan et al. (2018, Proposition 5) showed that the error from assuming a constant fulfillment cost diminishes to zero in the asymptotic case where there are infinite number of locations while holding positive safety stock. This is because, as the number of inventory nodes in a given area increases, the chance that a unit of unfulfilled demand from one location is fulfilled from a close-by location is high. In our case, a similar intuition applies: the probability that fulfillment happens in higher levels is low, as there is enough supply to fulfill the pooled demand in lower levels. As a result, the error contribution to the expected total shipping cost from higher levels of fulfillment is low.

fulfillment cost structure, and empirically show that this nested structure tightly approximates the actual expected total fulfillment cost under a variety of distributions.

3.4 Closed-form bounds

As can be seen from the closed-form expression in Lemma 3.3.1, solving a distributionally robust newsvendor-type problem requires finding a tight bound on the expected unmet demand. In previous works on the classical newsvendor, tight closed-form bounds have been derived that rely on specific parameters of the demand distribution. The most well-known of these uses mean and variance (Scarf, 1958; Gallego and Moon, 1993), which we state below for completeness.

Lemma 3.4.1 (Scarf 1958; Gallego and Moon 1993) *If $\tilde{D}$ has mean m and standard deviation σ, then for any y,*

$$\mathbb{E}\left(\tilde{D} - y\right)^{+} \leq \frac{1}{2}\left(m - y + \sqrt{(m-y)^2 + \sigma^2}\right). \tag{3.6}$$

It has been shown that this bound is tight, in the sense that there exists a demand distribution where (3.6) holds with equality.

The challenge in deriving tight closed-form bounds in our setting is that there are multiple demands which interact, not only because they are correlated, but also from the nested terms resulting from inventory pooling. From Lemma 3.3.1, a naïve bound for the expected cost is to simply use the Scarf result to get the bound:

$$\eta_{\ell,k} \cdot \mathbb{E}\left(\mathbf{e}_{\ell,k}^{\top}\mathbf{D} - \mathbf{e}_{\ell,k}^{\top}\mathbf{y}\right)^{+} \leq \frac{\eta_{\ell,k}}{2}\left(\mathbf{e}_{\ell,k}^{\top}\mathbf{m} - \mathbf{e}_{\ell,k}^{\top}\mathbf{y} + \sqrt{\left(\mathbf{e}_{\ell,k}^{\top}\mathbf{m} - \mathbf{e}_{\ell,k}^{\top}\mathbf{y}\right)^2 + \mathbf{e}_{\ell,k}^{\top}\boldsymbol{\Sigma}\mathbf{e}_{\ell,k}}\right), \tag{3.7}$$

where $\mathbf{e}_{\ell,k}^{\top}$ is the k^{th} row of matrix $\mathbf{E}_{\ell}$. Combining these for all ℓ, k results in a bound for the expected total cost. While this bound uses covariance information, what it loses however is the information from the nested hierarchy of fulfillment. Hence, for a nested network with multiple locations, using the Scarf bound would result in a loose bound for the expected total cost.

We demonstrate this precisely for the special case of two identical regions, i.e., they have the same fulfillment cost parameters ($s_{12} = s_{21} = s > s_0$) and the demands $\tilde{d}_1, \tilde{d}_2$, though correlated, have the same mean and standard deviation. Suppose that m and σ are the mean and standard deviation, respectively, of both demands. We denote the correlation coefficient by ρ. We refer to the set of all joint distributions with these statistics as $\mathcal{F}^{m\sigma\rho}$.

The symmetry of the problem yields the following lemma.

Lemma 3.4.2 *If there are two identical regions, then there exists optimal inventory levels in each node that are also identical of the form* $\mathbf{y} = (y, y)^\intercal$ *for some* $y \geq 0$.

The proof is relegated to the Appendix, and relies on the convexity of the min-max robust problem in (3.5). This property aids us in simplifying the derivation, since we can restrict our analysis to only inventory decisions where the quantity is identical in both locations. Note that from the definition of $\mathcal{F}^{m\sigma\rho}$, we allow the demands to take all values in $\Re^2$, including negative values.[3] The later sections would assume a nonnegative support.

For the case with two identical locations, each with an initial inventory level y, we are able to derive a *tight* bound on the combined unmet demand, which we state as a proposition below.

Proposition 3.4.1 *(Nested bound) For any* $\zeta > 0$, y, *and* $f \in \mathcal{F}^{m\sigma\rho}$,

$$\mathbb{E}_f\left[(\tilde{d}_1 - y)^+ + (\tilde{d}_2 - y)^+ + \zeta\left(\tilde{d}_1 + \tilde{d}_2 - 2y\right)^+\right] \leq (\zeta + 1)\left(m - y + \sqrt{(m - y)^2 + \gamma\sigma^2}\right) \tag{3.8}$$

where $\gamma := \frac{\zeta+1+\zeta\rho}{2\zeta+1} \in (0, 1]$. *Moreover, if* $\gamma(\nu^2 + 1) \geq 2$, *where* $\nu := \frac{3\zeta+1}{\zeta+1}$, *then this bound is tight for some probability distribution* $f_y^* \in \mathcal{F}^{m\sigma\rho}$ *with six support points.*

The left-hand side is the cost of unmet demand in a network with two regions, where the inventory level in each region is y. If there is not enough inventory within a region, a penalty cost (normalized to 1) is incurred, with an additional penalty cost ζ when demand remains unmet after pooling.

We prove Proposition 3.4.1 in the appendix, however, we provide a discussion here. The supremum in $\mathcal{F}^{m\sigma\rho}$ of the left-hand side of (3.8) is a moment problem which we show to be equivalent to the dual problem:

$$\begin{aligned} \max_{t,u,r,v} \quad & t + 2mr + 2(m^2 + \sigma^2)u + (m^2 + \rho\sigma^2)v \\ \text{s.t.} \quad & g(d_1, d_2; t, u, r, v) \leq q(d_1, d_2) \quad \forall (d_1, d_2) \in \Re^2 \end{aligned} \tag{3.9}$$

[3]If we assume a nonnegative support for the distribution f, the derivation of the analytic expression becomes more complicated; to simplify the derivation, we assume that the support is $\Re^2$ for the two-location problem.

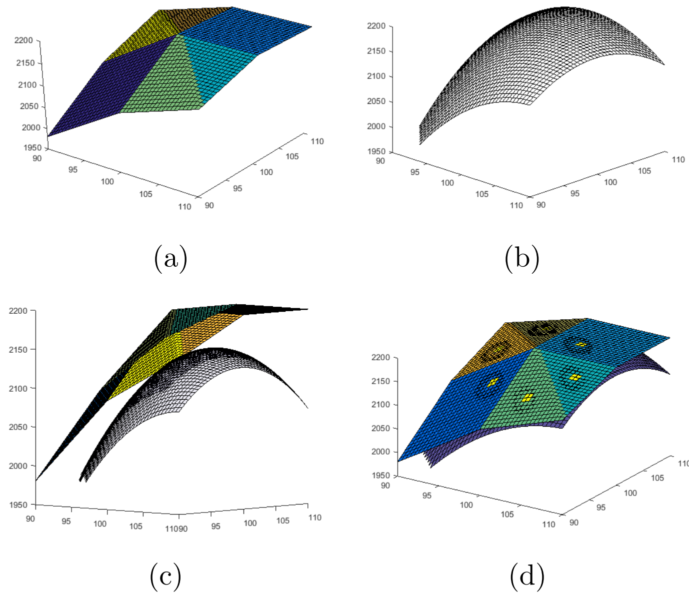

Figure 3.3: Illustration of the dual program. (a) The piecewise-planar function $q(d_1, d_2)$, (b) The quadratic function $g(d_1, d_2)$, (c) The functions corresponding to a dual feasible solution, (d) The functions corresponding to the dual optimal solution.

where $q(d_1, d_2)$ is a piecewise linear function in (d_1, d_2) with six pieces (shown in Figure 3.3a) and $g(d_1, d_2) := t + r(d_1 + d_2) + u(d_1^2 + d_2^2) + vd_1d_2$ is a quadratic function in (d_1, d_2) (shown in Figure 3.3b). The dual variables t, u, r, v are the parameters of the quadratic function.

Note that the dual program (3.9) has infinitely many constraints. A dual solution (t, u, r, v) is feasible if the quadratic function is bounded above by the piecewise-planar function in all of $\Re^2$ (Figure 3.3c shows these functions corresponding to a dual feasible solution). The optimal dual solution results in the two functions touching at exactly six points in $\Re^2$ (shown as the bright points in each face of the piecewise-planar function in Figure 3.3d). The points where the two functions touch are the support points of the demand distribution where the bound (3.8) holds with equality.

We note that the cap function $q(d_1, d_2)$ is a linear transformation of the function inside the expectation in (3.8). Therefore, the information about the hierarchical nesting is reflected in this cap function. Because of this additional structure, we can easily check that the bound in (3.8) can be much smaller than the corresponding Scarf bound. This is especially true for small values of y when demand spillover is more likely to occur (hence, the pooling term is more likely to be positive), as illustrated in Figure 3.4.

Therefore, this motivates the need to develop a new solution framework that uti-

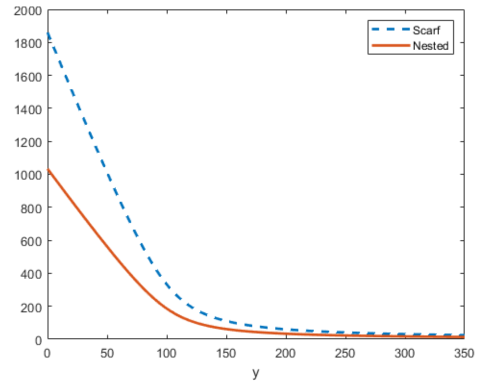

Figure 3.4: Scarf bound and nested bound (Proposition 3.4.1) for two identical locations with $\zeta = 4, m = 100, \sigma = 50, \rho = 0$.

lizes the cascading structure of the cost function (3.3.1) under a hierarchical nested fulfillment network. Unfortunately, the general problem is NP-hard, and only in the simplest case (two identical locations) do we have an exact solution due to the existence of a closed-form tight bound (though the proof is cumbersome). In Section 3.5, we propose such a solution framework that works on any hierarchical nested fulfillment network.

An implication of Proposition 3.4.1 is that the worst-case expected cost of inventory level $\mathbf{y} = (y, y)$: $\max_{f \in \mathcal{F}^{m\sigma\rho}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] = \bar{C}(y)$ is attained under a six-point distribution f_y^*.

Theorem 3.4.1 *If there are two identical regions, then for any* $\mathbf{y} = (y, y)$*,*

$$\sup_{f \in \mathcal{F}^{m\sigma\rho}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] \le \bar{C}(y) := 2s_0 m - (p - h - s_0)(y - m) + (p + h - s_0)\sqrt{(y - m)^2 + \gamma\sigma^2}, \tag{3.10}$$

where $\gamma := \frac{(p+h-s)(1+\rho)+s-s_0}{2(p+h)-s-s_0}$*. Moreover, if* $\gamma(\nu^2 + 1) \ge 2$*, where* $\nu := \frac{3(h+p-s_0)-2(s-s_0)}{h+p-s_0}$*, then* $\mathbb{E}_{f_y^*}[C(\mathbf{y}, \tilde{\mathbf{D}})] = \bar{C}(y)$ *for some probability distribution* $f_y^* \in \mathcal{F}^{m\sigma\rho}$ *with six support points.*

Proof. From Lemma 3.3.1, we know that

$$\begin{aligned} C(\mathbf{y}, \mathbf{D}) = {} & h(2y - d_1 - d_2) + s_0(d_1 + d_2) \\ & + \sum_{j=1,2} (s - s_0)(d_j - y)^+ + (p + h - s)(d_1 + d_2 - 2y)^+. \end{aligned}$$

Taking the expectation, the theorem immediately follows from Proposition 3.4.1 with

$\zeta = \frac{p+h-s}{s-s_0}$. ■

A corollary of Theorem 3.4.1 is an analytic expression for the robust optimal inventory level under the two-location nested fulfillment network. The proof is relegated to the appendix.

Proposition 3.4.2 *If there are two identical regions, and if $\gamma(\nu^2+1) \geq 2$ where $\gamma := \frac{(p+h-s)(1+\rho)+s-s_0}{2(p+h)-s-s_0}$ and $\nu := \frac{3(h+p-s_0)-2(s-s_0)}{h+p-s_0}$, then the inventory levels that minimize the maximal expected cost over distributions in $\mathcal{F}^{m\sigma\rho}$ is $\mathbf{y}^* = (y^*, y^*)$, where*

$$y^* = m + \left(\frac{p-h-s_0}{2} \sqrt{\frac{\gamma}{h(p-s_0)}} \right) \sigma. \tag{3.11}$$

The minmax expected cost is $C^* = 2s_0 m + 2\sigma\sqrt{\gamma h(p-s_0)}$.

If demand spillover is not allowed or if $s > p+h$, then inventory is decentralized with each location solving a classical single-location newsvendor problem. In this case, the minmax robust inventory level is y^*_{Sc} (where Sc denotes the Scarf solution), given by:

$$y^*_{Sc} = m + \left(\frac{p-h-s_0}{2\sqrt{h(p-s_0)}} \right) \sigma.$$

Since $\gamma \leq 1$, it directly follows that $y^* \leq y^*_{Sc}$ whenever $p - s_0 \geq h$, and that $y^* \geq y^*_{Sc}$ whenever $p - s_0 \leq h$. Since $p - s_0$ is the underage cost without pooling, then when $p - s_0 \geq h$, the decentralized solution y^*_{Sc} is large due to the high underage cost. On the other hand, in a centralized system, unmet demand can be fulfilled by inventory from any location, so the solution y^* is lower when $p - s_0 \geq h$. For a similar reason, when $p - s_0 \leq h$, the decentralized solution is low due to the high overage cost. In a centralized system, excess inventory in one location can be used elsewhere, so the solution y^* is higher when $p - s_0 \leq h$. Hence, the fact that $|y^* - m| \leq |y^*_{Sc} - m|$ is because of inventory risk sharing resulting from pooling, mirroring similar results from stochastic systems. Indeed, as $\rho \to 1$, we have $\gamma \to 1$, and as a result, y^* converges to the decentralized solution y^*_{Sc}.

3.4.1 The Effect of Cost and Demand Parameters on the Worst-case Distribution

We provide in the Appendix the expressions for the support points and the probabilities of the worst-case discrete distribution f^*_y (the subscript y is to emphasize that the

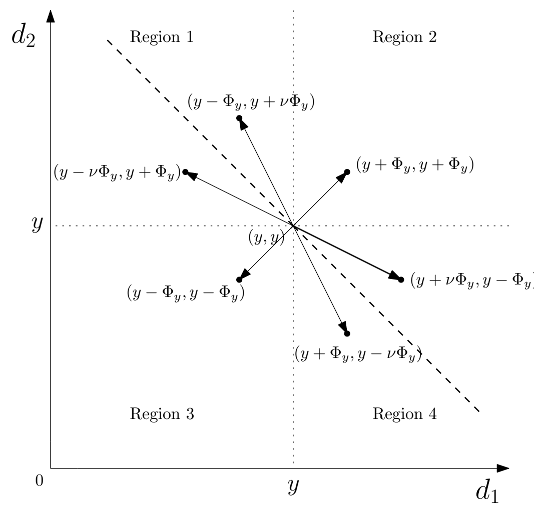

Figure 3.5: The six support points of f_y^*. Each support point lies on one of three solid lines passing through (y, y), where the distance from this point increases in proportion to $\Phi_y = \sqrt{(y-m)^2 + \gamma\sigma^2}$. The dashed line corresponds to a perfect balance between demand and supply after demand spillover.

worst-case distribution depends on the inventory level).

Figure 3.5 shows the support points of f_y^*, where $\Phi_y := \sqrt{(y-m)^2 + \gamma\sigma^2}$. From the figure, we observe that each support point lies on one of three solid lines that pass through point (y, y). When we either increase $|y - m|$, increase σ, or increase ρ, the distance of the support points to (y, y) increases proportionally in the direction indicated by the arrows in Figure 3.5.

The dashed line in the figure ($d_1+d_2 = 2y$) corresponds to a perfect balance between demand and inventory after demand spillover. Two of the solid lines converge to this dashed line as $\nu \to 1$, which occurs when $h + p$ decreases to the limit s. Thus, when the overage or the underage cost is high (or the fulfillment cost s is low), the solid lines pivot further away from the dashed line, and hence the support points of the worst-case distribution result in very large excess inventory or unmet demand after pooling.

We call a system imbalanced if, after in-location fulfillment, there is leftover inventory in one location and unfulfilled demand at the other location. In Figure 3.5, we divide the demand region into four quadrants (Regions 1 through 4 demarcated by the two dotted lines), and observe that when the demand realizations are in Regions 1 and 4, there is imbalance in the system. In these regions, we measure the magnitude of imbalance as the sum of the leftover inventory and unfulfilled demand. Mathematically, this is the L^1 distance between the support point and (y, y). We find that the magnitude of imbalance induced by the support points is increasing in $|y - m|$, σ and ρ.

The probability of imbalance (sum of probabilities of realizations in Regions 1 and 4), however, is no more than $\gamma\sigma^2/\Phi_y^2$. When the decision maker chooses inventory levels with high safety stock (as is typically done in practice), the worst case distribution causes imbalance across locations with low probability, but of large magnitude. Thus retailers should strive to eliminate such low-probability extreme situations, which can be done by adopting various demand-shaping strategies that prevent imbalances in the customer locations, such as strategic location-specific product display, recommendations and flash promotions.

3.5 Heuristic for general nested fulfillment networks

Our objective in this section is to develop a computationally tractable solution method that is able to find an upper bound on

$$\mathbb{E}_f\left[\sum_{\ell=0}^{L-1}\eta_\ell^\top\left(\mathbf{E}_\ell\tilde{\mathbf{D}}-\mathbf{E}_\ell\mathbf{y}\right)^+\right], \tag{3.12}$$

requiring only the mean and the covariance of $\tilde{\mathbf{D}}$, where $\Xi = \{\mathbf{E}_\ell\}_{\ell=0}^{L-1}$ defines the structure of a nested fulfillment network. Importantly, the solution method should preserve the cascading nature of the hierarchical nested fulfillment network, resulting in a tighter bound than by simply using Scarf-type bounds (3.7). This provides a computationally tractable heuristic that results in a good approximation to the robust newsvendor problem under a general nested fulfillment network

Before developing our solution method, we first introduce the following lemma.

Lemma 3.5.1 *If $\mathcal{F}$ is the set of all joint probability distributions with mean vector $\mathbf{m}$ and covariance matrix $\mathbf{\Sigma}\succ 0$, then for any $\mathbf{y}\in\Re^n$,*

$$\begin{aligned}\sup_{f\in\mathcal{F}}\mathbb{E}_f\left[\sum_{\ell=0}^{L-1}\eta_\ell^\top\left(\mathbf{E}_\ell\tilde{\mathbf{D}}-\mathbf{E}_\ell\mathbf{y}\right)^+\right] = \underset{t,\mathbf{r},\mathbf{Y}}{\text{minimize}}\quad & t+\mathbf{r}^\top\mathbf{m}+\langle\mathbf{Y},\mathbf{\Sigma}+\mathbf{m}\mathbf{m}^\top\rangle\\ \text{subject to:}\quad & \begin{pmatrix}\mathbf{Y} & \frac{1}{2}(\mathbf{r}-\mathbf{a})\\ \frac{1}{2}(\mathbf{r}-\mathbf{a})^\top & t+\mathbf{a}^\top\mathbf{y}\end{pmatrix}\succeq 0,\quad \forall\mathbf{a}\in\mathcal{L},\end{aligned} \tag{3.13}$$

where $\mathcal{L} := \Big\{\mathbf{a}\mid\mathbf{a}=\sum_{\ell=0}^{L-1}\mathbf{E}_\ell^\top\left(\eta_\ell\odot\mathbf{e}_{A_\ell}\right)$ *for some* $(A_0,A_1,\cdots,A_{L-1})\in 2^{[n_0]}\times 2^{[n_1]}\times$

$\cdots \times 2^{[n_{L-1}]}\Big\}$ *and* $\mathbf{e}_{A_\ell}$ *is an* n_ℓ*-dimensional binary vector where* $e_{A_\ell,k} = 1$ *if and only if* $k \in A_\ell$.

The proof is relegated to the Appendix and relies on strong duality in moment problems.

Note that because of the form of $C(\mathbf{y}, \mathbf{D})$ under an L-level nested fulfillment network (Lemma 3.3.1), the implication from Lemma 3.5.1 is that an exact solution $\mathbf{y}^*$ to the minmax robust problem can be found through solving a semidefinite program (SDP) with 2^N semidefinite constraints, where $N := \sum_{\ell=0}^{L-1} n_\ell$ is the total number of nodes in the tree representation of the nested fulfillment structure (similar to Figure 3.2). This is done by defining $\mathbf{y}$ as a decision variable in the SDP.

Semi-definite programs, much like linear programs, can be solved through interior point methods which have polynomial time worst-case complexity (Vandenberghe and Boyd, 1996). However, the SDP in (3.13) is not computationally tractable beyond small values for n, since it involves $\mathcal{O}(2^N)$ constraints that each require a $(n+1) \times (n+1)$ matrix to be positive semidefinite. In the worst case, $N = \frac{n(n+1)}{2}$ if the number of levels is $L = n$ and $n_\ell = n - \ell$. This motivates the need for tractable approximations to the tight bound for larger values of n.

We begin our discussion by reformulating the expectation of the nested function (3.12) as follows:

$$\mathbb{E}_f \left[\max_{\substack{\mathbf{x}^{(0)} \in \{0,1\}^{n_0}, \\ \mathbf{x}^{(1)} \in \{0,1\}^{n_1}, \cdots, \\ \mathbf{x}^{(L-1)} \in \{0,1\}^{n_{L-1}}}} \sum_{\ell=0}^{L-1} \left(\mathbf{x}^{(\ell)} \odot \eta_\ell\right)^\top \left(\mathbf{E}_\ell \tilde{\mathbf{D}} - \mathbf{E}_\ell \mathbf{y}\right) \right] \tag{3.14}$$

where $\mathbf{x}^{(\ell)}$ is an n_ℓ-dimensional binary vector, for $\ell = 0, 1, \ldots, L-1$, and $\odot$ is the element-wise product operator.

We observe that formulation (3.14) preserves the hierarchical structure of the network since the implied optimization problem inside the square brackets finds a solution $\{\mathbf{x}^{(\ell)}(\mathbf{D})\}_{\ell=0}^{L-1}$ that obeys the nested structure for a particular demand realization $\mathbf{D}$. To see this, note that since $\eta_{\ell,k} \geq 0$ for any ℓ and $k \in n_\ell$, we have that $x_k^{(\ell)}(\mathbf{D}) = 0$ if and only if $\mathbf{e}_{\ell,k}^\top \tilde{\mathbf{D}} < \mathbf{e}_{\ell,k}^\top \mathbf{y}$. In words, $x_k^{(\ell)}(\mathbf{D}) = 0$ will be chosen only if there is inventory remaining after level ℓ fulfillment of regions in $\mathcal{I}_k^{(\ell)}$. Note that this can only occur if there exists at least one path in the tree from level 0 to level ℓ (for example, the path through $\{\mathcal{I}_{k_0}^{(0)}, \mathcal{I}_{k_1}^{(1)}, \ldots, \mathcal{I}_{k_\ell}^{(\ell)}\}$) where there is excess inventory in all sets in this path.

This makes sense because level ℓ fulfillment will not occur unless there was excess inventory after all the less costly lower level fulfillments. Mathematically, if $x_k^{(\ell)}(\mathbf{D}) = 0$, then there must exist at least one path where $x_{k_0}^{(0)}(\mathbf{D}) = x_{k_1}^{(1)}(\mathbf{D}) = \cdots = x_{k_{\ell-1}}^{(\ell-1)}(\mathbf{D}) = 0$. The inherent interdependence among the values in $\{\mathbf{x}^{(\ell)}(\mathbf{D})\}_{\ell=0}^{L-1}$ is due to the nested structure of the fulfillment network.

Since the value of the maximizer depends on the specific realization of random demands, the maximizers are random variables, which we denote as $\{\tilde{\mathbf{x}}^{(\ell)}\}_{\ell=0}^{L-1}$. Hence, (3.14) is equivalent to

$$\mathbb{E}_f \left(\sum_{\ell=0}^{L-1} \left(\tilde{\mathbf{x}}^{(\ell)} \odot \eta_\ell \right)^\top \left(\mathbf{E}_\ell \tilde{\mathbf{D}} - \mathbf{E}_\ell \mathbf{y} \right) \right), \tag{3.15}$$

where we use tilde on the binary variables to emphasize that they are stochastic variables.

Note that reformulation (3.15) has cross products of random variables, where the underlying uncertainty has a joint distribution f with nonnegative support and with mean $\mathbf{m}$ and covariance $\mathbf{\Sigma}$. A method to relax a bilinear function into a linear formulation is to introduce new variables, which *lifts* the problem to a higher dimensional space (see for instance Sherali and Alameddine 1992). This method was used in developing heuristics for the multidimensional robust newsvendor problems in Natarajan and Teo (2017) and Natarajan et al. (2017), where such cross products of random variables occur, though these papers study simpler networks without demand spillover.

Consider the N-dimensional random vector $\tilde{\mathbf{x}} := \begin{pmatrix} \tilde{\mathbf{x}}^{(L-1)\top} & \tilde{\mathbf{x}}^{(L-2)\top} & \cdots & \tilde{\mathbf{x}}^{(0)\top} \end{pmatrix}^\top$. We linearize the function (3.15) by lifting it to a higher dimensional space by introducing the following new variables:

$$\mathbf{x} := \mathbb{E}_f \left(\tilde{\mathbf{x}} \right) \in \Re^N, \tag{3.16}$$

$$\mathbf{Q} := \mathbb{E}_f \left(\tilde{\mathbf{x}} \tilde{\mathbf{D}}^\top \right) \in \Re^{N \times n}, \tag{3.17}$$

$$\mathbf{R} := \mathbb{E}_f \left(\tilde{\mathbf{x}} \tilde{\mathbf{x}}^\top \right) \in \Re^{N \times N}. \tag{3.18}$$

Defining $N_{L-1} := 0$ and $N_\ell := \sum_{m=0}^{L-\ell-2} n_{L-1-m}$ for $\ell = 0, 1, \ldots, L-2$, we now have that (3.15) is equivalent to the linear function

$$\sum_{\ell=0}^{L-1} \sum_{k \in [n_\ell]} \eta_{\ell,k} \cdot \left(\sum_{i \in \mathcal{I}_k^{(\ell)}} Q_{N_\ell + k, i} - \sum_{i \in \mathcal{I}_k^{(\ell)}} x_{N_\ell + k} \cdot y_i \right). \tag{3.19}$$

Therefore, an upper bound to (3.12) can be found by taking the maximum of the linear function (3.19) with respect to the variables $(\mathbf{x}, \mathbf{Q}, \mathbf{R})$ that satisfy some feasibility constraints which are consistent with their definition in (3.16)–(3.18). First, since the joint distribution f has mean $\mathbf{m}$ and covariance $\mathbf{\Sigma}$, we have the constraint that the symmetric matrix

$$\mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix}\begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix}^\top\right) = \begin{pmatrix} 1 & \mathbf{m}^\top & \mathbf{x}^\top \\ \mathbf{m} & \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top & \mathbf{Q}^\top \\ \mathbf{x} & \mathbf{Q} & \mathbf{R} \end{pmatrix},$$

is positive semidefinite. Second, since the support of f is nonnegative, then it must follow that $\mathbf{Q} \geq \mathbf{0}$. A third necessary condition follows from the fact that $\tilde{\mathbf{x}} \in \{0,1\}^N$, implying that

$$\begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{R} \end{pmatrix} := \mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix}\begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix}^\top\right) \in \operatorname{conv}\left\{\begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix}\begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix}^\top : \mathbf{w} \in \{0,1\}^N\right\}. \tag{3.20}$$

That is, the left-hand side matrix is a convex combination of a set of Boolean matrices where each matrix is itself a product of Boolean variables. The convex hull of such matrices is often refered to as the Boolean quadric polytope (Padberg, 1989).

Computing the convex hull is a difficult problem since unconstrained binary quadratic programming is NP-hard in general. We instead use a simple linear relaxation of this polytope. Note that if the matrix on the left-hand side of (3.20) belongs to the Boolean quadric polytope, then $\mathbf{R} \geq \mathbf{0}$ and, for all $i, j \in [N]$, $R_{ii} = x_i$, $R_{ij} \leq x_i$, and $R_{ij} \geq x_i + x_j - 1$.

Therefore, we have the following proposition.

Proposition 3.5.1 *If f is a joint probability distribution with nonnegative support and* *with mean vector* $\mathbf{m}$ *and covariance matrix* $\boldsymbol{\Sigma} \succ 0$*, then for any* $\mathbf{y} \in \Re^n$*,*

$$
\begin{aligned}
\mathbb{E}_f\left[\sum_{\ell=0}^{L-1} \eta_\ell^\top \left(\mathbf{E}_\ell \tilde{\mathbf{D}} - \mathbf{E}_\ell \mathbf{y}\right)^+\right] \;\leq\; \max_{\mathbf{x},\mathbf{Q},\mathbf{R}} & \sum_{\ell=0}^{L-1} \sum_{k\in[n_\ell]} \eta_{\ell,k} \cdot \left(\sum_{i\in\mathcal{I}_k^{(\ell)}} Q_{N_\ell+k,i} - \sum_{i\in\mathcal{I}_k^{(\ell)}} x_{N_\ell+k} \cdot y_i\right) \\
\text{s.t.} & \begin{pmatrix} 1 & \mathbf{m}^\top & \mathbf{x}^\top \\ \mathbf{m} & \boldsymbol{\Sigma} + \mathbf{m}\mathbf{m}^\top & \mathbf{Q}^\top \\ \mathbf{x} & \mathbf{Q} & \mathbf{R} \end{pmatrix} \succeq 0 \\
& R_{ii} = x_i \qquad i \in [N] \\
& R_{ij} \leq x_i \qquad i \in [N], j \in [N] \\
& R_{ij} \geq x_i + x_j - 1 \qquad i \in [N], j \in [N] \\
& \mathbf{Q}, \mathbf{R} \geq 0
\end{aligned} \tag{3.21}
$$

The proposition states that the expected unmet demand can be bounded by the optimal value of a semidefinite program with a *single* semidefinite constraint of size $(N + n + 1) \times (N + n + 1)$. Contrast this with the tight bound in Lemma 3.5.1 from an SDP with $\mathcal{O}(2^N)$ semidefinite constraints of size $(n + 1) \times (n + 1)$.

The dual of the SDP (3.21) is a minimization problem. Therefore, from Lemma 3.3.1, an upper bound to the minmax cost C^* defined in (3.5), is

$$
\begin{aligned}
C^{\mathrm{H}} := \min_{\substack{\mathbf{y},t_0,\mathbf{t},\mathbf{Y},\mathbf{u},\\ \mathbf{B},\mathbf{W},\mathbf{U},\mathbf{V}}} \; & h \cdot \mathbf{e}^\top(\mathbf{y} - \mathbf{m}) + \mathbf{s}_0^\top \mathbf{m} + t_0 + \mathbf{t}^\top \mathbf{m} + \langle \mathbf{Y}, \boldsymbol{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle + \mathbf{e}^\top \mathbf{B}\mathbf{e} \\
\text{s.t.} \; & \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top \\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top \\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U} \end{pmatrix} \succeq 0 \\
& \mathbf{u} = -\mathbf{W}\mathbf{e} + \left(\mathbf{B} + \mathbf{B}^\top\right)\mathbf{e} + \mathbf{P}\mathbf{y} \\
& \mathbf{V} \geq \mathbf{P} \\
& \mathbf{U} \leq \mathbf{W} - \mathbf{B} \\
& \mathbf{W}, \mathbf{B} \geq 0 \\
& t_0 \in \Re,\; \mathbf{t} \in \Re^n,\; \mathbf{u} \in \Re^N, \\
& \mathbf{Y} \in \Re^{n\times n},\; \mathbf{B}, \mathbf{W}, \mathbf{U} \in \Re^{N\times N},\; \mathbf{V} \in \Re^{N\times n}.
\end{aligned} \tag{3.22}
$$

where $\mathbf{P} := \left(\mathbf{E}_{L-1}^\top \mathrm{diag}(\eta_{L-1}) \quad \mathbf{E}_{L-2}^\top \mathrm{diag}(\eta_{L-2}) \quad \cdots \quad \mathbf{E}_0^\top \mathrm{diag}(\eta_0)\right)^\top \in \Re^{N\times n}$. Thus, a computationally tractable heuristic for finding robust inventory levels under a nested

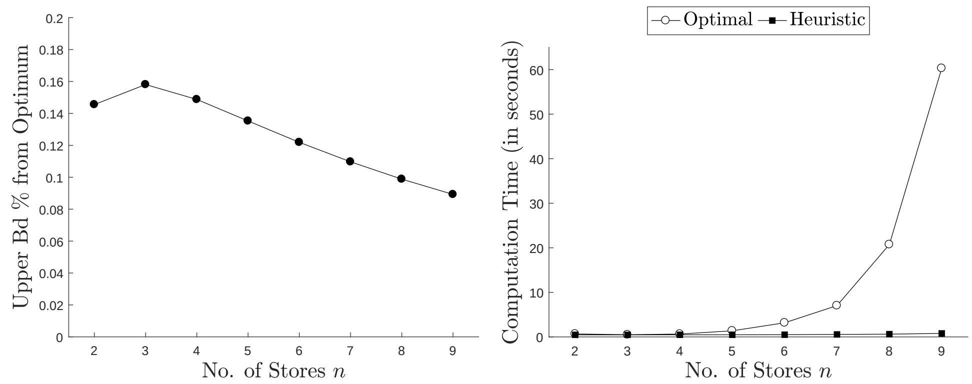

(a) (b)

Figure 3.6: Experiments with the optimal minmax cost and the optimized upper bound cost. (a) The gap between the two optimized costs is small, revealing that optimizing the upper bound is a good heuristic for the minmax robust problem. (b) The computational time of the heuristic is significantly smaller than the optimal SDP.

fulfillment network is to choose the inventory levels based on the optimal solution $\mathbf{y}$ to the SDP (3.22), which we denote as $\mathbf{y}^{\mathrm{H}}$ (recall that the true robust inventory levels are $\mathbf{y}^*$).

Whether $\mathbf{y}^{\mathrm{H}}$ is a close approximation to $\mathbf{y}^*$ depends on whether the SDP (3.21) can tightly approximate (3.13). We demonstrate this in computational experiments (since there is no support restriction in Lemma 3.5.1, we allow negative values for $\mathbf{q}$ and $\mathbf{Q}$, so the third constraint in (3.22) is an equality constraint). Our experiment is on a 2-level structure: in level 0, in-location fulfillment has a per-unit cost s_0, and in level 1, inventory from any location in the network can be used to meet unfulfilled demand at a per-unit cost s.

Figure 3.6 shows the results of the experiment as the number of locations is varied, with $h = 1$, $p = 100$, $s = 1$, $s_0 = 0$, with identical marginal distributions (mean $m = 100$ and standard deviation $\sigma = 50$), and with each pair of locations having a correlation coefficient of $\rho = 0.25$. Figure 3.6a shows the gap between the minmax cost C^* and the optimal value C^{H} of the SDP (3.22). The plot shows that C^{H} is close to the minmax cost (within 0.2%) and, in general, this gap decreases with the number of locations. Therefore, the upper bound (3.21) is empirically tight in the neighborhood of $\mathbf{y}^{\mathrm{H}}$ and $\mathbf{y}^*$. This reveals that our proposed solution method provides a good heuristic for approximating the robust inventory levels in a nested fulfillment network.

Figure 3.6b shows the computational tractability of this heuristic compared to solving for the robust inventory levels through (3.13). We observe that the computational time of the heuristic is significantly smaller compared to the optimal SDP, as demonstrated in the figure.[4] As a result, our proposed heuristic solves a general class of problems with nested fulfillment cost structures, with significantly less computational burden than the optimal SDP.

3.6 Approximating distance-based fulfillment costs with a nested structure

We now describe how to approximate a general distance-based fulfillment cost structure by a nested fulfillment structure. Suppose that the fulfillment cost is a function of the distance between two regions. In particular, for any two regions i, j, suppose $s_{ij} = \phi(r_{ij})$, where r_{ij} is the distance between the regions and ϕ is an increasing function. We denote by $\mathcal{R} = (r_{ij})_{ij}$ the distance matrix and $\mathbf{S} = (s_{ij})_{ij}$ the fulfillment cost matrix.

A decomposition of a general cost structure with n locations into an n-level nested fulfillment structure is done by hierarchical agglomerative clustering, which has been extensively studied in literature, dating back to Johnson (1967). We outline the procedure in Algorithm 3.

Algorithm 3 Hierarchical Agglomerative Clustering Algorithm

1: Let $\mathcal{S} = \{1, 2, \ldots, n\}$. Set $\bar{\mathcal{R}} = (\bar{r}_{ij})_{i,j\in\mathcal{S}} = \mathcal{R}$.
2: **while** $|\mathcal{S}| > 1$ **do**
3: Choose the two closest nodes $i^*, j^* = \text{argmin}_{i,j\in\mathcal{S}}\, \bar{r}_{ij}$.
4: Cluster i^*, j^* into a single node: $\mathcal{S} \leftarrow \mathcal{S} + \{i^*, j^*\} - \{i^*\} - \{j^*\}$
5: Recalculate distance matrix $\bar{\mathcal{R}} = (\hat{R}_{ij})_{i,j\in\mathcal{S}}$

The algorithm proceeds by progressively clustering the two closest nodes into a single node starting from the n leaf nodes corresponding to locations, until there remains only one cluster node which encompasses all the locations. In each step of the algorithm, the number of nodes is reduced by 1. Note that in order to choose the two closest nodes in each step, we need a notion of distance between clusters of nodes. A variety of measures can be considered to define the distance between two clusters, namely the minimum or maximum or average distance between the nodes in the two clusters, Ward's method, distance between the center of masses of the two clusters, etc.

[4] Even for $n = 9$, the optimal robust solution required around 20,000 unique SDP variables involved in the SDP constraints, whereas the heuristic required only 20 unique SDP variables, and the disparity is clearly seen in the computational times. The heuristic could solve up to $n = 100$ in under an hour, whereas the optimal solution could not be evaluated even for $n = 10$ due to memory constraints.

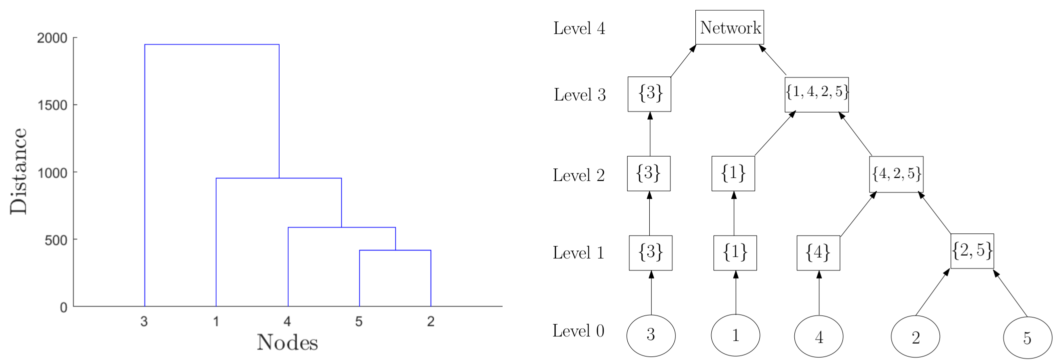

(a) (b)

Figure 3.7: Approximating a distance-based fulfillment cost structure with a nested fulfillment structure. Panel (a) is the dendrogram obtained by hierarchical clustering based on the distance matrix $\mathcal{R}$. Panel (b) is the corresponding 5-level nested fulfillment structure.

We demonstrate the algorithm using an example, where define the distance between two clusters as the average distance measure, namely the Unweighted Pair Group Method with Arithmetic Mean (UPGMA). The UPGMA distance between clusters $\mathcal{I}_1, \mathcal{I}_2$ is $\hat{r}_{\mathcal{I}_1,\mathcal{I}_2} = \frac{1}{|\mathcal{I}_1|} \cdot \frac{1}{|\mathcal{I}_2|} \cdot \sum_{i \in \mathcal{I}_1, j \in \mathcal{I}_2} r_{ij}$. In other words, it is the average distance between any two pairs of locations in $\mathcal{I}_1$ and $\mathcal{I}_2$.

Consider the following distance matrix among 5 nodes, and the corresponding fulfillment cost matrix constructed by the equation $s_{ij} = 10 + 0.005 \cdot r_{ij}$:

$$\mathcal{R} = \begin{bmatrix} 0 & 1,220 & 1,411 & 770 & 872 \\ 1,220 & 0 & 2,404 & 624 & 420 \\ 1,411 & 2,404 & 0 & 1,785 & 2,187 \\ 770 & 624 & 1,785 & 0 & 557 \\ 872 & 420 & 2,187 & 557 & 0 \end{bmatrix}, \quad \mathbf{S} = \begin{bmatrix} 10.0 & 16.1 & 17.1 & 13.8 & 14.4 \\ 16.1 & 10.0 & 22.0 & 13.1 & 12.1 \\ 17.1 & 22.0 & 10.0 & 18.9 & 20.9 \\ 13.8 & 13.1 & 18.9 & 10.0 & 12.8 \\ 14.4 & 12.1 & 20.9 & 12.8 & 10.0 \end{bmatrix}$$

Applying Algorithm 3 with UPGMA as the distance metric, we obtain the dendrogram shown in Figure 3.7a. The dendogram depicts the clustering in each step (level). $L = 5$ levels are shown, where level $\ell = 0$ consists of $\{3, 1, 4, 5, 2\}$, $\ell = 1$ consists of $\{3, 1, 4, \{5, 2\}\}$ and so on, with $\ell = n - 1$ corresponding to the cluster of all nodes. The dendogram also depicts the distance between the entities being clustered: in level 2 for instance, $\{4\}$ is clustered with $\{5, 2\}$ at the UPGMA distance $\hat{r}_{4,\{5,2\}} = 590.6$. This means that any fulfillment between nodes 4 and 5, or between nodes 4 and 2 takes place at a cost $\hat{s}_{4,\{5,2\}} = 10 + 0.005 \cdot \hat{r}_{4,\{5,2\}}$. Thus we have an approximation of the

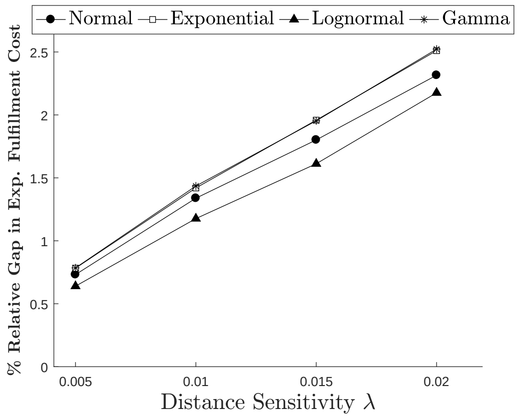

(a)

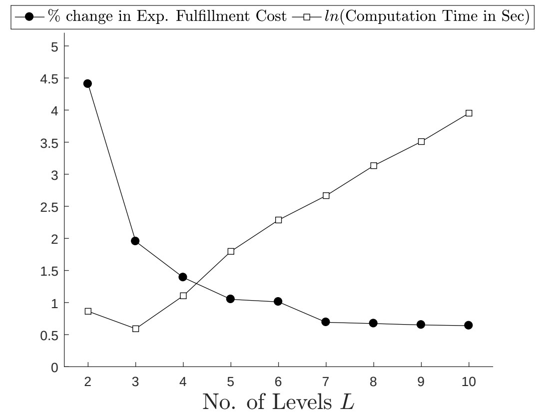

(b)

Figure 3.8: The approximation of the expected total fulfillment cost by a nested fulfillment structure. Panel (a) shows the relative gap in expected total fulfillment cost under a variety of distributions as fulfillment costs become more sensitive to distance. Panel (b) shows the quality of approximation and computational time as the number of levels in the nested structure is varied.

distance and fulfillment cost matrices:

$$\hat{\mathcal{R}} = \begin{bmatrix} 0 & 954 & 1,947 & 954 & 954 \\ 954 & 0 & 1,947 & 591 & 420 \\ 1,947 & 1,947 & 0 & 1,947 & 1,947 \\ 954 & 591 & 1,947 & 0 & 591 \\ 954 & 420 & 1,947 & 591 & 0 \end{bmatrix}, \quad \hat{\mathbf{S}} = \begin{bmatrix} 10.0 & 14.8 & 19.7 & 14.8 & 14.8 \\ 14.8 & 10.0 & 19.7 & 13.0 & 12.1 \\ 19.7 & 19.7 & 10.0 & 19.7 & 19.7 \\ 14.8 & 13.0 & 19.7 & 10.0 & 13.0 \\ 14.8 & 12.1 & 19.7 & 13.0 & 10.0 \end{bmatrix}$$

Algorithm 3 always results in a nested fulfillment structure with $L = n$ levels (see Figure 3.7b). We can also generate a nested fulfillment structure for any general L. We do this by cutting the dendrogram horizontally at $L - 2$ places across the y-axis: the ℓ^{th} line from the bottom gives rise to a partition of the set of locations by cutting through links that cluster these partitioned sets further. This partition gives us the clusters at level ℓ. We provide a detailed example in Appendix B.3.

3.6.1 Quality of nested approximations

We next study the nested network approximation in realistic fulfillment networks inspired from U.S. data. In these experiments, $n = 10$. (Further details of the experimental setup can be found in Section 3.7.2.) In particular, we calculate the relative gap in expected total fulfillment costs under the nested structure and under the actual fulfillment structure where $s_{ij} = s_0 + \lambda \cdot r_{ij}$. Given the inventory levels, if the gap is found to be small, we can confidently use the proposed heuristic to approximate the robut inventory levels for general networks with distance-based fulfillment costs.

Indeed, we see that the nested cost structure with $L = n$ tightly approximates the expected total fulfillment cost for a variety of distributions such as Normal, exponential, lognormal and gamma (Figure 3.8a). The approximation is detriorating as fulfillment costs become more sensitive to distance, however, even for high values of the distance sensitivity factor λ, the relative gap in expected total fulfillment costs is less than 3%. (We note that the distance sensitivity factors for the UPS Ground shipping is 0.0005 as estimated by Jasin and Sinha (2015), which is ten times lower than the lowest value of distance sensitivity considered in this experiment.) As a result, the nested fulfillment structure can serve as a good approximation for most distance-based shipping alternatives seen in practice.

Larger number of levels L best approximates general fulfillment networks, however it is at the expense of increased computational effort in solving the SDP (3.22). We next test how the approximation and the computational time are affected by the number of levels. The results under an exponential distribution are shown in Figure 3.8b. As expected, increasing L improves approximation quality, albeit marginally for $L > 7$, however, the computational time increases exponentially. Hence nominal values of L can achieve good approximations in relatively shorter time.

3.7 Numerical Analysis

We conduct multiple experiments on the proposed heuristic solutions. First, we compare the heuristic solution to stochastic optimal solutions in the 2-level structure (constant fulfillment cost for spillover demand) for various distributions to understand the expected value of additional information (EVAI) in a pooling context, which is defined as the loss incurred due to incomplete information about the distribution, as the heuristic solution only uses mean and covariance information. We then conduct experiments on simulated data to illustrate the superiority of the robust heuristic solution compared to stochastic solutions.

3.7.1 Experiments with a 2-level nested network

We begin by studying the performance of our proposed heuristic in a 2-level nested network where fulfillment costs for spillover demands are constant ($s_{ij} = s > s_0$ for all $i \neq j$). The simple 2-level structure is useful because we can isolate the effect of

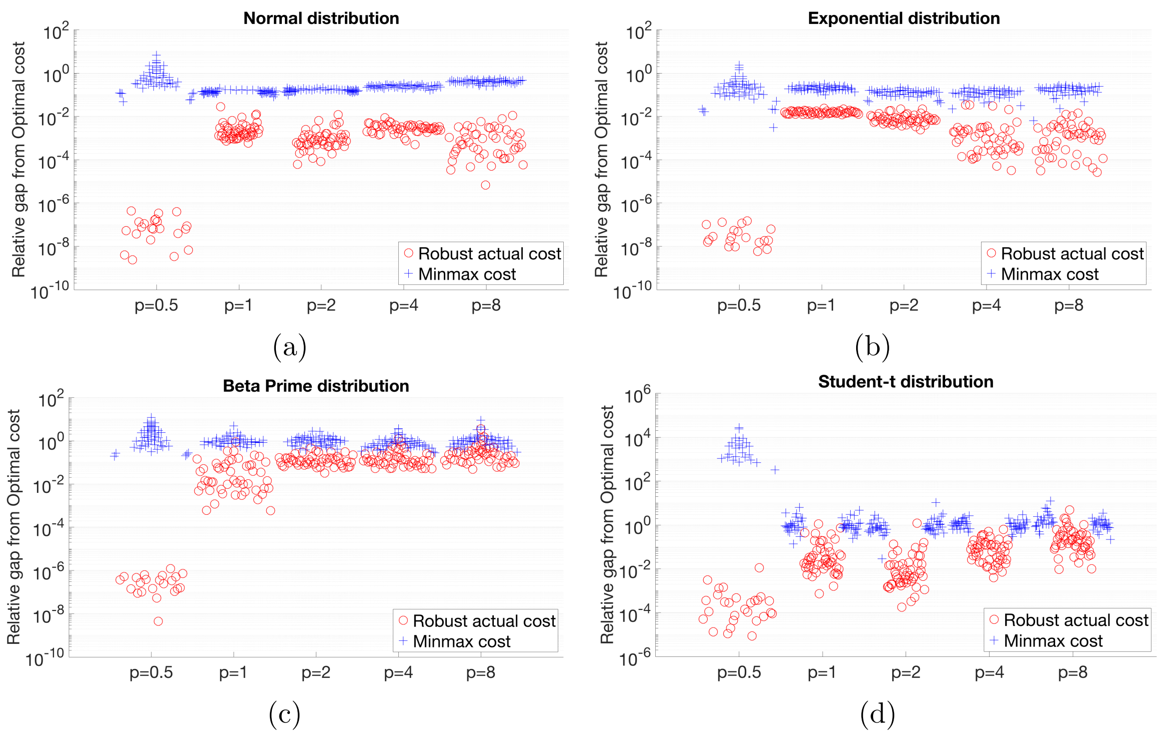

Figure 3.9: Relative gaps $(C^{\mathrm{H}}_f - C^*_f)/C^*_f$ (robust actual cost) and $(C^{\mathrm{H}} - C^*)/C^*$ (minmax cost). The x-axis spread of the data around each value of p is solely for visual clarity.

pooling from the network structure.

We randomly generate distribution parameters for the following parametric families: Normal, exponential, beta prime, and Student-t (the details for the parameter generation are given in Appendix B.4). Given a specific joint demand distribution f, we estimate the optimal expected cost $C^*_f := \min_{\mathbf{y}\geq 0} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$ using sample average approximation with 10^4 samples of the demand vector. Given the mean and the covariance of the random demand, we use our heuristic to approximate the robust inventory levels with $\mathbf{y}^{\mathrm{H}}$. We then compute the expected cost of the heuristic solution under the known true distribution f (which we denote as C^{H}_f).

Figure 3.9 illustrates the gap $C^{\mathrm{H}}_f - C^*_f$ (the $\circ$ markers), which represents the performance of the heuristic under a specific distribution f. In these experiments, $n = 5$, $h = 1$, $s = 1$, $s_0 = 0.5$, and we vary $p \in \{0.5, 1, 2, 4, 8\}$. We observe that the performance of the heuristic depends on p, seen from the small optimality gap for small values of p. If the distribution is either Normal or exponential, the heuristic has an actual expected cost that is close to optimal even for high values of p, with relative gaps in the order of .1% or 1%. For the beta prime and the Student-t distribution families, we observe that the relative gap can be as high as the order of 10%.

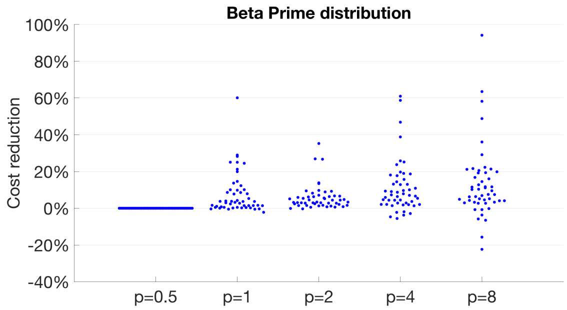

(a)

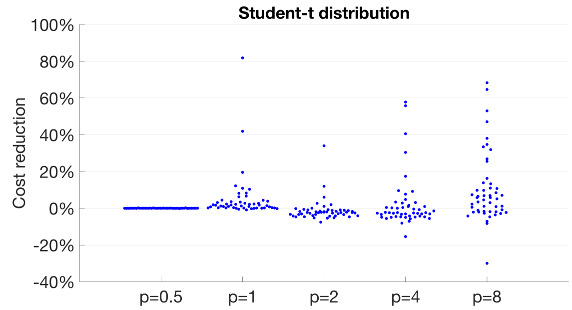

(b)

Figure 3.10: Reduction in expected cost (under the true distribution) of the robust inventory levels with partitioned statistics information.

The figure also shows the gap $C^{\mathrm{H}} - C_f^*$ (the + markers). Since under beta prime or Student-t distributions, the circle markers are close to the plus markers, we can infer that the expected cost under these distributions is close to the worst-case expected cost in the neighborhood of $\mathbf{y}^{\mathrm{H}}$. (Note that the plus markers are always above the circle markers since $C_f^{\mathrm{H}} \leq C^{\mathrm{H}}$.) We next discuss how the performance of a distributionally robust heuristic can be improved for these cases.

Since there are multiple joint ditributions in $\mathcal{F}$, then the range of possible values of $\mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})]$ for a given $\mathbf{y}$ could potentially be wide. This ambiguity may result in the true optimal solution to be different from the robust solution under some distributions (e.g. under beta prime or Student-t). A way to reduce ambiguity is by further restricting the distribution set, which can be accomplished by adding more information to $\mathcal{F}$. This can be done with *partitioned statistics* information, specifically, the mean and covariance of random vector $(\tilde{\mathbf{D}}^+, \tilde{\mathbf{D}}^-)$ whose i^{th} elements are $(\tilde{d}_i - m_i)^+$ and $(m_i - \tilde{d}_i)^+$, respectively. Partitioned statistics measures asymmetry of the distribution that is not represented by covariance alone (Natarajan et al., 2017). Moreover, we can utilize the techniques from this paper, hence adapt Proposition 3.5.1, for a distributionally robust heuristic under this additional information (see Appendix B.5 for the complete formulation). Figure 3.10 shows that the additional information could significantly reduce the expected cost of the distributionally robust inventory levels. It is no surprise that asymmetry information is important to estimate the impact of pooling, as this is in line with results from the pooling literature, specifically Yang and Schrage (2009) who show that right-skewed demand distributions can cause inventory levels to increase rather than decrease under pooling.

3.7.2 Experiments with realistic networks

3.7.2.1 Network Setup.

We now use realistic networks of fulfillment centers located in mainland US to study the performance of our heuristic solutions. We build these networks based on publicly available data from Chen (2017), who use unofficial data of a US-based online retailer's fulfillment center network. The dataset contains information about locations of the fulfillment centers, population of the US by zipcode, and estimated shipping costs based on UPS Ground and UPS Next Day Air from the fulfillment centers to customer locations.

We consider networks of size $n = 10$, by choosing n random fulfillment centers from the 87 fulfillment centers available in the data. Since the results of our experiments depend on network characteristics, we take a sample of 10^2 networks and conduct our experiments for each sample, reporting the mean values over all the networks for the metrics considered.

For each network, the mean demands at customer locations (approximated by zipcodes) are the population in millions, and each customer location demand is assigned to the nearest fulfillment center. That is, the fulfillment centers can fulfill demands from their assigned customer locations at the in-location fulfillment cost. The coefficient of variation is taken to be equal to 1. We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012), such that the maximum correlation coefficient has an absolute value less than 0.4. We take 10^3 samples of the demand vector for sample average approximations.

Similar to Jasin and Sinha (2015) and Lei et al. (2018), we take the fulfillment costs to be linear functions of the distance. Specifically we have $s_{ij} = s_0 + \lambda \cdot r_{ij}$, where r_{ij} is the distance in miles, $\lambda = 0.005$ is the distance sensitivity factor, with in-location fulfillment done at cost $s_0 = \$10$. This gives fulfillment costs in the range of $[\$10, \$23.6]$ for the entire network. The overage and underage cost parameters are taken to be: $h = \$10,\ p = \$50$. We use Algorithm 3 to generate the L-level nested fulfillment structure with $L = n$ as the base case.

3.7.2.2 Misspecifying the Distribution.

We study the effect of misspecifying demand distributions. In particular, we compare the expected costs under the nested fulfillment cost structure achieved by the following two inventory solutions:

1. The robust inventory solution $\mathbf{y}^{\mathrm{H}}$ derived from our proposed heuristic

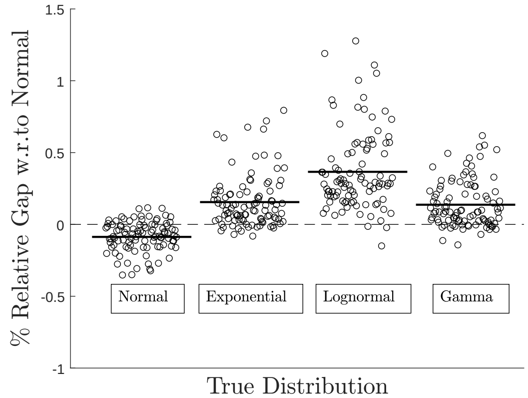

Figure 3.11: The relative gap in expected costs under the true distribution, of the robust solution $\mathbf{y}^{\mathrm{H}}$ and the Normal solution $\mathbf{y}^{\mathrm{N}}$. The x-axis spread of the data is solely for visual clarity.

2. The stochastic inventory solution $\mathbf{y}^{\mathrm{N}} = \mathrm{argmin}_{\mathbf{y}\geq 0}\mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})]$ that assumes f is the Normal distribution (the solution is approximated using sample average approximation).

The expected costs $\mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})]$ are calculated under the true distributions, where f is a Normal, exponential, lognormal, or gamma distribution (details are provided in Appendix B.4).

The results are shown in Figure 3.11. Each circle corresponds to a randomly chosen network of size $n = 10$. Indeed, if the true demand distribution were Normal, then $\mathbf{y}^{\mathrm{N}}$ will be the true optimal solution, in which case the relative gap in expected cost achieved by the robust solution is negative. However, this is not usually the case in reality, as the real joint distribution of demands can seldom be accurately predicted. We see that for certain networks, when the true distribution is non-Normal, significant reduction in expected costs can be realized by using the robust solution instead of the Normal solution. The savings are likely to be higher for larger networks as the normal distribution perceives higher pooling benefits which may not be the case under the true distribution.

3.7.2.3 Dynamic Myopic Fulfillment.

By considering a dynamic setting where demands arrive at random, we study the quality of the single-period assumption made in our study. We model random arrivals in the following fashion: we generate the single period demand vector, and randomize the sequence in which each unit of demand arrives. The decision on which location

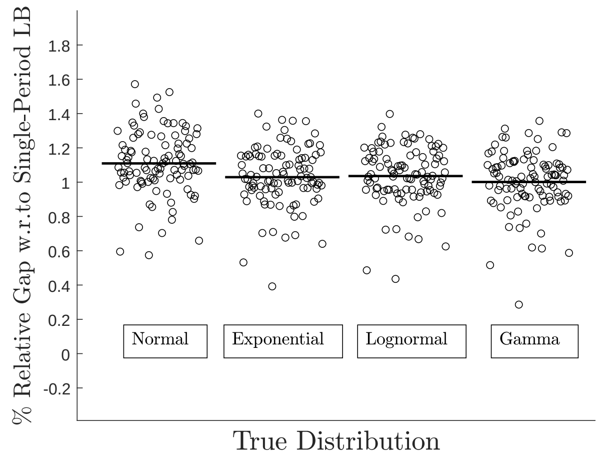

Figure 3.12: Figure showing the relative gap in expected costs between the dynamic setting under myopic fulfillment and the single-period lower bound for various distributions.

fulfills an incoming unit of demand is taken myopically – the nearest location with available inventory is chosen to fulfill the demand, which is the fulfillment norm in practice.

The starting inventory levels are set by the robust heuristic. Note that given any inventory levels, the single period cost is a *hindsight optimal* lower bound for the cost under the dynamic setting. We see in Figure 3.12 that the relative gap in expected costs between the single-period and dynamic settings is less than 2%, and hence the single-period expected cost can serve as a good approximation for the expected cost under a dynamic setting. Note that the myopic strategy need not be the optimal fulfillment strategy in a dynamic setting, and hence the actual relative gap in expected costs will be less than 2%. Similar results were observed when the nested fulfillment structure was used in place of the actual fulfillment costs.

3.8 Extensions

3.8.1 Location-specific Demand Classes

In the previous sections, we made the assumption that all demands can be fulfilled by inventory in any node, regardless of the demand location or the inventory node location. However, in some settings, there may be classes of demand that cannot be fulfilled by inventory nodes in a different location. One example is an omni-channel store network; in each location, there are two types of customers: those purchasing from the local brick-and-mortar store, and those placing an order through the online store. Demand from store customers can only be met with inventory that is located in the local store. On the other hand, demand from an online customer can be fulfilled from any store location, through what is known in the retail industry as *ship-from-store* fulfillment.

Note that there are two different types of inventory risk pooling involved here. First, within a location, store demand and local online demand are pooled since they deplete from the same store inventory. Second, online demand across locations are pooled since they are fulfilled from inventory in the store network. While ignoring the first type (for instance, by keeping a separate inventory for store customers) simplifies the problem to one explored in the previous section, it results in suboptimal inventory levels since it is likely that local demands are highly correlated.

Detailed analyses of this setting can be found in Appendix B.6. We show that this problem can also be analyzed in a similar fashion to Section 3.5, except that there are nested piecewise linear terms of the form $(x-(y-z)^+)^+$ in the objective. Introducing integer variables similar to what was done in (3.19) yields products of integer variables. We deal with this complication by introducing new integer variables to replace these product terms, and we obtain the heuristic in the form of an SDP of increased size, though still polynomial in the number of nodes).

3.8.2 Uncertainty in Moment Information

As e-commerce demands are highly volatile, there may be uncertainty in the moment information estimated from the data. Such uncertainty may be in the form of confidence intervals constructed around the moment information through empirical estimation from the data, or in the form of more complicated uncertainty sets. These can be incorporated easily into our models by simply including the constraint $(\mu, \Sigma) \in \mathcal{U}$, where $\mathcal{U}$ is a non-empty, closed and convex uncertainty set for the estimated mean and covariance matrices, and allowing μ and Σ to be variables that are constrained in the above fashion, rather than parameters (Natarajan et al., 2011).[5]

Natarajan et al. (2011) provide two examples of uncertainty set representations:

1. Linear: $\mathcal{U} = \{(\mu, \Sigma) : \mu_L \leq \mu \leq \mu_U, \Sigma_L \leq \Sigma \leq \Sigma_U\}$. This can simply be incorporated as linear constraints, for which the dual can be taken easily.

[5]Note that this modification is to be made before taking the dual SDP of the inner robust problem – the constraint $(\mu, \Sigma) \in \mathcal{U}$ is included in the SDP relaxation of the moment problem in maximization form in (3.21).

2. Ellipsoidal (Delage and Ye, 2010): $(\mu-\mu_0)^\intercal\Sigma_0^{-1}(\mu-\mu_0) \le \gamma_1, \quad \Sigma - 2\mu\mu_0^\intercal + \mu_0\mu_0^\intercal \preceq \gamma_2\Sigma_0$, where μ_0, Σ_0 are the estimated mean and covariance matrices, and γ_1, γ_2 are parameters. Notice that the first constraint is non-linear, but can be expressed as the following semi-definite constraint:

$$\begin{pmatrix} \gamma_1 & (\mu-\mu_0)^\intercal \\ (\mu-\mu_0) & \Sigma_0 \end{pmatrix} \succeq 0$$

We note that any uncertainty set that can be characterized by linear or semi-definite constraints can be included, as they easily yield themselves to dualizing.

3.9 Conclusion

Robust strategies are gaining importance in retail due to the increase in complexity arising from innovations. Particularly for e-commerce demands, incorrect forecasting may lead to disastrous results, as inventory planning is done at the network level. We provide a framework to analyze the distributionally robust newsvendor network problem where there are network flows after realization of uncertainty. We solve the two-location setting to yield a closed-form nested bound that serves as an analogue to the Scarf bound for a system with inventory pooling.

For the multi-location case, we provide a heuristic approximation and upper bound for the case where the fulfillment costs exhibit a nested fulfillment structure, where the cost function can be written as the sum of piecewise linear terms. We show how any general fulfillment cost structure can be approximated by this nested fulfillment structure through simple agglomerative clustering algorithms, and that the approximation of the expected total fulfillment cost is empirically tight for commonly seen distance-based shipping cost structures under various distributions.

Following Natarajan et al. (2017), we show that the value of asymmetry information is significant for a system with pooling, which also echoes results from pooling literature which state that the shape of the distributions have a significant effect on pooling benefits. We also demonstrate that a distributionally robust solution can significantly outperform stochastic inventory solutions that assume a particular demand distribution.

Multiple directions for future work exist. A multi-period formulation can be considered, where actions in the current period affect the future state. While tractable formulations can elude us, we can approximate future stages through an affine approximation, where the future actions are restricted to be affine functions of the corresponding data (Ben-Tal et al., 2004). Under such settings, robust fulfillment decisions can be analyzed which can yield helpful tools for practitioners to fulfill online demands. Our heuristic also yields the probability of stockout at the end of the period for each node in the nested fulfillment structure, which can also be used to guide dynamic fulfillment. Another natural extension is to consider how the network should look like in the first place: the solution from the inventory optimization can inform network design decisions, which is an important unexplored area in e-commerce.

CHAPTER 4

The Value of Personalized Return Policies in Retail

4.1 Abstract

With increasing online purchases and retail firms competing to provide lenient return policies, customer returns has become an important problem with roughly 10% of all the products being returned. In this paper, we consider a firm that jointly decides the inventory level and return policy at the start of the selling season. In particular, the return policy consists of two components – a return window and a return fee. The return policy affects the hassle experienced by the customer in returning the product, and customers are heterogeneous in how they perceive this hassle. In the absence of information about customers' hassle types, the firm can only offer blanket return policies, or a menu of return policies for customers to choose from; however, knowing customers' hassle types, the firm can offer personalized return policies. We show that the firm can achieve higher profits by offering a menu of policies as compared to blanket policies. When the firm provides full refunds and personalizes return windows, we show that, consistent with industry practice, low-hassle types (customers that return frequently under lenient policies) must be banned from returning. However when the firm can also personalize return fees, we show that the firm should prioritize sales to low-hassle types by offering them strict return policies (short return window and high return fee). We also show that personalization based on customer behavior data wipes out the customer surplus, providing implications for usage of customer behavior data in enforcing operational policies.

4.2 Introduction

The advent of e-commerce has infused intense competition in the retail industry, with firms competing for customers in both physical stores as well as digital platforms. In this competitive landscape, customers returns has emerged as a significant problem – in 2018, the total merchandise returns accounted for nearly $369 billion in lost sales for US retailers, which is around 10% of the total sales (Appriss Retail, 2018). Customer returns have been allowed by retail firms to incentivize purchases by alleviating customers' uncertainty over product *fit*, and this feature has become crucial for online purchases due to the lack of direct interaction with the product at the time of purchase.

Return rates in e-commerce sales are much higher (more than 30%) compared to brick-and-mortar sales (around 10%) (Saleh, 2016), due to lenient return policies which is one of the most important drivers of online purchases, next only to free shipping (Walker Sands, 2018). With online sales channels becoming a principal part of most retail firms, high return rates are accompanied by several problems; the product can be returned outside the selling season, the product can be in a state (damaged or used) that is not suitable for reselling, restocking costs, etc. In addition, many e-commerce firms also bear the brunt of return shipping fees, which is estimated to cost $550 billion by 2020, representing an increase of 75% over the costs in 2016 (Statista, 2018).

Retailers have adopted various strategies to reduce the cost of returns due to the problems mentioned above. Some firms have reduced the *return window* – firms such as LL Bean and Bed, Bath and Beyond, which had an unlimited return window (customers could return items years after purchase and did not need receipts to do so), have shortened the return window to one year, citing return abuse by certain customers (Rosato, 2018). While full refunds are common, some retailers charge customers upon return of certain items (e.g. electronic products) in the form of a flat restocking fee or *partial refunds*. To reduce return shipping costs on online purchases, firms may pass on the return shipping costs (wholly or partially) to the customer, or in the case of omnichannel retailers such as Macy's and Target (with physical stores), customers are allowed to return to one of their physical stores free of cost (which is later transported to return processing centers through internal logistics). While refunds and return windows are common levers in return policies, Janakiraman and Syrdal (2016) list other aspects of returns which firms can focus on, such as leniency on the cause of return, selectivity based on products and customers (members vs non-members), and starting the return policy later by allowing trial windows.

While these initiatives have helped reduce cases of misuse by customers, they still suffer from a fundamental problem – these policies are agnostic to variations in customers' return behaviors. Abbey et al. (2018) segment customers based on profitability using transaction data from a retailer into three groups – non-returners (47.7% of customers), legitimate returners (52% of customers, 23% return rate) and abusive returners (0.4% of customers, 60% return rate). The profit contributions were $1,445, $222 and -$1,254 per year respectively. It is clear that in such a case, a one-size-fits-all policy that aims to target the small fraction of abusive returners can hurt profits by adversely affecting the majority of profitable customers.

Indeed, strategies that target different segments of customers relies on the firm's knowledge of its customers. This is particular easy for e-commerce firms as huge amounts of data are collected when users interact with online platforms. Several brick-and-mortar retailers also track their customers with the use of third-party companies such as The Retail Equation, which creates a risk report for customers based on their historical return behavior (Safdar, 2018). Equipped with this data, firms can tailor return policies to their customers. In fact, many firms have already started engaging in this practice – Amazon.com (Safdar and Stevens, 2018) and Costco (Hanbury, 2018) have banned several customers that they identified as fraudulent returners, and 61% of US retailers are ready to ban frequent returners from shopping on their websites (Brightpearl, 2018). Motivated by this practice of targeting customers based on historical data, we address the following research questions in this paper:

1. Knowing information about customers' return behavior, how should a firm personalize return policies to maximize profits?

2. What is the value in collecting information about customers to offer personalized policies?

3. What is the effect of personalized policies on consumer welfare and surplus?

We employ a stylized, single-period newsvendor model to analyze the problem, where there is uncertainty in the aggregate demand as well as customers' valuation for the product. As return policies (partial refunds, return windows, return shipping fees) primarily affect the customers' *disutility* once they decide to return (and in turn their purchasing decision itself), we choose to model demand as a continuous mass of customers that are heterogeneous in their perceived hassle in returning items. Perceived hassle may include the physical hassle (distance or access to the nearest store or shipping center), as well as psychological hassle (mental characteristics inhibit returning items). Incidentally, the hassle types of customers can also be thought of as a proxy for return rates under lenient return policies – low-hassle customers are more likely to return items due to lower disutility in doing so, whereas high-hassle customers are less likely to exercise the return option.

We assume that the return policy consists of two components: 1) a return window (short or long), and 2) a refund fee. The firm benefits from a shorter return window, as returned items are less likely to be damaged and more likely to be resold during the selling season (modeled by a higher salvage price), whereas customers are inconvenienced by shorter windows (modeled by increase in their return hassle).

First, when the firm offers only a blanket policy for all customers, we show that for products that lose value quickly, short return window policy leads to higher profits by minimizing the loss due to returns. However, for products that have a relatively stable value over time (e.g. non-perishable goods), long return windows are more profitable as they increase sales by lowering the hassle due to returns. We also show that the firm can achieve higher profits compared to the blanket policy case by offering customers a menu of policies to choose from.

When the firm provides full refunds but personalizes the return window for each customer, we show that low-hassle customers (who return frequently under lenient return policies) must be banned from returning, consistent with industry practice of banning serial returners. However, when refunds can also be personalized, we show that when inventory is limited, the firm benefits by prioritizing sales to low-hassle customers under strict return policies (short return window and high return fees). When firms can identify customers based on their return behavior, we show that customer surplus is wiped out. This provides implications for customers' privacy in retail settings, and the value of data about customer behavior.

The rest of the paper is organized as follows. We discuss the relevant literature in Section 4.3. We then introduce the stylized model in Section 4.4, detailing customer decisions and the firm's profit maximization problem. In Section 4.5, we analyze retailer return policies in the absence of personalized information about customers' return behavior. In Section 4.6, we analyze the optimal personalized return policy. We conclude with extensions in Section 4.7 and future directions in Section 4.8.

4.3 Literature Review

Our research primarily belongs to the stream of literature on consumer return policies. In particular, there are several studies on the refund amount: Davis et al. (1995) and Che (1996) consider full-refund and no-refund return policies when customers face uncertainty in their valuation of the product. Lenient return policies have been shown to be useful tools – Moorthy and Srinivasan (1995) show that full refund policies can be used as a signal for high product quality, and Petersen and Kumar (2009) find that positive return experiences build trust in the firm, and in turn leads to positive behavioral outcomes.

In studying return policies, several studies jointly consider pricing and refund decisions. Su (2009) shows the optimality of partial return policies, with the optimal refund amount being equal to the salvage price, and also examine the impact of consumer returns on supply chain performance. Shulman et al. (2009) consider a two-product setting where products can also be exchanged, and find that the optimal refund need not be the same as the salvage price. Altug and Aydinliyim (2016) show that the optimal refund is bounded by the clearance price when customers are strategic when purchasing under full price, and offer explanations for the commonly seen "no-restocking-fee" return policies. Shang et al. (2017) consider price and restocking fee decisions when customers indulge in 'wardrobing' (misuse of trial periods intended to identify fit), and find that the optimal price and refund are decreasing in the extent of wardrobing among the customers.

Some researchers have also studied return policies in the form of return windows (time to return the product to obtain refund). Ülkü et al. (2013) consider a firm that decides the price and return window, where customer valuations and salvage price of inventory is affected by the return window. Ülkü and Gürler (2018) also considers the return window decision, however in the context of a newsvendor that decides inventory levels rather than price. Both these studies explicitly model fraudulent customers that buy products with no intention of keeping them, a feature also seen in Hess et al. (1996) and Chu et al. (1998). We take a different approach by attributing high return rates to customers who perceive little or no hassle in returning items that they deem unfit.

We model customer behavior and firm's policy based on the return hassle perceived by the customer. Davis et al. (1998) models hassle as a decision variable that the firm can set, and find that the retailer should offer low-hassle policies when the salvage value is high, or when there are opportunities for cross-selling, or the products' benefits cannot be consumed during a short period of time. There are two studies that explicitly model customer heterogeneity in hassle costs. Hsiao and Chen (2012) compare two cases: first, under full returns, the firm sets the optimal amount of hassle for the customers, and second, a hassle-free policy with partial returns. Hsiao and Chen (2014) study the interplay between price, return policy and quality risk, where the firm decides the quality of the product in addition to price and return policy. In both studies, customers are segmented into two groups – high hassle low hassle customers. In contrast, we use a more general model for hassle heterogeneity where the firm may not know the hassle-types or valuations of customers, and we focus on inventory decisions as compared to price optimization.

While it is common to examine pricing decisions along with return policies, there are several papers that place return policies in other contexts. Alptekinoğlu and Grasas (2014) study the optimal retail assortment when consumer returns are allowed. Similar to our setting, Akçay et al. (2013) and Ketzenberg and Zuidwijk (2009) consider inventory decisions along with pricing when product returns can be salvaged or resold. There have also been studies that analyze supply chain interactions due to consumer returns. Crocker and Letizia (2014) studies return policies between a manufacturer and retailer who receives customer returns; Shulman et al. (2010) study how retailer's optimal return policy is affected by the reverse channel structure, namely whether returns are salvaged by the manufacturer or by the retailer; and Ferguson et al. (2006) address the problem of reducing false failure returns through supply chain coordination methods, as manufacturers primarily incur the cost of these returns, whereas efforts to reduce returns are primarily taken by retailers.

There has been a growing literature recently on role of return policies specifically in the context of e-commerce and omnichannel firms. Nageswaran et al. (2017) study the pricing and return policy decisions of an omnichannel retailer, whose customers can purchase products in-store (and return in-store) or online (and return online or in-store). They find that generous refunds observed in practice are driven by customers' channel choice and the convenience of returning in-store. Ofek et al. (2011) consider the strategies of competing retailers with respect to opening an internet outlet, as online purchases lead to higher likelihood of costly product returns. Some firms like Jet.com have started offering the no-returns option to customers at a discounted price. Najafi and Duenyas (2018) consider a firm's decision to offer the no-returns option, while offering full refunds when purchased at the full price. Hsiao and Chen (2015) show that allowing retailers to implement the no-return option can sometimes improve supply chain efficiency by eliminating the manufacturer's attempt to induce inefficient cosumer returns.

Finally, we note that although customer returns has been well-researched, we provide a new perspective on this problem by considering the case where the firm can personalize return policies based on historical customer behavior. There have been papers that study personalization in other areas in retail, such as pricing (Fudenberg and Villas-Boas, 2006; Liu and Zhang, 2006; Choe et al., 2017), assortment (Golrezaei et al., 2014; Gallego et al., 2016), and advertising (Bleier and Eisenbeiss, 2015). To the best of our knowledge, our paper is the first to study behavior-based personalization in return policies in retail.

4.4 Model

We consider a monopolistic firm selling finite inventory of a single product at a fixed price p. We model the demand as a continuous mass of customers with a mean valuation V for the product, and their true valuation can only be realized after purchase of the product. This is a common assumption in literature, and is justified in the context of e-commerce retail where customers can only experience the product after purchase. If a customer chooses to buy the product, she realizes the true valuation $V + \epsilon$, where ϵ is drawn from a distribution $\mathcal{G}$. Once the product is purchased and the true valuation is realized, the customer can choose to return the product, in which case she incurs a hassle cost (this models the customer's proximity to store or ease of returning the product, or the customer's psychological hassle in returning the product), as well as a return fee (equivalent to partial refunds). [1] If the customer chooses not to buy the product, she leaves the system and receives a utility of 0.

Customers are characterized by their *hassle type* θ, which influences their perceived hassle in returning a product. In particular, a customer of type θ returning a product incurs a hassle cost that is linear in θ. We assume that each customer knows their hassle-type θ, and their purchasing decision takes into account the hassle cost of returns. However, the valuation uncertainty ϵ is only realized after purchase. We assume that θ and ϵ are independent and uniformly distributed: $\theta \sim U[0, \bar{\theta}]$, and $\epsilon \sim U[\underline{\epsilon}, \bar{\epsilon}]$. [2]

In the base case where the firm does not have knowledge of each customer's hassle type, we assume that the distribution of the hassle types $\mathcal{H}$ across customers is known to the firm. The aggregate demand is assumed to be a random variable D following distribution $\mathcal{F}$. In addition to uncertainty from the aggregate demand, the firm's decisions also depend on customer behavior with respect to purchasing and return decisions, which in turn depend on customers' hassle types and valuations. Thus, the firm has knowledge of the distributions $\mathcal{F}$, $\mathcal{G}$ and $\mathcal{H}$.

At the start of the period, the firm decides the starting inventory level $y \geq 0$ (purchased at a per-unit cost c), as well as the return policy π, which maps each

[1] The return fee can also include any return shipping fees that the customer bears.

[2] We note here that the distribution of ϵ is constrained such that the resultant valuation $V + \epsilon$ is non-negative, and we also have $p < V + \bar{\epsilon}$ so that the price is not too high so that no one buys.

customer of hassle-type θ to a combination of a return window and a return fee: $\pi(\theta) = (T^{\pi}(\theta), r^{\pi}(\theta))$. The return policy affects the customers' disutility when they choose to return a product after purchase. In particular, when a customer of type θ purchases the product under policy $\pi(\theta)$ and chooses to return the product, they face a disutility comprised of two components – a perceived hassle cost of $\frac{\theta}{T^{\pi}(\theta)}$, and the return fee $r^{\pi}(\theta)$. [3] We note that in the case where the firm does not have access to each customer's type θ, the return window and return fees are necessarily independent of θ.

To simplify analysis, we assume that $T^{\pi}(\theta)$ can only take one of three values $T^{\pi}(\theta) \in \{T_{\emptyset}, T_S, T_L\}$, with $0 = T_{\emptyset} < T_S < T_L$ where T_S and T_L correspond to short and long return windows respectively, and $T_{\emptyset}$ corresponds to the no-return option. We make the implicit assumption that shorter return windows inconvenience the customer, and hence are associated with a higher hassle cost, whereas longer return windows provide flexibility to the customer, and hence are associated with lower hassle cost. We assume that the no-return option increases the perceived hassle cost to $-\infty$, and hence returning is not an option for any customer that is offered $T_{\emptyset}$.

In practice, returned items can rarely be resold at the full price as they lose some of their value, primarily due to two reasons – first, the item may be damaged during the time that it spends with the customer, and second, the item may be returned at a time outside the normal selling season. In such cases, the firm either spends restocking efforts to sell the product at a discount, or directly salvages the product (Phillips, 2018). Consistent to a single period setting, we assume that items that are returned within the short window carry a higher salvage price of s' compared to the salvage price s for unsold items or items that are returned later under the long window policy, i.e. $s' > s$.

Finally, we note that charging a return fee of r is equivalent to a refund amount of $p - r$. Thus, $r = 0$ is equivalent to full refunds, and $r > 0$ is equivalent to partial refunds. We assume that the firm cannot profit off a return: that is, the profit from a customer keeping the product is higher than the profit if she returns it. Mathematically, this means that $p - r^{\pi}(\theta) \geq s'$, or $r^{\pi}(\theta) \leq p - s'$.

4.4.1 Customer Decisions

We first analyze the decision of a customer once she arrives and observes the return policy offered to her: $(T^{\pi}(\theta), r^{\pi}(\theta))$. The customer decision tree is shown in Figure 4.1.

[3] We assume that the values of $T^{\pi}(\theta)$ are normalized to make the fraction $\frac{\theta}{T^{\pi}(\theta)}$ comparable to the return fee

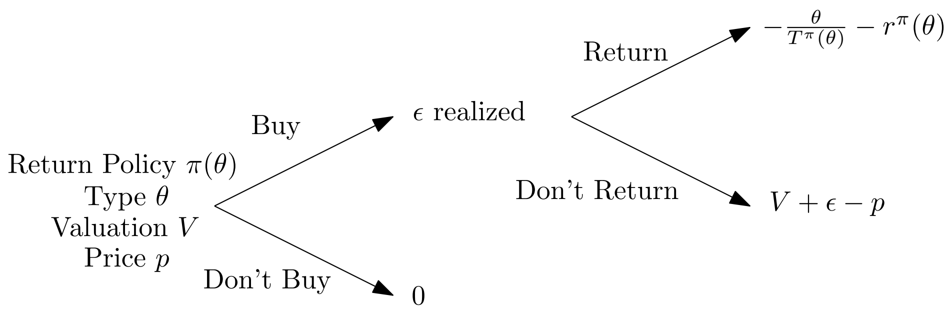

Figure 4.1: Figure showing the decision tree for customer of hassle-type θ when offered a return policy $\pi(\theta)$.

If a customer decides to purchase the product, she realizes her random valuation $V + \epsilon$ for the product. The customer will choose to keep the product if and only if the utility of keeping the product $(V + \epsilon - p)$ is greater than the utility from returning the product.

In our model, when the customer returns the product, she gets $p - r^{\pi}(\theta)$ as refund towards her purchase cost of p, but she also incurs a hassle cost of $\frac{\theta}{T^{\pi}(\theta)}$. Therefore, the utility from returning the product is $-\frac{\theta}{T^{\pi}(\theta)} - r^{\pi}(\theta)$ (we refer to the negative of this value as the *disutility of product returns*).

Suppose that the return disutility is equal to K. Then, the expected utility of purchasing the product is

$$U(K) := \mathbb{E}_{\epsilon}\left[\max\left(V + \epsilon - p, -K\right)\right]. \tag{4.1}$$

A customer will only purchase the product if the expected utility is at least equal to zero.

We note here that if $p \leq V$, the firm is better off not offering the return option to customers. This is because, all customers will buy the product irrespective of whether the return option is offered or not since

$$U(K) \geq \mathbb{E}_{\epsilon}\left[\max\left(V + \epsilon - V, -K\right)\right] \ \geq \ \mathbb{E}_{\epsilon}\left[\epsilon\right] = 0$$

The firm is better off if customers keep their products $(r^{\pi}(\theta) + s' \leq p)$, and hence allowing returns only serves to decrease the profits in this case. Offering the return option allows the firm to charge a higher price, as customers now have a recourse in case they are dissatisfied with the product. Hence, we will only assume that $p > V$, and we discuss the no-returns option as an extension in Section 4.7.

Since the customer type θ only influences the hassle cost, hence, the disutility of returns, we have a simple threshold rule to characterize customers' buying behavior. This is formalized in the following proposition.

Proposition 4.4.1 *There exists a threshold $\bar{K}$ (independent of θ) such that a customer buys the product if and only if her return disutility does not exceed this threshold. That is, a customer type θ will buy if and only if $\frac{\theta}{T^{\pi}(\theta)} + r^{\pi}(\theta) \leq \bar{K}$.*

The Proposition states that $\bar{K}$ is the maximum return disutility that customers are willing to accept before purchasing, and that this threshold is independent of θ. This threshold exists because the expected utility of buying the product (4.1) is non-increasing in K. Since ϵ is independent of the customer type θ, there exists a threshold $\bar{K}$ that is independent of θ and where $U(K) \geq 0$ for all $K \leq \bar{K}$ and $U(K) < 0$ for all $K > \bar{K}$.

An implication of the Proposition is that $\bar{K}$ is the maximum return fee that the firm will charge to any customer. While charging a return fee higher than $\bar{K}$ results in a higher profit for returns, this benefit will never be realized since the customer never buy the product. Therefore, we can assume that $\bar{K} \leq p - s'$. This is due to our earlier assumption that the return fee $r^{\pi}(\theta)$ charged to a customer cannot exceed $p - s'$.

Suppose that customers with type θ are offered policy $(r^{\pi}(\theta), T^{\pi}(\theta))$. There will be demand for the product from this customer type if the threshold rule in Proposition 4.4.1 is satisfied. Further, given that this customer purchases the product (contingent on inventory being available), the probability of a product return is

$$\mathbb{P}_{\epsilon}\left(V + \epsilon - p < -\frac{\theta}{T^{\pi}(\theta)} - r^{\pi}(\theta) \;\middle|\; \theta, \text{ buys}\right) = \mathcal{G}\left(p - V - \frac{\theta}{T^{\pi}(\theta)} - r^{\pi}(\theta)\right), \quad (4.2)$$

which is the probability that the valuation is too low to justify keeping the product.

4.4.2 Firm's Problem

We now formulate the firm's expected profit maximization problem. As a result of Proposition 4.4.1, customers that are offered $T_{\emptyset}$ do not buy the product (since $V < p$) and do not affect the profit function, and hence we can ignore these customers while formulating the profit function.

The salvage value of returned products depends on when the product is returned. Therefore, knowing when the product is returned is important for evaluating the firm's expected profit. We denote by ξ_S^{π} and ξ_L^{π} the fraction of customers that demand the product and are offered a short window and a long window, respectively. Mathematically,

$$\xi_S^\pi = \int\limits_{\theta \in \Omega_S^\pi} h(\theta) d\theta, \qquad \xi_L^\pi = \int\limits_{\theta \in \Omega_L^\pi} h(\theta) d\theta, \tag{4.3}$$

where

$$\Omega_S^\pi := \left\{ \theta \in [0, \bar{\theta}] : \ T^\pi(\theta) = T_S, \ r^\pi(\theta) \leq \bar{K} - \frac{\theta}{T_S} \right\}, \tag{4.4}$$

$$\Omega_L^\pi := \left\{ \theta \in [0, \bar{\theta}] : \ T^\pi(\theta) = T_L, \ r^\pi(\theta) \leq \bar{K} - \frac{\theta}{T_L} \right\}. \tag{4.5}$$

are the set of customers offered a short (long) return window and have a nonnegative utility for purchasing the product. Note that the total demand rate is $\xi^\pi = \xi_S^\pi + \xi_L^\pi$.

If there is infinite inventory, all customers with positive purchase utility will be able to buy the product. If this is the case, the rate of early returns and late returns are given by:

$$\int_{\theta \in \Omega_S^\pi} \mathcal{G}\left(p - V - \frac{\theta}{T^\pi(\theta)} - r^\pi(\theta) \right) h(\theta) d\theta, \quad \int_{\theta \in \Omega_L^\pi} \mathcal{G}\left(p - V - \frac{\theta}{T^\pi(\theta)} - r^\pi(\theta) \right) h(\theta) d\theta. \tag{4.6}$$

Note that the integration is over all customer types who generate demand for the product (Proposition 4.4.1). In the case of infinite inventory, these customers are able to purchase the product, and (4.2) is the return probability conditional on purchasing the product. Thus, the expression for the total salvage value is straightforward from (4.6).

On the other hand, if inventory is finite, the rate of early or late returns depends on how the inventory is allocated between customers offered a short return window and those offered a long window. For example, suppose that all inventory has been allocated to customers given a long return window. Then even if there is demand from customers with a short return window, the early return rate will be zero. Hence, expressing the total salvage value requires keeping track of which customers possess each unit of inventory.

In the interest of tractability, we make the following reasonable assumption: we assume that inventory is sold at the same rate of purchase for the two classes of customers Ω_S^π and Ω_L^π in the following fashion: if y is the inventory level, $\frac{\xi_S^\pi}{\xi^\pi} \cdot y$ is available for customers belonging to set Ω_S^π, and likewise, $\frac{\xi_L^\pi}{\xi^\pi} \cdot y$ for customers belonging to set Ω_L^π. Under this assumption, the firm's expected total sales is

$$\mathbb{E}\min\left(\frac{\xi_S^\pi}{\xi^\pi}y, \xi_S^\pi D\right) + \mathbb{E}\min\left(\frac{\xi_L^\pi}{\xi^\pi}y, \xi_L^\pi D\right), \tag{4.7}$$

where D is the random total demand.

We require a final assumption to be able to formulate the expression for the rate of early and late returns. This is due to the fact that the return behavior depends on the *specific* customer types in Ω_S^π and Ω_L^π that receive the product. For example, though the salvage value of returns from types θ_1 and θ_2 (where $\theta_1 < \theta_2$ and $\theta_1, \theta_2 \in \Omega_S^\pi$) is the same, selling the product to the former customer type will result in a lower probability of return. Hence, as before, expressing the rate of returns requires keeping track of which specific customer types receive the product. To address this issue, we assume that customers in Ω_S^π have equal chances of receiving the product, and similarly, customers in Ω_L^π have equal chances of receiving the product.

Under this last assumption, the expected return probability ψ_S^π of a customer selected at random among all who purchased in Ω_S^π is:

$$\begin{aligned}\psi_S^\pi &:= \mathbb{E}_\theta\left[\mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^\pi(\theta)\right) \;\Big|\; \theta \in \Omega_S^\pi\right] \\ &= \frac{1}{\xi_S^\pi} \cdot \int_{\theta\in\Omega_S^\pi} \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^\pi(\theta)\right) h(\theta)d\theta \end{aligned}\tag{4.8}$$

Similarly, the expected return probability ψ_L^π of a customer selected randomly from those who purchased in Ω_L^π is:

$$\begin{aligned}\psi_L^\pi &:= \mathbb{E}_\theta\left[\mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^\pi(\theta)\right) \;\Big|\; \theta \in \Omega_L^\pi\right] \\ &= \frac{1}{\xi_L^\pi} \cdot \int_{\theta\in\Omega_L^\pi} \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^\pi(\theta)\right) h(\theta)d\theta \end{aligned}\tag{4.9}$$

Note that, due to our assumption, these expected return probabilities do not depend on the actual sales from Ω_S^π or Ω_L^π. Hence, the expected total returns is equal to:

$$\psi_S^\pi \cdot \mathbb{E}\min\left(\frac{\xi_S^\pi}{\xi^\pi}y, \xi_S^\pi D\right) + \psi_L^\pi \cdot \mathbb{E}\min\left(\frac{\xi_L^\pi}{\xi^\pi}y, \xi_L^\pi D\right). \tag{4.10}$$

Since (4.10) distinguishes between early returns (the first term which has salvage value s') and late returns (the second term whose salvage value is s), we are able to compute the expected total salvage value.

To complete the specification of the firm's expected profit, we need to determine the expected total refunds that the firm will issue for product returns. Let R_S^π and R_L^π are the expected return fees collected from a customer selected at random from those who purchased in Ω_S^π and Ω_L^π, respectively, and who have returned the product. Mathematically,

$$\begin{aligned}
R_S^\pi &= \mathbb{E}_{\theta,\epsilon}\left[r^\pi(\theta) \;\middle|\; \theta \in \Omega_S^\pi,\ V + \epsilon - p < -\frac{\theta}{T_S} - r^\pi(\theta) \right] \\
&= \frac{1}{\xi_S^\pi \psi_S^\pi} \int\limits_{\theta \in \Omega_S^\pi} r^\pi(\theta) \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^\pi(\theta) \right) h(\theta) d\theta \qquad (4.11)
\end{aligned}$$

$$\begin{aligned}
R_L^\pi &= \mathbb{E}_{\theta,\epsilon}\left[r^\pi(\theta) \;\middle|\; \theta \in \Omega_L^\pi,\ V + \epsilon - p < -\frac{\theta}{T_L} - r^\pi(\theta) \right] \\
&= \frac{1}{\xi_L^\pi \psi_L^\pi} \int\limits_{\theta \in \Omega_L^\pi} r^\pi(\theta) \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^\pi(\theta) \right) h(\theta) d\theta \qquad (4.12)
\end{aligned}$$

Putting everything together, the firm's expected total profit is given by:

$$\begin{aligned}
\Pi(y,\pi) = &\underbrace{-c \cdot y + p \cdot \mathbb{E} \min\left(\frac{\xi_S^\pi}{\xi^\pi} y, \xi_S^\pi D \right) + p \cdot \mathbb{E} \min\left(\frac{\xi_L^\pi}{\xi^\pi} y, \xi_L^\pi D \right)}_{\text{profit from sales}} \\
&- \underbrace{(p - s' - R_S^\pi)\psi_S^\pi \mathbb{E} \min\left(\frac{\xi_S^\pi}{\xi^\pi} y, \xi_S^\pi D \right)}_{S \text{ returns salvaged at } s'} \\
&- \underbrace{(p - s - R_L^\pi)\psi_L^\pi \mathbb{E} \min\left(\frac{\xi_L^\pi}{\xi^\pi} y, \xi_L^\pi D \right)}_{L \text{ returns salvaged at } s} \\
&+ \underbrace{s \cdot \mathbb{E}\left(\frac{\xi_S^\pi}{\xi^\pi} y - \xi_S^\pi D \right)^+ + s \cdot \mathbb{E}\left(\frac{\xi_L^\pi}{\xi^\pi} y - \xi_L^\pi D \right)^+}_{\text{unsold items salvaged at } s}
\end{aligned}$$

Simplifying, we get:

$$\begin{aligned}
\Pi(y,\pi) = &(s - c) \cdot y \\
&+ \left(p - s - \frac{1}{\xi^\pi} \cdot \left[(p - s' - R_S^\pi)\, \xi_S^\pi \psi_S^\pi + (p - s - R_L^\pi)\, \xi_L^\pi \psi_L^\pi \right] \right) \cdot \mathbb{E} \min\left(y, \xi^\pi D \right)
\end{aligned} \qquad (4.13)$$

The firm's problem is to choose (y, π) so as to maximize its expected profit $\Pi(y, \pi)$ in (4.13).

We next discuss how the return policy affects the firm's expected profit. To do this, let us derive an alternative expression for the profit function by first defining

$$\tilde{p}^{\pi} := \frac{1}{\xi^{\pi}} \cdot \left[(p - s' - R_S^{\pi})\, \xi_S^{\pi} \psi_S^{\pi} + (p - s - R_L^{\pi})\, \xi_L^{\pi} \psi_L^{\pi} \right]. \tag{4.14}$$

Then, we have:

$$\Pi(y, \pi) = (s - c) \cdot y + (p - \tilde{p}^{\pi} - s) \cdot \mathbb{E} \min\left(y, \xi^{\pi} D\right) \tag{4.15}$$

We will assume that for all return strategies, $\tilde{p}^{\pi} \leq p - s$ and $c \leq p - \tilde{p}^{\pi}$, as otherwise, it is optimal to keep zero inventory. Comparing (4.15) with the profit in a traditional newsvendor setting, it seems at first glance that implementing a return policy reduces a firm's profit since it reduces the effective price at which the product is sold by $\tilde{p}^{\pi}$, while also reducing the demand. However, this is not true since the traditional newsvendor setting does not model customer choice, so all customers are willing to buy the product at price p even without the option of returns. In our setting, if price $p > V$ is offered without the option for returns, no customer will buy the product so the expected profit is zero. [4] Given a fixed price, the effect that a return policy has on the expected profit is two fold: first, it limits the effective demand, as some customers may choose to not buy the product due to higher price in spite of the return option; and second, allowing returns reduces the effective revenue, as the firm loses the sale and only obtains the salvage value of the product and any return fee it may have charged.

These effects demonstrate the tradeoff in setting a return policy – imposing a lenient return policy (i.e., by increasing the return window or by reducing the return fee) encourages more people to buy (reflected in a higher ξ^{π}); however, it also increases the revenue loss $\tilde{p}^{\pi}$ by either a lower salvage value or a higher rate of return. Strict return policies can serve to keep the revenue loss due to returns low, however it reduces the effective demand due to increased hassle in returns. Janakiraman and Syrdal (2016) advise that retailers "should approach return policies as a balancing act between increasing demand and limiting returns"; we thus capture the essence of return policies through our model as a tradeoff between limiting demand and limiting the revenue loss from returns.

[4] Offering returns allows the firm to charge a higher price $p > V$ than in the case where no returns are offered. For the purpose of this study, we assume that the price is exogenous to the model (we discuss optimizing the price when no-returns option is also offered in Section 4.7).

We have the following Proposition that establishes the structure of the profit function given a return strategy π.

Proposition 4.4.2 *Given a return strategy π, the expected profit function $\Pi(y, \pi)$ is concave in y, with the optimal inventory level $y^*(\pi)$ given by:*

$$y^*(\pi) = \xi^\pi \cdot F^{-1}\left(1 - \frac{c-s}{p-s-\bar{p}^\pi}\right), \tag{4.16}$$

Thus, the problem of optimizing the expected profit can be restated as a problem where the return strategy π is the only variable. However, optimizing $\Pi(y^*(\pi), \pi)$ can be complicated. In fact, even when the inventory level y is fixed, it is not immediately clear what the optimal return policy $\pi^*(y)$ looks like. In the following sections, we solve these problems under various cases:

1. Blanket Return Policies (π^{SB}, and π^{LB}) where all customers are offered the same return policy,

2. Menu of Return Policies (π^M) where all customers can choose among a set of return policies, and this set is the same for all customers,

3. Personalized Return Policies (π^*) where the return policy is specific to each customer type

4.5 Return Policies without Customer-Specific Information

In this section, we assume that all customers regardless of their type are offered the same return policy. This can be because either the firm does not have information about the individual customer hassle costs, or because the firm chooses not to implement a customer-specific return policy. Mathematically, we assume that the firm only knows the distribution of customer types $\mathcal{G}$, but that the return policy $\pi = (T, r)$ is independent of the customer type θ. As in Section 4.4.2, we ignore the no-returns option ($T_\emptyset$) since it is inconsequential to the profit function in these cases.

4.5.1 Blanket Policies

A typical practice seen in retail is to offer the same return policy $\pi = (T, r)$ to every customer. A natural question is: if the firm imposes a blanket return policy, should the return window be short or long? By shortening the return window, the firm can earn a higher salvage value for returned products; however, this is at the cost of a lower demand rate since the option of buying the product with a short return window is less attractive. In this section, we will investigate the conditions under which a firm should impose a short return window or a long return window.

We start with a simple result. The following corollary of Proposition 4.4.1 establishes the customer types that are willing to buy under a blanket return policy $\pi = (T, r)$:

Corollary 4.5.1 *The set of customer types willing to buy the product under blanket return policy $\pi = (T, r)$ is $\Omega = \left\{0 \leq \theta \leq \left(\bar{K} - r\right) T\right\}$.*

We next compare two blanket return policies:

1. The Short Blanket Policy, where all customers are offered $\pi^{SB} = (T_S, r_S)$, and
2. The Long Blanket Policy, where all customers are offered $\pi^{LB} = (T_L, r_L)$

Using Corollary 4.5.1 to calculate customer parameters that should be substituted in Equation 4.15, the expected profits of the firm given inventory level y and return fees r_{SB}, r_{LB} are given by:

$$\Pi(y, \pi^{SB}) = (s - c) \cdot y + \left(p - s - \tilde{p}^{SB}\right) \cdot \mathbb{E} \min\left(\xi^{SB} D, y\right) \tag{4.17}$$

$$\Pi(y, \pi^{LB}) = (s - c) \cdot y + \left(p - s - \tilde{p}^{LB}\right) \cdot \mathbb{E} \min\left(\xi^{LB} D, y\right) \tag{4.18}$$

where:

$$\tilde{p}^{SB} = (p - s' - r_{SB}) \cdot \psi^{SB} \tag{4.19}$$

$$\tilde{p}^{LB} = (p - s - r_{LB}) \cdot \psi^{LB} \tag{4.20}$$

Note that early returns under π^{SB} are salvaged at price s', whereas late returns under π^{LB} are salvaged at price s, however unsold inventory are salvaged at price s in both cases.

Given an inventory level $y \geq 0$, let $\Pi^{SB}(y)$ and $\Pi^{LB}(y)$ denote the optimal expected profit under a short-blanket and a long-blanket policy, respectively. That is,

$$\Pi^{SB}(y) = \max_{r_S \in [0,\bar{K}]} \Pi(y, T_S, r_S) \tag{4.21}$$

$$\Pi^{LB}(y) = \max_{r_L \in [0,\bar{K}]} \Pi(y, T_L, r_L) \tag{4.22}$$

We have the following Proposition which sheds a light on the choice of blanket policy based on the salvage price s':

Proposition 4.5.1 *Let $\Delta\Pi(y) = \Pi^{SB}(y) - \Pi^{LB}(y)$ denote the difference in the optimal expected profits under a short blanket policy and a long blanket policy. Then,*

1. *For low values of s' (close to s), $\Delta\Pi(y) \leq 0$, so the firm should impose a long blanket policy, and*
2. *$\Delta\Pi(y)$ is monotonically increasing in s'.*

The proposition implies that shortening return windows can lead to higher profits for products that lose value quickly with the time they spend with the customer. For products that have a relatively stable value over time (e.g. non-perishable consumer goods), longer return windows are more profitable as they boost sales by lowering the hassle cost of return. However for products that do not lose value and can be resold efficiently, shorter return windows may be preferable due to lower rates of return and increased values for returned goods.

While Proposition 4.5.1 provides guidance for the firm to choose between short and long window policies, it does not provide any information about the optimal return fee that the firm has to charge to maximize profits. Let $y^*(r)$ denote the optimal inventory level for a blanket policy that charges a return fee of r with a return window of T, given by Equation 4.16:

$$y^*(r) = \xi(r) \cdot F^{-1}\left(1 - \frac{c - s}{p - s - \tilde{p}(r)}\right). \tag{4.23}$$

The effect of increasing r on $y^*(r)$ is characterized by two effects: first, the demand effect, which is due to the decrease in the demand rate $\xi(r)$, and second, the revenue effect, which is due to the decrease in the revenue loss $\tilde{p}(r)$. The following Proposition establishes the structure of the expected profit function in terms of the return fee r.

Proposition 4.5.2 *Given a return window T, when the demand effect of increasing return fees is dominant (i.e. $y^*(r)$ is decreasing in r), then $\Pi(y^*(r), T, r)$ is unimodal*

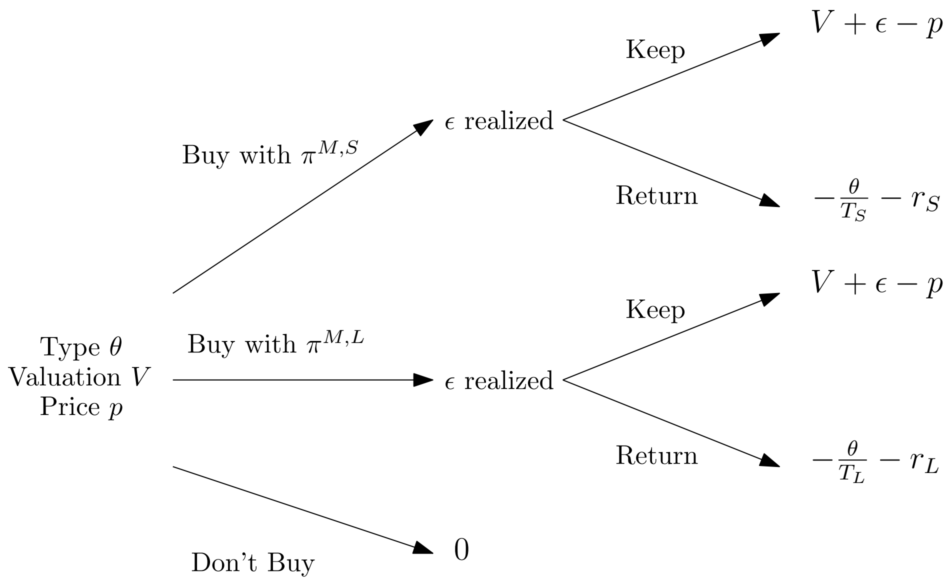

Figure 4.2: Figure showing the decision tree for a customer of hassle-type θ, when offered a menu of policies $(\pi^{M,S}, \pi^{M,L})$.

in r. Furthermore, there exists a threshold $\bar{c}$ such that whenever $c > \bar{c}$, it is optimal for the firm to charge a return fee $r^* > 0$. For low values of c, it is only optimal for the firm to offer full refunds ($r^* = 0$) when the revenue loss due to returns is not too large.

4.5.2 Menu of Policies

While blanket policies are commonly seen in practice, in the absence of customer-specific information, the firm can also offer a menu of policies and allow customers to choose the policy that suits them best. We analyze the case where the firm offers two policies $\pi^{M,S} : (T_S, r_S)$ and $\pi^{M,L} : (T_L, r_L)$.

The decision tree of the customer's decision is shown in Figure 4.2.

As the customers are free to choose their return policy, they choose the policy that maximizes their expected utility. First, following Proposition 4.4.1, we have the following Corollary:

Corollary 4.5.2 *There exists thresholds $\theta_{M,S}$ and $\theta_{M,L}$ where*

$$\theta_{M,S} = \left(\bar{K} - r_S\right) \cdot T_S$$
$$\theta_{M,L} = \left(\bar{K} - r_L\right) \cdot T_L,$$

such that customers with $\theta > \theta_{M,S}$ will never buy the product under $\pi^{M,S}$, and similarly

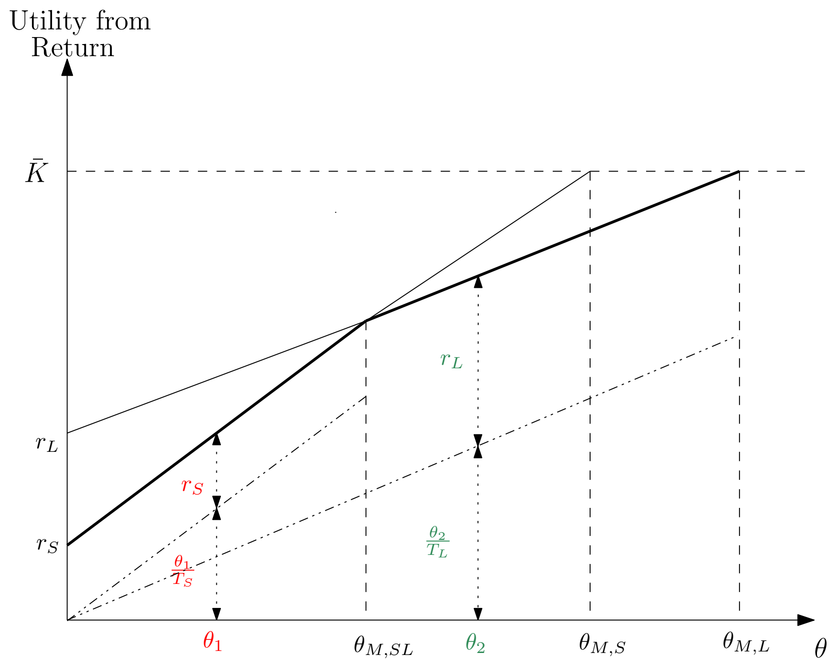

Figure 4.3: Figure showing the customer adoption under a menu of policies $(\pi^{M,S}, \pi^{M,L})$. Customer θ_1 (in red) buys with policy $\pi^{M,S}$, and customer θ_2 (in green) buys with policy $\pi^{M,L}$. The darker lines represent the choice of customers corresponding to the option that yields the highest expected utility (or lowest return disutility).

customers with $\theta > \theta_{M,L}$ will never buy the product under $\pi^{M,L}$.

Corollary 4.5.2 is illustrated in Figure 4.2. Note that a customer type θ will only buy if her disutility from product returns is less than the threshold $\bar{K}$. The figure plots the return disutility as a linear function of θ (each return policy in the menu is associated with one disutility function). $\theta_{M,S}$ and $\theta_{M,L}$ are the points where these disutility functions cross $\bar{K}$. The figure also shows that the fee charged for product returns coincides with the vertical intercept of the disutility function. So increasing the fee charged for product returns will result in a shift up of the disutility function (i.e., less customers would be willing to buy).

Given the option between $\pi^{M,S}$ and $\pi^{M,L}$, a customer will choose the return policy that results in the lower disutility of product returns. Hence, we can use the lower envelope of the two return disutility functions in Figure 4.3 to determine which return policy each customer type will choose. This is formalized in the following proposition:

Proposition 4.5.3 *Let $\Omega^{M,S}$ and $\Omega^{M,L}$ denote the sets of customers who buy the product under $\pi^{M,S}$ and $\pi^{M,L}$, respectively, when offered the menu of return policies. Assuming that $r_S, r_L \in [0, \bar{K}]$ and $r_S \leq r_L \leq r_S \cdot \frac{T_S}{T_L} + \bar{K} \cdot \left(1 - \frac{T_S}{T_L}\right)$, then*

1. $\Omega^{M,S} = \{\theta \leq \theta_{M,SL}\}$,

2. $\Omega^{M,L} = \{\theta_{M,SL} \leq \theta \leq \theta_{M,L}\}$

where $\theta_{M,SL} = \frac{r_L - r_S}{\frac{1}{T_S} - \frac{1}{T_L}} \in [0, \theta_{M,S}]$.

The conditions of Proposition 4.5.3 are needed so that the customer choice is non-trivial. Note that we require $r_S \leq r_L$ since, otherwise, all customers will choose $\pi^{M,L}$ over $\pi^{M,S}$. Further if r_L is too high, no customer will ever choose $\pi^{M,L}$. Hence, we need $\frac{\theta_{M,S}}{T^L} + r_L \leq \bar{K}$.

Let $\Pi(y, \pi^M)$ denote the expected profit under the menu of policies π^M, and let $\Pi^M(y)$ denote the maximum expected profit given an inventory level $y \geq 0$, given by:

$$\Pi^M(y) = \max_{\substack{r_S \in [0, \bar{K}], \\ r_L \in \left[r_S, r_S \cdot \frac{T_S}{T_L} + \bar{K} \cdot \left(1 - \frac{T_S}{T_L}\right)\right]}} \Pi(y, \pi^M) \tag{4.24}$$

Proposition 4.5.4 $\Pi^M(y) \geq \max\left(\Pi^{SB}(y), \Pi^{LB}(y)\right), \quad \forall y \geq 0.$

Interestingly, Proposition 4.5.4 shows that if the firm is able to optimize the refund fees, its profit from offering a menu of return policies is higher than what it can earn when it imposes a blanket return policy. Thus the firm can achieve higher expected profits by allowing customers to choose from a menu of policies, rather than offering a single blanket policy.

4.6 Personalized Return Policies with Customer-Specific Information

In this Section, we establish important structural results about the optimal personalized return strategy that the firm should follow to maximize its expected profits. In order to understand the optimal policy better, and to contrast with industry practices, we first analyze the case where the retailer offers full refunds (i.e., $r = 0$), but can personalize the return window according to the customer types.

4.6.1 Personalized Return Policies with Full Refunds

In the case of the firm offering full refunds, we have $r^\pi(\theta) = 0$ for all θ, but the firm can choose $T^\pi(\theta) \in \{T_\emptyset, T_S, T_L\}$. First, we rewrite the expected profit function from

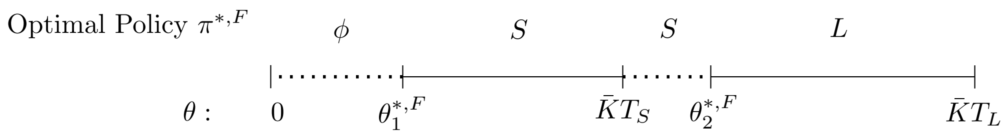

Figure 4.4: Figure showing the optimal personalized policy under full refunds. The dotted regions correspond to customers not buying the product.

Equation 4.13 for the case of full refunds as follows:

$$\Pi(y, \pi) = (s - c) \cdot y + \left(p - s - \frac{1}{\xi^\pi} \cdot \left[(p - s')\, \xi_S^\pi \psi_S^\pi + (p - s)\, \xi_L^\pi \psi_L^\pi\right]\right) \cdot \mathbb{E} \min\left(y, \xi^\pi D\right) \tag{4.25}$$

We have the following Proposition that establishes the structure of the optimal policy $\pi^{*,F}$:

Proposition 4.6.1 *Given inventory level $y \geq 0$, there exists an optimal policy such that:*

1. *For any $\theta \leq \bar{K}T_S$, if $T^{*,F}(\theta) = T_\emptyset$, then for any other $\theta' < \theta$, $T^{*,F}(\theta') = T_\emptyset$.*
2. *For any $\theta \leq \bar{K}T_L$, if $T^{*,F}(\theta) = T_S$, then for any other $\theta' < \theta$, $T^{*,F}(\theta') = T_S$ or $T^{*,F}(\theta') = T_\emptyset$.*
3. *If $\theta \leq \bar{K}T_S$, $T^{*,F}(\theta) \neq T_L$.*

We restate the proposition in the following Corollary:

Corollary 4.6.1 *Given an inventory level $y \geq 0$, the optimal return policy $\pi^{*,F}$ that maximizes the expected cost has a threshold structure:*

$$\pi^{*,F} = \begin{cases} T_\emptyset, & \text{if } 0 \leq \theta \leq \theta_1^{*,F} \\ T_S, & \text{if } \theta_1^{*,F} < \theta \leq \theta_2^{*,F} \\ T_L, & \text{if } \theta_2^{*,F} < \theta \leq \bar{K}T_L \end{cases} \tag{4.26}$$

where $0 \leq \theta_1^{*,F} \leq \bar{K}T_S \leq \theta_2^{*,F} \leq \bar{K}T_L$.

This paints an intuitive picture of the optimal policy, shown in Figure 4.4. We see that the optimal policy bans customers of low-hassle types from buying the product, which is consistent with the strategies seen in practice by Amazon.com and Costco, who identify customers who return frequently and ban them from purchasing products.

We also note that there is another set of customers that may not end up purchasing the products – as opposed to low-hassle customers for whom offering the option of return is too lenient, we have customers for whom offering the long-return window is too lenient. As a result, they are only offered the short-return window, and hence they do not purchase the product.

4.6.2 Personalized Return Policies with General Refunds

We now move on to the general case where return fees can be applied to customers. We first establish the importance of return fees. It is apparent that the expected profit can be increased, if the effective revenue loss due to returns $\tilde{p}^{\pi}$ can be decreased (keeping ξ^{π} the same). One way of doing this is to increase the return fee for customers – increasing return fees can reduce the rate of returns, as the utility in returning the item decreases. Additional revenue is also obtained from returning customers in the form of increased return fees. The following lemma establishes the dependency of the effective revenue loss on the return policy:

Lemma 4.6.1 *Let π_1 and π_2 be two return policies such that*

1. $T^{\pi_1}(\theta) = T^{\pi_2}(\theta) \ \big(:= T^{\pi}(\theta)\big)$, $\forall \theta$.
2. *for any θ with $r^{\pi_2}(\theta) \leq \bar{K} - \frac{\theta}{T^{\pi}(\theta)}$, it is also true that $r^{\pi_2}(\theta) \leq r^{\pi_1}(\theta) \leq \bar{K} - \frac{\theta}{T^{\pi}(\theta)}$.*

Then, $\tilde{p}^{\pi_1} \leq \tilde{p}^{\pi_2}$.

The lemma solidifies the intuition behind the benefits of increasing the return fees. Keeping the set of purchasing customers fixed, increasing the return fee for these customers can increase the expected profit. It is to be noted that the firm cannot charge customers drastically high return fees, as this can dissuade the customers from purchasing the product in the first place.

Let π^* be the optimal personalized return strategy which maximizes the expected profit given any inventory level y (π^*'s dependence on y is abbreviated for notational clarity), and π^* assigns each θ to a return window $T^*(\theta)$ and a return fee $r^*(\theta)$. Let $\Omega^{\pi} = \Omega_S \cup \Omega_L$ denote the set of customers that will buy the product.

With the help of Lemma 4.6.1, we have the following Proposition which establishes the structure of the optimal personalized return policy π^* given any inventory level y:

Proposition 4.6.2 *For any given y, there exists a profit maximizing return policy π^*, corresponding to return fee $r^*(\theta)$ and return window $T^*(\theta)$ for every θ such that the following are true:*

1. $\forall\theta \in \Omega^\pi$, $-\frac{\theta}{T^*(\theta)} - r^*(\theta) = -\bar{K}$.

2. For $\theta \leq \bar{K}T_S$, *if* $\theta \in \Omega^\pi$, *then* $\theta \in \Omega^S$ *if and only if* $\theta \leq \frac{s'-s}{\frac{1}{T_S}-\frac{1}{T_L}}$.

3. If $\theta \in \Omega^\pi$, *then for all* $\theta' < \theta$, $\theta' \in \Omega^\pi$.

Proposition 4.6.2 (1) reiterates the result from Lemma 4.6.1, by stating the customers should be charged the maximum return fee that they are willing to accept to purchase the product. However, it stops short of addressing which customers ought to be allowed to purchase the product. This is addressed by Proposition 4.6.2 (2) and (3), which state that customers with lower values of θ are preferred by the firm to purchase the product. To understand these better, we have the following Corollary:

Corollary 4.6.2 *An optimal policy that has the properties in Proposition 4.6.2 can be obtained by optimally choosing* $\theta^* \leq \bar{K}T_L$, *and implementing the following policy:*

$$\pi^*(\theta^*) = \begin{cases} \left(T_S,\ \bar{K} - \frac{\theta}{T_S}\right), & \text{if } \theta \leq \min(\theta^*, \theta_1^*) \\ \left(T_L,\ \bar{K} - \frac{\theta}{T_L}\right), & \text{if } \min(\theta^*, \theta_1^*) < \theta \leq \theta^* \\ T_\emptyset, & \text{if } \theta > \theta^* \end{cases} \tag{4.27}$$

where

$$\theta_1^* = \min\left(\frac{s'-s}{\frac{1}{T_S} - \frac{1}{T_L}},\ T_S \cdot \bar{K}\right) \tag{4.28}$$

Note that for any $\theta \leq \tilde{\theta}_S$ (where $\tilde{\theta}_S = T_S \cdot \bar{K}$), two policies can satisfy the condition in Proposition 4.6.2 (1): $\left(T_S, \bar{K} - \frac{\theta}{T_S}\right)$ and $\left(T_L, \bar{K} - \frac{\theta}{T_L}\right)$. Both these policies yield the same return disutility to the customer, and hence the probability of return is also the same. When the customer returns under T_S, a higher salvage value s' is obtained, and a return fee of $\bar{K} - \frac{\theta}{T_S}$ is collected. However, when the customer returns under T_L, a lower salvage value s is obtained, but a higher return fee $\bar{K} - \frac{\theta}{T_L}$ is collected. Thus, θ_1^* is the value at which these two are equal, which yields Equation 4.28.

Thus, the profit-maximizing policy is given by the optimization problem:

$$\Pi^*(y) = \max_{\theta^* \leq \bar{K}T_L} \Pi(y, \theta^*) \tag{4.29}$$

Finally, since any return policy is a feasible solution to the personalized return policy, we have the following Proposition:

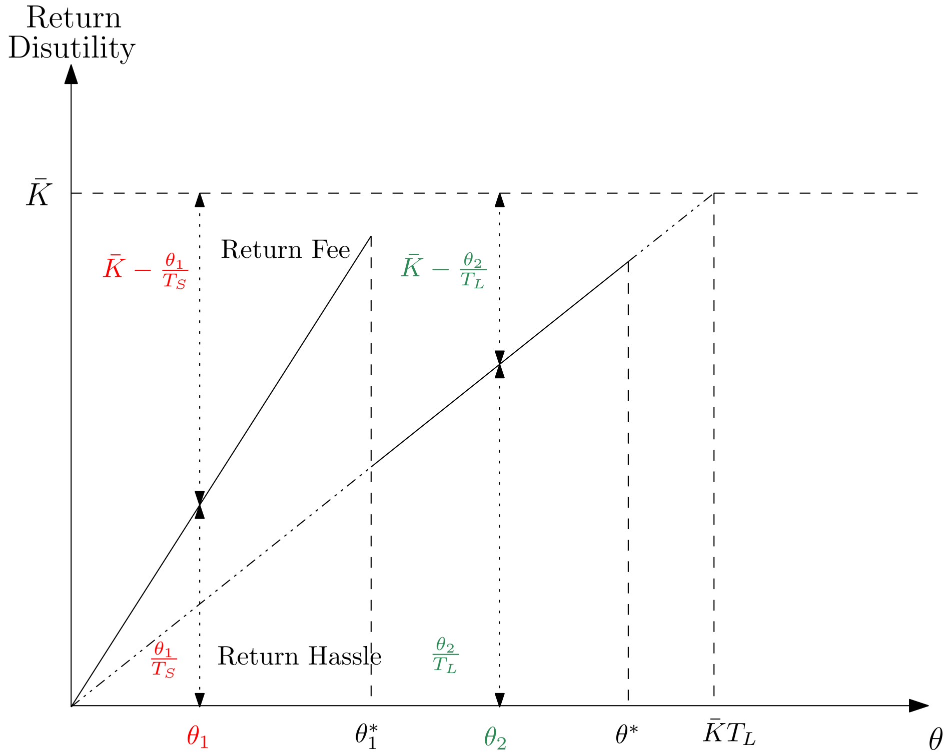

Figure 4.5: Figure showing the structure of the optimal personalized policy with general refunds. First, a threshold θ^* is chosen, which determines the demand rate. Then, customer θ_1 (in red) is offered a policy $(T_S, \bar{K} - \frac{\theta_1}{T_S})$, and customer θ_2 (in green) is offered a policy $(T_L, \bar{K} - \frac{\theta_2}{T_L})$.

Proposition 4.6.3 $\Pi^*(y) \geq \Pi^M(y),\ \forall y \geq 0.$

Thus from Propositions 4.5.4 and 4.6.3, we have:

$$\max\left(\Pi^{SB}(y), \Pi^{LB}(y)\right) \leq \Pi^M(y) \leq \Pi^*(y), \quad \forall y \geq 0 \tag{4.30}$$

Additionally, since customers are offered the maximum return fee that they are willing to pay, the firm extracts all the surplus from the customers. The structure of the optimal strategy is illustrated in Figure 4.5.

Note that the optimal policy prioritizes low-hassle type customers as opposed to high-hassle customers, which is in direct contrast to the case where full refunds are offered. This is because, when partial refunds are allowed, it is more profitable to extract the surplus from low-hassle customers than to ban them from purchasing the product. This suggests that firms such as Amazon.com and Costco may be better off in offering strict return policies to customers who return frequently rather than banning them outright, as if the firm can identify these customers, the firm can also extract all their surplus.

4.7 Extensions

4.7.1 No Returns Option

Note that we have trivialized the no-returns option so far, as whenever $p > V$, when customers are offered the no-returns option they will choose to not buy the product. However, we now consider a case where the no-returns option is provided at a lower price equal to V. This is common practice for retailers like Jet.com that provide the no-return option at a lower price compared to the regular option of buying the item with a return option.

We can extend our analyses easily to the case where customers have an option to buy the product with the no-return option at a lower price V, compared to the option of buying the product at the regular price of $p > V$ under a return policy π. In this case, any customer that did not buy in the original formulation will now buy the product under the no-returns option at price V, as this gives them a non-negative (zero) expected utility.

In Section 4.4.2, we thus include an additional set of buying customers:

$$\xi_\emptyset^\pi = \int_{\theta\in\Omega_\emptyset^\pi} h(\theta)d\theta \tag{4.31}$$

where $\Omega_\emptyset^\pi := (\Omega_S^\pi \cup \Omega_L^\pi)^\complement$. Note that the total demand rate is now $\xi^\pi = \xi_S^\pi + \xi_L^\pi + \xi_\emptyset^\pi = 1$. That is, $\xi_\emptyset^\pi = 1 - \xi_S^\pi - \xi_L^\pi$. The firm's expected total profit is thus:

$$\begin{aligned}
\Pi(y,\pi) = &-c\cdot y + p\cdot\mathbb{E}\min\left(\xi_S^\pi y, \xi_S^\pi D\right) + p\cdot\mathbb{E}\min\left(\xi_L^\pi y, \xi_L^\pi D\right) + V\cdot\mathbb{E}\min\left(\xi_\emptyset^\pi y, \xi_\emptyset^\pi D\right)\\
&-(p-s'-R_S^\pi)\psi_S^\pi\mathbb{E}\min\left(\frac{\xi_S^\pi}{\xi^\pi}y, \xi_S^\pi D\right)\\
&-(p-s-R_L^\pi)\psi_L^\pi\mathbb{E}\min\left(\frac{\xi_L^\pi}{\xi^\pi}y, \xi_L^\pi D\right)\\
&+s\cdot\mathbb{E}\left(\xi_S^\pi y - \xi_S^\pi D\right)^+ + s\cdot\mathbb{E}\left(\xi_L^\pi y - \xi_L^\pi D\right)^+ + s\cdot\mathbb{E}\left(\xi_\emptyset^\pi y - \xi_\emptyset^\pi D\right)^+
\end{aligned}$$

Simplifying, we get:

$$\Pi(y,\pi) = (s-c)\cdot y + (p - \tilde{p}^\pi - s)\cdot\mathbb{E}\min\left(y, D\right) \tag{4.32}$$

where $\tilde{p}^{\pi} = [(p - V) \cdot (1 - \xi_S^{\pi} - \xi_L^{\pi}) + (p - s' - R_S^{\pi})\, \xi_S^{\pi}\psi_S^{\pi} + (p - s - R_L^{\pi})\, \xi_L^{\pi}\psi_L^{\pi}]$.

We note that *all* the results in our paper continue to hold in this setting. In addition, as can be seen from Equation 4.32, the inventory and return policy decisions are decoupled, as the total demand rate is equal to 1. Thus, we can first optimize the return policy by solving the problem

$$\min_{\pi} \tilde{p}^{\pi} \tag{4.33}$$

and substituting the optimal policy π^* into:

$$y^* = F^{-1}\left(1 - \frac{c - s}{p - s - \tilde{p}^{\pi^*}}\right) \tag{4.34}$$

We also note that in some cases (such as personalized return policy with partial refunds), where the optimal return policy has a simple structure, we can obtain the optimal threshold in closed-form in terms of the price. This also implies that we can find the optimum price that the firm should charge to maximize profits, as the price optimization can also be done independent of the inventory situation.

4.8 Conclusion

With recent innovations in retailing due to the rapid rise in sales conducted over the internet, retail firms are increasingly focusing on personalizing customer experiences through collected historical data about customer behavior. Personalized policies have been implemented in various areas such as pricing, promotions, recommendations, bundle deals, etc., and firms are competing to innovate in this space. We address the important problem of consumer returns in retail by analyzing the value of personalized return policies that are tailored based on customers' historical return behavior.

When the firm does not have access to customer's individual behavior information, we show that the firm can achieve higher expected profits by offering customers a menu of policies as compared to a single blanket policy. Equipped with information about individual customers, the firm can offer personalized return policies – we show that when the firm offers full refunds, customers with low-hassle types (those that return frequently under lenient policies) ought to be banned from returning items, as seen in practice with firms like Amazon.com and Costco banning customers for returning too frequently. However, when the firm can personalize refunds, we find that it is better to sell to the low-hassle customers, as they can be charged higher return fees as compared to high-hassle customers.

Several future directions exist. One stream of research that can follow is personalizing other features that affect customers' experience. One such feature where blanket policies are common is shipping policies, where firms apply one policy to every customer irrespective of their locations. Another avenue for customization is fulfillment options, where the firm can strategically offer incentives for customers to select fulfillment options that reduce the firm's fulfillment costs.

While our model tries to capture different features such as product value uncertainty and hassle due to return policies, customer behavior in reality can be quite complex. For instance, Janakiraman and Syrdal (2016) show that increasing the return window can lead to reduced return rates, which is attributed to endowment effect – customers' valuation for the product increases with the time that they spend keeping the product. This information asymmetry cannot be captured in our model, as customers do not anticipate the endowment effect while deciding to purchase the item, but the firm does. An additional difficulty that arises in this feature is heterogeneity in customer's endowment effects arising from keeping the product for longer times. We leave this as an opportunity for future research.

As more and more firms are getting access to data about their customers from various sources (in-house, social media, third-party companies, etc.), personalization is at the forefront of retail innovation. While personalization can improve customer satisfaction in certain areas (product recommendations), as we show in our study, personalization is not always beneficial to customers – the firm's effort to maximize profits through personalized strategies based on data about customers' behavior comes directly at the cost of customer surplus. This provides policy implications for usage of customer-centric information, especially in monopolistic markets, where customer choice is limited.

APPENDIX A

Appendix to Chapter 2

A.1 Proofs

A.1.1 Proof of Proposition 2.5.1

We first observe that given a realization of the demands, the optimal cost can be obtained using a linear program. The proof follows in similar fashion to Seifert et al. (2006, Proposition 1). Consider the linear program $P(y_1, y_2, \tilde{D})$, where z_i represents the amount of inventory at R_i used to fulfill its in-store demand, and z_{ij} represents the amount of inventory of R_i used to fulfill online demand from region j.

$$
\begin{aligned}
P(y_1, y_2, \tilde{D}) = \quad \min_{z_i, z_{ii}, z_{ij}} \sum_i h(y_i - z_i - & \sum_j z_{ij}) + \sum_i p_s(D_{is} - && z_i) \\
& + \sum_i p_o(D_{io} - \sum_j z_{ji}) && + \sum_i s z_{ii} + \sum_i \sum_{j \neq i} s_{ij} z_{ij} \\
\text{subject to } z_i + \sum_j z_{ij} & \leq y_i, && \forall i \\
z_i & \leq D_{is}, && \forall i \\
\sum_j z_{ji} & \leq D_{io}, && \forall i \\
z_i, z_{ij} & \geq 0, && \forall i, j
\end{aligned}
\tag{A.1}
$$

To show that the function P represents C^{IIP} for a given demand $\tilde{D}$, notice that the coefficients of the decision variables z_i, z_{ii}, $z_{ij,(j \neq i)}$ in the objective function follow $(-h - p_s) < (s - h - p_o) < (s_{ij} - h - p_o)$, under the conditions in Ψ in Equation 2.4. The linear program can be solved greedily, and it is easy to see that the optimal solution is given by $z_i = \min(y_i, D_{is})$, $z_{ii} = \min\big((y_i - D_{is})^+, D_{io}\big)$, $z_{ij} = \min\Big(\big((y_i - D_{is})^+ - D_{io}\big)^+, \big(D_{jo} - (y_j - D_{js})^+\big)^+\Big)$.

The sequence of fulfillment is clear: in-store demand is fulfilled first, followed by online demand from the same region, and finally cross-shipment to other regions. Hence, we have $C^{IIP}(y_1, y_2) = \mathbb{E}_{\tilde{D}}\left(P\left(y_1, y_2, \tilde{D}\right)\right)$. The objective function is linear and the constraint set in (A.1) is a polyhedral convex set with linear constraints, and hence by Heyman and Sobel (2003, Proposition B-4), P is jointly convex in y_1, y_2, $\tilde{D}$. As the expectation of a convex function is convex, it follows that $C^{IIP}(y_1, y_2)$ is jointly convex in y_1 and y_2.

The structure of C^{IIP} as an expectation of a linear program draws direct comparison with the value function in newsvendor networks (van Mieghem and Rudi, 2002). Similar to Proposition 2 in Harrison and van Mieghem (1999), the gradient of the function $P(y_1, y_2, \tilde{D})$ with respect to (y_1, y_2) can be written as:

$$\nabla_{y_1,y_2} P\left(y_1, y_2, \tilde{D}\right) = (h, h)^T - \lambda\left(y_1, y_2, \tilde{D}\right)$$

where $\lambda(y_1, y_2, \tilde{D})$ is the dual-price vector corresponding to the constraints with y_1 and y_2 in (A.1). The 4-dimensional demand space can be divided into domains $\Omega_k\left(y_1, y_2\right)$ such that in each domain, the optimal values of the decision variables z_i, z_{ii} and z_{ij} are linear in y_1 and y_2, and hence the dual-price vector $\lambda(y_1, y_2, \tilde{D})$ is constant (refer to Appendix B for a discussion). The first-order conditions are:

$$\nabla_{y_1,y_2} C^{IIP}(y_1, y_2) = 0 = \nabla_{y_1,y_2} \mathbb{E}_{\tilde{D}}\left(P\left(y_1, y_2, \tilde{D}\right)\right) \tag{A.2}$$

We can interchange the gradient and expectation on the right hand side of Equation A.2 (see Harrison and van Mieghem (1999) for a proof), and thus Equation A.2 becomes

$$\begin{aligned}\nabla_{y_1,y_2} C^{IIP}(y_1, y_2) = 0 = \mathbb{E}_{\tilde{D}} \nabla_{y_1,y_2} P\left(y_1, y_2, \tilde{D}\right) &= (h, h)^T - \mathbb{E}_{\tilde{D}}\lambda\left(y_1, y_2, \tilde{D}\right) \\ &= (h, h)^T - \sum_k \lambda^k \mathbb{P}\left(\Omega_k\left(y_1, y_2\right)\right)\end{aligned}$$

where λ^k is the constant $\lambda\left(y_1, y_2, \tilde{D}\right)$ for $\tilde{D} \in \Omega_k\left(y_1, y_2\right)$. $\square$

A.1.2 Proof of Proposition 2.5.2

Based on the approximation used to formulate C^{LB}, the difference in costs between C^{IIP} and C^{LB} is:

$$
\begin{aligned}
&C^{IIP}(\mathbf{y}) - C^{LB}(\mathbf{y}) \\
&= (h + p_o - s_{12})\,\mathbb{E}\Big[\Big(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\Big)^+ + \sum_i (D_{is} - y_i)^+ - \Big(D - \sum_i y_i\Big)^+\Big] \\
&\geq (h + p_o - s_{12})\,\mathbb{E}\Big[\Big(\sum_i D_{io} - \sum_i (y_i - D_{is})^+ + \sum_i (D_{is} - y_i)^+\Big)^+ - \Big(D - \sum_i y_i\Big)^+\Big] \\
&= 0
\end{aligned}
$$

The first inequality follows from : $a^+ + b^+ \geq (a+b)^+$, and further simplification uses $x^+ - (-x)^+ = x$. □

The proof follows for any number of stores, as long as the cross-shipping cost is a constant and $s_{12} < h + p_o$. The proof also follows when s_{12} is reduced to s, as done in Equation 2.14.

A.1.3 Proof of Proposition 2.5.3

A similar result is proved in Dong and Rudi (2004, Lemma 1), who consider the case of traditional transshipment. Substituting $\mathbf{y^{DIP}}$ into the first order condition for C^{LB} in Equation 2.12, we have:

$$
\begin{aligned}
(h + po - s_{12})F_D\left(\sum_j y_j^{DIP}\right) + (s_{12} - s)F_{D_i}(y_i^{DIP}) + (p_s - p_o + s)F_{D_{is}}((y_i^{DIP}) - p_s \\
= (h + p_o - s_{12})\left(\Phi\left(z^{DIP}\sum_i \sigma_i/\sigma\right) - \Phi\left(z^{DIP}\right)\right)
\end{aligned}
$$

where Φ is the CDF of the standard normal distribution. The equality follows from the fact that $\mathbf{y^{DIP}}$ satisfies Equation 2.7, and the normality of demands, as we can write $y_i^{DIP} = \mu_i + z^{DIP}\sigma_i$, where $D_i \sim \mathcal{N}(\mu_i, \sigma_i)$, and $D \sim \mathcal{N}(\mu, \sigma)$. As $\sum_i \sigma_i/\sigma \geq 1$, it follow that the gradient of C^{LB} at $\mathbf{y^{DIP}}$ is $\geq 0 (\leq 0)$ whenever $z^{DIP} \leq (\geq) \mu_i$. Also, writing $\sigma = \sqrt{\sum_i \sigma_i^2 + \sum_j 2\rho_l \sigma_i \sigma_j}$, where ρ_l is the correlation coefficient between locations, y^{DIP} is optimal to C^{LB} and C^{IIP} when $\rho_l = 1$. □

A.1.4 Proof of Proposition 2.5.4

Due to similarities to Dong and Rudi (2004), we have a similar solution where the optimal inventory at each location is at the same critical fractile of the location's demands. Equation 2.15 can be written as:

$$y_i^{LBN} = F_{D_{is}}^{-1}\left(\frac{m}{p_s - p_o + s}\right), \quad \forall i \in \mathcal{S}_{so} \tag{A.3}$$

where $m = p_s - (h + p_o - s)F_{D_\mathcal{S}}(\sum_{j\in\mathcal{S}} y_j^{LBN})$. Substituting Equation A.3 into the definition of m, we have:

$$\sum_{j\in\mathcal{S}} F_{D_{is}}^{-1}\left(\frac{m}{p_s - p_o + s}\right) = F_{D_\mathcal{S}}^{-1}\left(\frac{p_s - m}{h + p_o - s}\right)$$

Solving this yields a unique solution for m, which in turn yields a unique solution $\mathbf{y^{LBN}}$, where each stores stocks at the same critical fractile of their in-store demand, as seen from Equation A.3. □

For OFCs ($i \in \mathcal{S}_o$), $y_i^{LBN} = 0$, as otherwise, the value of m is forced to be $p_s - p_o + s$, which renders Equation A.3 to infinity.

A.1.5 Proof of Proposition 2.5.5

Consider a square of unit area in which N stores are uniformly distributed. Let the square be divided into $\sqrt{N}$ identical cells, such that each cell contains $\sqrt{N}$ stores. The dimensions of each cell are thus $\frac{1}{N^{\frac{1}{4}}} \times \frac{1}{N^{\frac{1}{4}}}$. The superscript l for a demand variable (e.g. D_{is}^l) denotes that the demand belongs to a store in cell l.

Let $C^{LB'}$ be the cost function obtained from C^{IIP} by lowering all cross-shipping costs to the within-region shipping cost s. Let C^{IIP_c} and $C^{LB'_c}$ be the functions obtained by restricting C^{IIP} and $C^{LB'}$ respectively, so that cross-shipments can only be made between two stores belonging to the same cell. Clearly, $C^{IIP}(y) \leq C^{IIP_c}(y)$ and $C^{LB'}(y) \leq C^{LB'_c}(y)$ for any $y \geq 0$. Let $g(y, N)$ denote the cost incurred by N stores starting with inventory y each, without the option of cross-shipping:

$$\begin{aligned} g(y, N) = \sum_{i=1}^{N} \Big[& h\,(y - D_i)^+ + p_s\,(D_{is} - y)^+ \\ & + p_o\left(D_{io} - (y - D_{is})^+\right)^+ + s \min\left(D_{io}, (y - D_{is})^+\right) \Big] \end{aligned}$$

Note that $g(y, N)$ represents the sum of costs incurred by individual stores, and hence,

$\mathbb{E}g(y,N) = \mathbb{E}\sum_{l=1}^{\sqrt{N}} g(y,\sqrt{N}) = \sqrt{N}g(y,\sqrt{N})$. Let $CS_{ij}(y,N)$ denote the cross-shipped quantity between stores i and j, when there are N stores with order-up-to quantity y each (CS_{ij}^l when defined within a cell). Note that both the functions g and CS_{ij} also depend on the demand vector, but the dependency is ignored for notational convenience. As the cells are identical in terms of demands and costs, we have:

$$
\begin{aligned}
C^{IIP_c}(y^{IIPH}) = \quad & \mathbb{E}\left(\sum_{l=1}^{\sqrt{N}}\left[g(y^{IIPH},\sqrt{N}) + \sum_{i=1}^{\sqrt{N}}\sum_{j=1,j\neq i}^{\sqrt{N}}(s_{ij}^l - h - p_o)CS_{ij}^l(y^{IIPH},\sqrt{N})\right]\right) \\
= \quad & \mathbb{E}g(y^{IIPH},N) + \mathbb{E}\left(\sum_{l=1}^{\sqrt{N}}\left(\sum_{i=1}^{\sqrt{N}}\sum_{j=1,j\neq i}^{\sqrt{N}}(s_{ij}^l - h - p_o)CS_{ij}^l(y^{IIPH},\sqrt{N})\right)\right) \\
C^{LB'}(y^{IIPH}) = \quad & C^{LB'_c}(y^{IIPH}) \\
& + (s-h-p_o)\mathbb{E}\left[\sum_{l=1}^{\sqrt{N}}\left(\sum_{i=1}^{\sqrt{N}} D_{io}^l - \left(y^{IIPH} - D_{is}^l\right)^+\right)^+ - \left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^+\right)^+\right] \\
= \quad & \mathbb{E}g(y^{IIPH},N) + \mathbb{E}\left(\sum_{l=1}^{\sqrt{N}}\left(\sum_{i=1}^{\sqrt{N}}\sum_{j=1,j\neq i}^{\sqrt{N}}(s - h - p_o)CS_{ij}^l(y^{IIPH},\sqrt{N})\right)\right) \\
& + (s-h-p_o)\left[\sqrt{N}\mathbb{E}\left(\sum_{i=1}^{\sqrt{N}} D_{io}^l - \left(y^{IIPH} - D_{is}^l\right)^+\right)^+ - \mathbb{E}\left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^+\right)^+\right]
\end{aligned}
$$

The expression for $C^{LB'}$ is written as the sum of $C^{LB'_c}$ which restricts cross-shipping to within each cell, and the cost of the additional cross-shipped units with this restriction removed. We know that $C^{LB}(y^{IIPH}) \le C^{LB'}(y^{IIPH}) \le C^{IIP}(y^{IIPH}) \le C^{IIP_c}(y)$, where the first inequality follows from Proposition 2.5.5. We first show that $\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} \to 1$ as $N \to \infty$. We have:

$$
\begin{aligned}
& \frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} - 1 \\
&= \frac{\mathbb{E}\left(\sum_{l=1}^{\sqrt{N}}\left(\sum_{i=1}^{\sqrt{N}}\sum_{j=1,j\neq i}^{\sqrt{N}}(s_{ij}^l - s)CS_{ij}^l(y^{IIPH},\sqrt{N})\right)\right)}{C^{LB'}(y^{IIPH})} \quad + \\
& \frac{(h+p_o-s)\left[\sqrt{N}\mathbb{E}\left(\sum_{i=1}^{\sqrt{N}} D_{io}^l - \left(y^{IIPH} - D_{is}^l\right)^+\right)^+ - \mathbb{E}\left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^+\right)^+\right]}{C^{LB'}(y^{IIPH})}
\end{aligned}
$$

We have $s_{ij}^l - s = f(d_{ij}^l) \le f\left(\frac{\sqrt{2}}{N^{\frac{1}{4}}}\right)$, as the maximum distance within a cell is $\frac{\sqrt{2}}{N^{\frac{1}{4}}}$. Thus, using $C^{LB'}(y^{IIPH}) \geq \mathbb{E}\left(\sum\limits_{l=1}^{\sqrt{N}}\left(\sum\limits_{i=1}^{\sqrt{N}}\sum\limits_{j=1,j\neq i}^{\sqrt{N}}(s)CS_{ij}^{l}(y^{IIPH},\sqrt{N})\right)\right)$ for the first term, and $C^{LB'}(y^{IIPH}) \geq s\mu_o N$ for the second term, we have

$$\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} - 1 \leq \frac{f\left(\frac{\sqrt{2}}{N^{\frac{1}{4}}}\right)}{s} + \left(\frac{h+p_o-s}{s\mu_o\sqrt{N}}\right)\mathbb{E}\left(\sum_{i=1}^{\sqrt{N}} D_{io} - \left(y^{IIPH} - D_{is}\right)^{+}\right)^{+} \tag{A.4}$$

The first term on the right hand side vanishes to zero as $N \to \infty$, as $f(d) \to 0$ as $d \to 0$. To simplify the second term, we need the following lemmas.

Lemma A.1.1 *If $h < p_o - s$, then $y^{IIPH} > \mu$ where $\mu = \mu_s + \mu_o$, and if additionally $h < (p_s - p_o + s)F_s(\mu)$,*

$$y^{IIPH} \to F_s^{-1}\left(\frac{p_s - p_o + s - h}{p_s - p_o + s}\right) \in (0, \infty), \text{ as } N \to \infty$$

Proof: Lemma 1 is proved from the optimality equations of C^{LBN} (Equation 2.15) for identical stores:

$$(h + p_o - s)\mathbb{P}\left(\sum_{i=1}^{N} D_i \leq N y^{IIPH}\right) + (p_s - p_o + s)F_{D_{1s}}(y^{IIPH}) = p_s$$

From the above equation, when $h < p_o - s$, we have

$$p_s < 2\,(p_o - s)\,\mathbb{P}\left(\sum_{i=1}^{N} D_i \leq N y^{IIPH}\right) + (p_s - p_o + s)$$

This simplifies to yield $y^{IIPH} > \mu$. Now, by applying the central limit theorem as $N \to \infty$ and $y^{IIPH} > \mu$, $\mathbb{P}\Big(\sum\limits_{i=1}^{N} D_i/N \leq y^{IIPH}\Big) \to 1$, and the result follows. Note that the asymptotic solution should also satisfy $y^{IIPH} > \mu$, which translates to the condition $h < (p_s - p_o + s)F_s(\mu)$. □

Lemma A.1.2 *When $h < p_o - s$ and $h < p_s - (p_o - s)$, and the demands are bounded above as $D_{is} \leq M_s$ and $D_{io} \leq M_o$ for all i,*

$$\mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right) \leq \exp\left\{\frac{-2\sqrt{N}(y^{IIPH} - \mu)^2}{M_o + M_s}\right\}$$

Proof:

$$
\begin{aligned}
&\mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right) \\
&= \mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} \left(D_i - \left(D_{is} - y^{IIPH}\right)^{+}\right) > \sqrt{N} y^{IIPH}\right) \leq \mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_i > \sqrt{N} y^{IIPH}\right) \\
&\leq \exp\left\{\frac{-2\sqrt{N}(y^{IIPH} - \mu)^2}{M_o + M_s}\right\} \to 0, \ as\ N \to \infty
\end{aligned}
$$

The final inequality follows from the Hoeffding bound for tail probabilities Hoeffding (1963), as $y^{IIPH} > \mu$ and demands are bounded, and the limit exists as y^{IIPH} approaches a finite positive quantity as $N \to \infty$ by Lemma 1. The expectation in the second term of Equation A.4 can be bounded as follows:

$$
\begin{aligned}
&\mathbb{E}\left(\sum_{i=1}^{\sqrt{N}} \left(D_{io} - \left(y^{IIPH} - D_{is}\right)^{+}\right)\right)^{+} \\
&= \mathbb{E}\left[\left(\sum_{i=1}^{\sqrt{N}} \left(D_{io} - \left(y^{IIPH} - D_{is}\right)^{+}\right)\right)^{+} \middle| \sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right] \mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right) \\
&\leq \mathbb{E}\left[\sum_{i=1}^{\sqrt{N}} D_{io} \middle| \sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right] \mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^{+}\right) \\
&\leq M_o\sqrt{N} \exp\left\{\frac{-2\sqrt{N}(y^{IIPH} - \mu)^2}{M_o + M_s}\right\}
\end{aligned}
$$

The last inequality follows from Lemma 2 and the boundedness of the demands as $D_{is} \leq M_s$, and $D_{io} \leq M_o$ for all i with $0 < M_s, M_o < \infty$. $\square$

Thus, we have:

$$
\begin{aligned}
\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} &\leq 1 + \frac{f\left(\frac{\sqrt{2}}{N^{\frac{1}{4}}}\right)}{s} + \left(\frac{h + p_o - s}{s\mu_o}\right)\left(M_o\sqrt{N} \exp\left\{\frac{-2\sqrt{N}(y^{IIPH} - \mu)^2}{M_o + M_s}\right\}\right) \\
&\to 1, \ as\ N \to \infty
\end{aligned} \tag{A.5}
$$

The next step is to show the C^{LBN} is off by a constant factor from the $C^{LB'}$. From the proof of Proposition 2.5.2, the difference simplifies to:

$$
\begin{aligned}
&C^{LB'}(y^{IIPH}) - C^{LBN}(y^{IIPH}) \\
&= (h + p_o - s)\,\mathbb{E}\Bigg[\left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^{+} \right)^{+} \\
&\qquad + \sum_{i=1}^{N} \left(D_{is} - y^{IIPH}\right)^{+} - \left(D - \sum_{i=1}^{N} y^{IIPH} \right)^{+} \Bigg]
\end{aligned}
$$

where $D = \sum_{i=1}^{N} D_{is} + D_{io}$.

Similar to what was done to bound the second term in Equation A.4, we can show that whenever the conditions in Lemma 2 are satisfied

$$
\mathbb{E}\left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^{+} \right)^{+} \leq M_o N \exp\left\{ \frac{-2N(y^{IIPH} - \mu)^2}{M_o + M_s} \right\}
$$

Thus, we have:

$$
\begin{aligned}
&C^{LB'}(y^{IIPH}) - C^{LBN}(y^{IIPH}) \\
&\leq (h + p_o - s) \left[M_o N \exp\left\{ \frac{-2N(y^{IIPH} - \mu)^2}{M_o + M_s} \right\} + \sum_{i=1}^{N} \left(D_{is} - y^{IIPH}\right)^{+} \right]
\end{aligned}
$$

Using $C^{LBN}(y^{IIPH}) \geq s\mu_o N$ and $C^{LBN}(y^{IIPH}) \geq (p_s - p_o + s) \sum_{i=1}^{N} \left(D_{is} - y^{IIPH}\right)^{+}$, we have:

$$
\frac{C^{LB'}(y^{IIPH})}{C^{LBN}(y^{IIPH})} - 1 \leq \left(\frac{h + p_o - s}{s\mu_o} \right) \left(M_o \exp\left\{ \frac{-2N(y^{IIPH} - \mu)^2}{M_o + M_s} \right\} \right) + \left(\frac{h + p_o - s}{p_s - p_o + s} \right) \tag{A.6}
$$

Thus, from Equations A.5 and A.6, as $N \to \infty$, we have

$$
\begin{aligned}
&\frac{C^{IIP_c}(y^{IIPH})}{C^{LBN}(y^{IIPH})} \leq 1 + \frac{h + p_o - s}{p_s - p_o + s} \\
&\Rightarrow \frac{C^{IIPH}}{C^{LBN}(y^{IIPH})} \leq \frac{h + p_s}{p_s - p_o + s}
\end{aligned}
$$

The final step follows from $C^{IIP_c}(y^{IIPH}) \geq C^{IIP}(y^{IIPH}) = C^{IIPH}$. □

The result may hold subject to some generalizations, such as the unit square can be replaced with any finite area, and non-identical cells as long as the number of stores in each cell grows to infinity as $N \to \infty$. The resulting cases may call for a more complicated proof, and is outside the scope of this study.

A.1.6 Proof of Proposition 2.6.1

Consider the single period case, where items are ordered at the start of the period, and online demands are fulfilled over T fulfillment epochs. Assume that $C_{T+1}(\mathbf{x^{T+1}}, \tilde{D}^{T+1}) = 0$ without loss of generality. Thus, $C_T(\mathbf{x^T}, \tilde{D}^T)$ is given by a simple linear program which is jointly convex in $(\mathbf{x^T}, \tilde{D}^T)$. This leads to the base case result that $C_T(\mathbf{x^T}, \tilde{D}^T)$ is convex in x^T given any $\tilde{D}^T$. By backward induction, we need to show that $C_t(\mathbf{x^t}, \tilde{D}^t)$ is convex in $\mathbf{x^t}$ for any given $\tilde{D}^t$, with the assumption that $C_{t+1}(\mathbf{x^{t+1}}, \tilde{D}^{t+1})$ is convex in $\mathbf{x^{t+1}}$ given any $\tilde{D}^{t+1}$. The cost-to-go function can be represented by $C_t(\mathbf{x^t}, \tilde{D}^t) = \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \mathcal{G}(\mathbf{x^t}, \tilde{D}^t, \mathbf{z^t}, \mathbf{Z^t})$, where

$$\mathcal{G}(\mathbf{x^t}, \tilde{D}^t, \mathbf{z^t}, \mathbf{Z^t}) = \Big[P(\mathbf{x^t}, \tilde{D}^t, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E}C_{t+1}(x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t, \tilde{D}^{t+1})\Big] \tag{A.7}$$

Consider any $\mu \geq 0$, and $\mathbf{x_1^t}, \mathbf{x_2^t} \geq 0$. Let $(\mathbf{z_i^t}, \mathbf{Z_i^t}) = \arg\min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \mathcal{G}(\mathbf{x_i^t}, \tilde{D}^t, \mathbf{z^t}, \mathbf{Z^t})$. Note that P is a linear function in its variables (Equation 2.2), and $\mathbb{E}C_{t+1}(\mathbf{x^{t+1}}, \tilde{D}^{t+1})$ is convex in $\mathbf{x^{t+1}}$, as expectation preserves convexity. Let $\bar{\mathbf{x}}^\mathbf{t} = \mu\mathbf{x_1^t} + (1-\mu)\mathbf{x_2^t}$, $\bar{\mathbf{z}}^\mathbf{t} = \mu\mathbf{z_1^t} + (1-\mu)\mathbf{z_2^t}$ and $\bar{\mathbf{Z}}^\mathbf{t} = \mu\mathbf{Z_1^t} + (1-\mu)\mathbf{Z_2^t}$. We have:

$$\begin{aligned}
&C_t(\bar{\mathbf{x}}^\mathbf{t}, \tilde{D}^t) \\
&= \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \Big[P(\bar{\mathbf{x}}^\mathbf{t}, \tilde{D}^t, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t, \tilde{D}^{t+1})\Big] \\
&\leq P(\bar{\mathbf{x}}^\mathbf{t}, \tilde{D}^t, \bar{\mathbf{z}}^\mathbf{t}, \bar{\mathbf{Z}}^\mathbf{t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{D}^{t+1}) \\
&\leq \mu P(\mathbf{x_1^t}, \tilde{D}^t, \mathbf{z_1^t}, \mathbf{Z_1^t}) + (1-\mu)P(\mathbf{x_2^t}, \tilde{D}^t, \mathbf{z_2^t}, \mathbf{Z_2^t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{D}^{t+1})
\end{aligned}$$

The first inequality follows from the feasibility of $\bar{\mathbf{z}}^\mathbf{t}, \bar{\mathbf{Z}}^\mathbf{t}$ in Δ, as $(\mathbf{z_1^t}, \mathbf{Z_1^t})$ and $\mathbf{z_2^t}, \mathbf{Z_2^t})$ are feasible in Δ. The second inequality follows from the convexity of P.

As $\mathbb{E}C_{t+1}(\mathbf{x^{t+1}}, \tilde{D}^{t+1})$ is convex in $\mathbf{x^{t+1}}$, we have:

$$
\begin{aligned}
&\mathbb{E}C_{t+1}\big(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{D}^{t+1}\big) \\
&= \mathbb{E}C_{t+1}\left[\mu\left(x_{1,i}^t - z_{1,i}^t - \sum_{j=1}^{N} Z_{1,ij}^t\right) + (1-\mu)\left(x_{2,i}^t - z_{2,i}^t - \sum_{j=1}^{N} Z_{2,i}^t\right), \tilde{D}^{t+1}\right] \\
&\leq \mu\mathbb{E}C_{t+1}\left[x_{1,i}^t - z_{1,i}^t - \sum_{j=1}^{N} Z_{1,ij}^t, \tilde{D}^{t+1}\right] + (1-\mu)\mathbb{E}C_{t+1}\left[x_{2,i}^t - z_{2,i}^t - \sum_{j=1}^{N} Z_{2,i}^t, \tilde{D}^{t+1}\right]
\end{aligned}
$$

Thus, from Equation A.7, we have:

$$
\begin{aligned}
C_t(\bar{\mathbf{x}}^{\mathbf{t}}, \tilde{D}^t) &\leq \mu\mathcal{G}(\mathbf{x_1^t}, \tilde{D}^t, \mathbf{z_1^t}, \mathbf{Z_1^t}) + (1-\mu)\mathcal{G}(\mathbf{x_2^t}, \tilde{D}^t, \mathbf{z_2^t}, \mathbf{Z_2^t}) \\
&= \mu C_t(\mathbf{x_1^t}, \tilde{D}^t) + (1-\mu)C_t(\mathbf{x_2^t}, \tilde{D}^t)
\end{aligned}
$$

The equality follows from the definitions of $(\mathbf{z_1^t}, \mathbf{Z_1^t})$ and $(\mathbf{z_2^t}, \mathbf{Z_2^t})$. □

A.1.7 Proof of Proposition 2.6.2

Let the single period cost function be given by $C^{IIP}(\mathbf{y}) = \mathbb{E}C_1(\mathbf{y}, \tilde{D})$, and let $\mathbf{y}^{IIP}$ be the optimal solution. When the initial level of inventory x_i at region i before ordering, the cost function is as follows:

$$V^{IIP}(\mathbf{x}) = \min_{\mathbf{y}\geq\mathbf{x}} C^{IIP}(\mathbf{y}) = C^{IIP}(\mathbf{y}^{IIP})$$

As $\mathbf{y}^{IIP}$ minimizes the function C^{IIP}, for any $\{\mathbf{x} : \mathbf{x} \leq \mathbf{y}^{IIP}\}$, it is optimal to order up to $\mathbf{y}^{IIP}$. We ignore cases where $x_i > y_i^{IIP}$ for some i, as eventually the system is brought to the state $\mathbf{x} \leq \mathbf{y}^{IIP}$.

For the multiple period case, we have M time periods: $m = 1, 2, .., M$. The in-store demands $\{D_{is}^m, m > 0\}$ and online demands $\{D_{io}^m, m > 0\}$ are assumed to be i.i.d. The available inventory at the end of a review period serves as the initial inventory for the next review period, and we assume zero purchasing costs. The discount factor is $\delta \in (0, 1]$.

The proof is by induction, and similar to the proof of Proposition 4 in van Mieghem and Rudi (2002). If we show that $V_m^{IIP}(\mathbf{x}^m)$, the expected cost-to-go function evaluated in review period m with the initial inventory $\mathbf{x}^m$, is convex and affine, a stationary base stock policy would be optimal. For the $M + 1^{th}$ review period, the cost function is $V_{M+1}^{IIP}(\mathbf{x}^{M+1}) = 0$ (assuming zero purchasing costs) which is trivially convex and affine in $\mathbf{x}^{M+1}$. Let V_{m+1}^{IIP} be convex and affine in $\mathbf{x}^{m+1}$. The cost function for review period m is:

$$V_m^{IIP}(\mathbf{x}) = \min_{\mathbf{y}\geq\mathbf{x}} \left[C^{IIP}(\mathbf{y}) + \delta\mathbb{E}V_{m+1}^{IIP}\left(f\left(\mathbf{y},\tilde{D}\right)\right)\right] = \min_{\mathbf{y}\geq\mathbf{x}} U_m^{IIP}(\mathbf{y})$$

where $f(\mathbf{y}, D)$ is the vector of ending inventories. $\tilde{D}$ is the demand vector constituting the in-store and online demands for both the regions. As taking expectation preserves convexity, and the sum of convex functions is convex, $U_m^{IIP}(\mathbf{y})$ is convex in $\mathbf{y}$. It only remains to be shown that V_m^{IIP} is affine in $\mathbf{x}$. To show this, consider any $\mathbf{y} \leq \mathbf{y}^{IIP}$, so that $f\left(\mathbf{y},\tilde{D}\right) \leq \mathbf{y} \leq \mathbf{y}^{IIP}$. We have

$$U_m^{IIP}(\mathbf{y}) = C^{IIP}(\mathbf{y}) + \delta\mathbb{E}V_{m+1}^{IIP}\left(f\left(\mathbf{y},\tilde{D}\right)\right) = C^{IIP}(\mathbf{y}) + \delta\mathbb{E}V_{m+1}^{IIP}\left(\mathbf{y}^{IIP}\right)$$

as V_{m+1}^{IIP} is affine in $\mathbf{x}^{m+1}$ and the purchasing cost is zero. Clearly, $\mathbf{y} = \mathbf{y}^{IIP}$ minimizes U_m^{IIP} for $\mathbf{y} \leq \mathbf{y}^{IIP}$. Thus, $V_m^{IIP}(\mathbf{x}) = \max_{\mathbf{y}\geq\mathbf{x}} U_m^{IIP}(\mathbf{y})$ is affine (constant) in $\mathbf{x}$ for all $\mathbf{x} \leq \mathbf{y}^{IIP}$, and hence a stationary base-stock policy $\mathbf{y}^{IIP}$ is optimal if $\mathbf{x} \leq \mathbf{y}^{IIP}$. If there is some i for which $x_i > y_i^{IIP}$, the optimal policy will be more complicated, but eventually, the system comes back to $\mathbf{x} \leq \mathbf{y}^{IIP}$. □

A.2 Demand Regions for the IIP Solution

We illustrate the identification of demand regions in which the dual vector λ is constant (as discussed in Section 2.5.1.2) and the calculation of the corresponding probabilities. For any given (y_1, y_2), the demand space $(D_{1s}, D_{1o}, D_{2s}, D_{2o})$ can be divided into a number of independent regions. Based on the values taken by the variables in the optimal solution in (A.1), Table A.1 shows the different cases that are possible given y_1 and y_2. From these cases, the independent demand regions are listed in Table A.2 along with the constant dual prices in those regions. The underlined cases are redundant, and can be discarded while calculating the probability for each region.

Table A.1: Table showing the various demand cases based on the values of y_1, y_2

	A	B	C	D
1	$y_1 < D_{1s}$	$D_{1s} \leq y_1 < D_1$	$D_1 \leq y_1 < D_1 + D_{2o}$	$y_1 \geq D_1 + D_{2o}$
2	$y_2 < D_{2s}$	$D_{2s} \leq y_2 < D_2$	$D_2 \leq y_2 < D_2 + D_{1o}$	$y_2 \geq D_2 + D_{1o}$
3	$y_1 + y_2 < D_1 + D_2$	$y_1 + y_2 \geq D_1 + D_2$		

The dual prices λ_1, λ_2 are the shadow prices of the constraints which contain y_1 and y_2

Table A.2: Table showing the various demand regions and the corresponding constant dual-prices.

Region	Case	λ_1	λ_2	Region	Case	λ_1	λ_2
Ω_1	A1,A2,$\underline{\text{A3}}$	$h+p_s$	$h+p_s$	Ω_{11}	C1,A2,$\underline{\text{A3}}$	$h+p_o-s_{12}$	$h+p_s$
Ω_2	A1,B2,$\underline{\text{A3}}$	$h+p_s$	$h+p_o-s$	Ω_{12}	C1,B2,A3	$h+p_o-s_{12}$	$h+p_o-s$
Ω_3	A1,C2,$\underline{\text{A3}}$	$h+p_s$	$h+p_o-s_{12}$	Ω_{13}	C1,B2,B3	0	$s_{12}-s$
Ω_4	$\underline{\text{A1}}$,D2,A3	$h+p_s$	0	Ω_{14}	C1,C2,$\underline{\text{B3}}$	0	0
Ω_5	A1,$\underline{\text{D2}}$,B3	$h+p_s$	0	Ω_{15}	C1,D2,$\underline{\text{B3}}$	0	0
Ω_6	B1,A2,$\underline{\text{A3}}$	$h+p_o-s$	$h+p_s$	Ω_{16}	D1,$\underline{\text{A2}}$,A3	0	$h+p_s$
Ω_7	B1,B2,$\underline{\text{A3}}$	$h+p_o-s$	$h+p_o-s$	Ω_{17}	$\underline{\text{D1}}$,A2,B3	0	$h+p_s$
Ω_8	B1,C2,A3	$h+p_o-s$	$h+p_o-s_{12}$	Ω_{18}	D1,B2,$\underline{\text{B3}}$	0	$s_{12}-s$
Ω_9	B1,C2,B3	$s_{12}-s$	0	Ω_{19}	D1,C2,$\underline{\text{B3}}$	0	0
Ω_{10}	B1,D2,$\underline{\text{B3}}$	$s_{12}-s$	0	Ω_{20}	D1,D2,$\underline{\text{B3}}$	0	0

respectively, namely the first set of constraints $z_i + \sum_{j=1}^{2} z_{ij} \leq y_i, \forall i$ in the linear program in (A.1), and can be obtain in a standard fashion from linear programming theory. For example, for the demand regions with the case D1, that is, $y_1 \geq D_1 + D_{2o}$, irrespective of the value of y_2, there will be inventory left over at retail store 1 at the end of the period. Thus the constraint $z_1 + \sum_{j=1}^{2} z_{1j} \leq y_1$ will not bind, and hence $\lambda_1 = 0$.

The probability for each region is calculated as follows, when demands follow normal distributions. The region is expressed as an inequality of the form $R_k\tilde{D} <= S_kY$, where $\tilde{D} = [D_{1s}, D_{1o}, D_{2s}, D_{2o}]^\intercal$ and $Y = [y_1, y_2]^\intercal$. For example, $\Omega_3 = (A1, C2) = \{y_1 < D_{1s}, D_2 \leq y_2 < D_2 + D_{1o}\}$. This can be expressed as:

$$\begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & -1 & -1 & -1 \end{bmatrix} \begin{bmatrix} D_{1s} \\ D_{1o} \\ D_{2s} \\ D_{2o} \end{bmatrix} \leq \begin{bmatrix} -1 & 0 \\ 0 & 1 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$$

$R_k\tilde{D}$ is multivariate normal with mean $R_k\mu$ and covariance matrix $R_k\Sigma\Sigma^\intercal R_k^\intercal$, where μ and Σ are the mean and covariance matrices of $\tilde{D}$. The probability of region k reduces to evaluating the cumulative distribution function of $A_k\tilde{D}$ at B_kY. For general demand distributions, numerical methods have to be employed.

A.3 Additional Details for Numerical Analyses

All numerical analyses were done on a desktop computer (i7-3770 CPU @3.7GHz, 16GB RAM). The total market is assumed to be the top 300 most populous cities in mainland US. The demands for the OFCs are calculated based on the population not covered by omnichannel stores. This online demand is allocated to each OFC based on the optimal throughput rates estimated by Chicago Consulting (2013).

A.3.1 Simulation Procedure

A brief overview of the simulation is listed below:

1. The parameters for demands in each fulfillment epoch are calculated based on review-period demands estimated from population data. The starting inventory level vectors $\mathbf{y}^{\mathbf{DIP}}$ and $\mathbf{y}^{\mathbf{IIPH}}$ are calculated using the demand information based on Equation 2.7 and Algorithm 1 respectively.

2. We generate a sample of size 10^4, where each sample is a realization of demands in a review period, although fulfillment decisions in each fulfillment epoch are made without knowing future demands. For each sample, we iterate over steps 3-7, and take the sample averages as approximations for expectations.

3. The fulfillment thresholds for the TF policy are calculated based on Equation 2.18. For the MF policy, these thresholds are set to zero.

4. For $t = 1, \ldots, T$, iterate over steps 5-6. The starting inventory levels are set based on the inventory policy followed (IIPH or DIP).

5. Implement Algorithm 2 based on the fulfillment policy followed (MF or TF) and the corresponding thresholds calculated in Step 3.

6. At the end of each fulfillment epoch, the holding, penalty and fulfillment costs are calculated. The ending inventory at a location becomes the starting inventory for the next epoch.

7. The total cost in a review period is the sum of the costs in each fulfillment epoch in that period.

APPENDIX B

Appendix to Chapter 3

B.1 Proofs

B.1.1 Proof of Lemma 3.4.2

The identical two location problem is described as:

$$\begin{aligned}
C^* &= \min_{y_1,y_2\geq 0} \sup_{f\in\mathcal{F}^{m\sigma\rho}} \mathbb{E}_f\left[C(y_1,y_2,\tilde{\mathbf{D}})\right] \\
&:= \min_{y_1,y_2\geq 0} G(y_1,y_2)
\end{aligned}$$

It is easy to see that $\mathbb{E}_f\left[C(y_1,y_2,\tilde{\mathbf{D}})\right]$ is jointly convex in y_1, y_2, as C can be expressed as a linear program, and expectation preserves convexity. Note that C is also symmetric with respect to y_1 and y_2 when the locations are identical. Thus, we have $G(y_1,y_2) = G(y_2,y_1)$. If we show that G is jointly convex in y_1, y_2, we are done, because if (y_1^*, y_2^*) is an optimal solution to C^*, then so is (y_2^*, y_1^*), and so is $\left(\frac{y_1^*+y_2^*}{2}, \frac{y_1^*+y_2^*}{2}\right)$.

To show joint convexity of G, consider two points: $(\hat{y}_1,\hat{y}_2)$ and $(\bar{y}_1,\bar{y}_2)$. Let $\lambda \in [0,1]$. We have:

$$\begin{aligned}
&G\left(\lambda\hat{y}_1+(1-\lambda)\bar{y}_1, \lambda\hat{y}_2+(1-\lambda)\bar{y}_2\right) \\
&\quad= \sup_{f\in\mathcal{F}^{m\sigma\rho}} \mathbb{E}_f\left[C\left(\lambda\hat{y}_1+(1-\lambda)\bar{y}_1, \lambda\hat{y}_2+(1-\lambda)\bar{y}_2, \tilde{\mathbf{D}}\right)\right] \\
&\quad\leq \sup_{f\in\mathcal{F}^{m\sigma\rho}} \left(\lambda\cdot\mathbb{E}_f\left[C(\hat{y}_1,\hat{y}_2,\tilde{\mathbf{D}})\right] + (1-\lambda)\cdot\mathbb{E}_f\left[C(\bar{y}_1,\bar{y}_2,\tilde{\mathbf{D}})\right]\right) \\
&\quad= \lambda\cdot\mathbb{E}_{f^*}\left[C(\hat{y}_1,\hat{y}_2,\tilde{\mathbf{D}})\right] + (1-\lambda)\cdot\mathbb{E}_{f^*}\left[C(\bar{y}_1,\bar{y}_2,\tilde{\mathbf{D}})\right] \\
&\quad\leq \lambda\cdot\sup_{f\in\mathcal{F}^{m\sigma\rho}}\left(\mathbb{E}_f\left[C(\hat{y}_1,\hat{y}_2,\tilde{\mathbf{D}})\right]\right) + (1-\lambda)\cdot\sup_{f\in\mathcal{F}^{m\sigma\rho}}\left(\mathbb{E}_f\left[C(\bar{y}_1,\bar{y}_2,\tilde{\mathbf{D}})\right]\right) \\
&\quad= \lambda G(\hat{y}_1,\hat{y}_2) + (1-\lambda)G(\bar{y}_1,\bar{y}_2)
\end{aligned}$$

where $f^* = \arg\sup_{f\in\mathcal{F}^{m\sigma\rho}} \left(\lambda \cdot \mathbb{E}_f \left[C(\hat{y}_1, \hat{y}_2, \tilde{\mathbf{D}})\right] + (1-\lambda) \cdot \mathbb{E}_f \left[C(\bar{y}_1, \bar{y}_2, \tilde{\mathbf{D}})\right]\right)$. The first inequality follows from joint convexity of $\mathbb{E}_f \left[C(y_1, y_2, \tilde{\mathbf{D}})\right]$. □

B.1.2 Proof of Proposition 3.4.1

Define

$$\begin{aligned} M(y) \ := \quad & \inf_f \ \mathbb{E}_f \left[\zeta \min\left(\tilde{d}_1 + \tilde{d}_2, 2y\right) + \sum_{j=1,2} \min\left(\tilde{d}_j, y\right)\right] \\ \text{s.t.} \quad & \mathbb{E}_f(1) = 1, \\ & \mathbb{E}_f(\tilde{d}_j) = m, \qquad j = 1, 2, \\ & \mathbb{E}_f(\tilde{d}_j^2) = m^2 + \sigma^2, \qquad j = 1, 2, \\ & \mathbb{E}_f(\tilde{d}_1 \tilde{d}_2) = m^2 + \rho\sigma^2, \\ & f(\mathbf{D}) \geq 0, \qquad \forall \mathbf{D} \in \Re^2. \end{aligned} \tag{B.1}$$

Using the relation $(a-b)^+ = a - \min(a, b)$, we observe that the left-hand side of (3.8) is equivalent to

$$(\zeta + 1)(2m) - \mathbb{E}_f \left[\zeta \min\left(\tilde{d}_1 + \tilde{d}_2, 2y\right) + \sum_{j=1,2} \min\left(\tilde{d}_j, y\right)\right].$$

Therefore, to prove the proposition, it suffices to show that

$$M(y) = (\zeta + 1)\left(y + m - \sqrt{(y-m)^2 + \gamma\sigma^2}\right),$$

and that the distribution that solves (B.1) has no more than six support points.

The dual of the semi-infinite linear program (B.1) is as follows:

$$\begin{aligned} \sup_{t, u_1, u_2, r_1, r_2, v} \quad & t + m(r_1 + r_2) + (m^2 + \sigma^2)(u_1 + u_2) + (m^2 + \rho\sigma^2)v \\ \text{s.t.} \quad & t + r_1 d_1 + r_2 d_2 + u_1 d_1^2 + u_2 d_2^2 + v d_1 d_2 \\ & \quad \leq \zeta \min(d_1 + d_2, 2y) + \min(d_1, y) + \min(d_2, y), \qquad \forall (d_1, d_2) \in \Re^2. \end{aligned}$$

A result by Smith (1995) is that strong duality holds for moment problems if the moment vector is an interior point of the set of feasible moments. For $\mathcal{F}^{m\sigma\rho}$, this is true for $\sigma > 0$ and $\rho \in (-1, 1)$.

Note that because $\tilde{d}_1$ and $\tilde{d}_2$ are interchangeable in the primal, r_1 and r_2 must be interchangeable in the dual. The same argument applies for u_1 and u_2 as well. This implies, $r_1 = r_2 = r$, and $u_1 = u_2 = u$. Thus, we have the following dual formulation:

$$\begin{aligned}
\sup_{t,u,r,v} \quad & t + 2mr + 2(m^2+\sigma^2)u + (m^2+\rho\sigma^2)v \\
\text{s.t.} \quad & t + r(d_1+d_2) + u(d_1^2+d_2^2) + vd_1d_2 \\
& \quad \le \zeta \min(d_1+d_2, 2y) + \min(d_1, y) + \min(d_2, y), \quad \forall (d_1, d_2) \in \Re^2.
\end{aligned}$$

The right hand side of the constraint is a piecewise linear function in $\Re^2$. For notational brevity, define the quadratic function $g(d_1, d_2; t, u, r, v) = t + r(d_1 + d_2) + u(d_1^2 + d_2^2) + vd_1d_2$. Hence, the dual formulation can be equivalently reformulated as

$$\begin{aligned}
\sup_{t,u,r,v} \quad & t + 2mr + 2(m^2+\sigma^2)u + (m^2+\rho\sigma^2)v \\
\text{s.t.} \quad & g(d_1,d_2;t,u,r,v) \le (\zeta+1)(d_1+d_2), && \forall d_1 \le y,\ d_2 \le y \\
& g(d_1,d_2;t,u,r,v) \le \zeta(d_1+d_2) + d_1 + y, && \forall d_1 \le y \le d_2,\ d_1+d_2 \le 2y \\
& g(d_1,d_2;t,u,r,v) \le \zeta(d_1+d_2) + y + d_2, && \forall d_2 \le y \le d_1,\ d_1+d_2 \le 2y \\
& g(d_1,d_2;t,u,r,v) \le (\zeta+1)(2y), && \forall d_1 \ge y,\ d_2 \ge y \\
& g(d_1,d_2;t,u,r,v) \le \zeta(2y) + d_1 + y, && \forall d_1 \le y \le d_2,\ d_1+d_2 \ge 2y \\
& g(d_1,d_2;t,u,r,v) \le \zeta(2y) + y + d_2, && \forall d_2 \le y \le d_1,\ d_1+d_2 \ge 2y.
\end{aligned} \tag{B.2}$$

Note that the dual feasible set is the set of all bi-quadratic functions $g(x_1, x_2)$ that are bounded above by a piecewise linear function with six facets (one for each constraint). Let $q_i(x_1, x_2)$ denote the linear function for facet i, i.e., the right hand side of the constraint i in model (B.2).

Let us consider the case where $g(d_1, d_2)$ touches the piecewise linear function at exactly 6 points, one on each facet. We will later show that this case corresponds to the dual optimal solution. To find these points, for each i, we equate $\nabla g(d_1, d_2) = \nabla q_i(d_1, d_2)$ and solve for (d_1^*, d_2^*) as a function of the dual variables t, u, r, v. Then, setting $g(d_1^*, d_2^*) = q_i(d_1^*, d_2^*)$ gives us a condition on the dual variables for which the two functions touch at exactly one point. We for now ignore the ranges of d_1, d_2 in which each constraint is valid (we will later use these ranges to establish constraints on the dual variables). Table B.1 gives, for each facet, the points of contact and the condition on dual variables t, u, r, v. Note that we have the following four equations that need to be satisfied for $g(d_1, d_2)$ to touch all six facets of the piecewise linear

Facet i	$\nabla g(d_1^*, d_2^*) = \nabla f_i(d_1^*, d_2^*)$	$g(d_1^*, d_2^*) = f_i(d_1^*, d_2^*)$
1	$(d_1^*, d_2^*) = \left(\frac{\zeta+1-r}{2u+v}, \frac{\zeta+1-r}{2u+v}\right)$	$t = \frac{(\zeta+1-r)^2}{2u+v}$
2	$(d_1^*, d_2^*) = \left(\frac{\zeta+1-r}{2u+v} + \frac{v}{4u^2-v^2}, \frac{\zeta+1-r}{2u+v} - \frac{2u}{4u^2-v^2}\right)$	$y(4u^2 - v^2) + u - (\zeta + 1 - r)(2u - v) = 0$
3	$(d_1^*, d_2^*) = \left(\frac{\zeta+1-r}{2u+v} - \frac{2u}{4u^2-v^2}, \frac{\zeta+1-r}{2u+v} + \frac{v}{4u^2-v^2}\right)$	$y(4u^2 - v^2) + u - (\zeta + 1 - r)(2u - v) = 0$
4	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v}, \frac{-r}{2u+v}\right)$	$t = \frac{r^2}{2u+v} + 2(\zeta + 1)y$
5	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v} + \frac{2u}{4u^2-v^2}, \frac{-r}{2u+v} - \frac{v}{4u^2-v^2}\right)$	$t = \frac{r(r-1)(2u-v)+u}{4u^2-v^2} + \zeta(2y) + y$
6	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v} - \frac{v}{4u^2-v^2}, \frac{-r}{2u+v} + \frac{2u}{4u^2-v^2}\right)$	$t = \frac{r(r-1)(2u-v)+u}{4u^2-v^2} + \zeta(2y) + y$

Table B.1: Points of contact of biquadratic with each facet, and conditions on (t, u, r, v) for biquadratic and facet to touch at exactly one point.

function:

$$t = \frac{(\zeta + 1 - r)^2}{2u + v}, \tag{B.3}$$

$$y(4u^2 - v^2) + u - (\zeta + 1 - r)(2u - v) = 0, \tag{B.4}$$

$$t = \frac{r^2}{2u + v} + 2(\zeta + 1)y, \tag{B.5}$$

$$t = \frac{r(r-1)(2u - v) + u}{4u^2 - v^2} + \zeta(2y) + y. \tag{B.6}$$

We use the following transformation of variables:

$$\theta = 2u - v \tag{B.7}$$

$$\phi = 2u + v \tag{B.8}$$

We convert all the dual variables into functions of θ and ϕ. It directly follows that: $u = \frac{1}{4}(\phi + \theta)$ and $v = \frac{1}{2}(\phi - \theta)$.

From (B.3) and (B.5), we have:

$$r = \frac{\zeta + 1}{2} - y\phi. \tag{B.9}$$

Using (B.9) in (B.3), we have:

$$t = \frac{\left(y\phi + \frac{\zeta+1}{2}\right)^2}{\phi}. \tag{B.10}$$

Using (B.9) and (B.10) in (B.6), we have the following:

$$\phi = \theta\,(2\zeta + 1)\,. \tag{B.11}$$

Note that we have not used (B.4) yet, but substituting (B.9)–(B.11) into (B.4), we find that (B.4) is satisfied already. That is, of the four equations (B.3)–(B.6), one of them is linearly dependent on other three.

We can now write all the dual variables t, u, v, r as a function of θ, summarized as follows:

$$r = \frac{\zeta + 1}{2} - y\theta\,(2\zeta + 1)\,, \tag{B.12}$$

$$t = \frac{\left(y\theta\,(2\zeta + 1) + \frac{1}{2}(\zeta + 1)\right)^2}{\theta\,(2\zeta + 1)}, \tag{B.13}$$

$$u = \frac{1}{2}\theta(\zeta + 1), \tag{B.14}$$

$$v = \theta\zeta. \tag{B.15}$$

Thus, we know that the dual variables need to be of this form so that the biquadratic touches all six facets. We still need to check whether the points at which the biquadratic touches each facet satisfies the corresponding ranges of d_1, d_2 in (B.2). Substituting the values (B.12)–(B.15) of the dual variables into the touching points in Table B.1, and observing that $\zeta > 0$, we find that the dual variables are feasible (i.e., the touching points are in the required range) for any $\theta < 0$ (see Table B.2).

Thus, we consider the following optimization program:

$$\sup_{\theta<0} \frac{1}{4(2\zeta + 1)} \left[a + b\theta + \frac{c}{\theta}\right] \tag{B.16}$$

where,

$$a = 4(y + m)(\zeta + 1)\,(2\zeta + 1)\,, \tag{B.17}$$

$$b = 4\,(2\zeta + 1)^2 \left[(y - m)^2 + \sigma^2 \left(\frac{\zeta + 1 + \zeta\rho}{2\zeta + 1}\right)\right], \tag{B.18}$$

$$c = (\zeta + 1)^2, \tag{B.19}$$

Note that the objective function is the objective of a dual feasible solution (B.12)–(B.15) parameterized by θ. The supremum is achieved at $\theta^* = -\sqrt{\frac{c}{b}}$, where $b > 0$ since we have that $\zeta > 0$, $\sigma > 0$, and $\rho \in (-1, 1)$. Let $\gamma := \frac{\zeta+1+\zeta\rho}{2\zeta+1} \in (0, 1]$. The

Facet i	**Contact points**	**Condition on θ**
1	$(d_1^*, d_2^*) = \left(y + \frac{1}{2}\frac{(\zeta+1)}{\theta(2\zeta+1)}, y + \frac{1}{2}\frac{(\zeta+1)}{\theta(2\zeta+1)}\right)$	$d_1^* \leq y,\ d_2^* \leq y \ \Leftrightarrow\ \theta < 0$
2	$(d_1^*, d_2^*) = \left(y + \frac{1}{2}\frac{3\zeta+1}{\theta(2\zeta+1)}, y - \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}\right)$	$d_1^* \leq y \leq d_2^*,\ d_1^* + d_2^* \leq 2y \ \Leftrightarrow\ \theta < 0$
3	$(d_1^*, d_2^*) = \left(y - \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}, y + \frac{1}{2}\frac{3\zeta+1}{\theta(2\zeta+1)}\right)$	$d_2^* \leq y \leq d_1^*,\ d_1^* + d_2^* \leq 2y \ \Leftrightarrow\ \theta < 0$
4	$(d_1^*, d_2^*) = \left(y - \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}, y - \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}\right)$	$d_1^* \geq y,\ d_2^* \geq y \ \Leftrightarrow\ \theta < 0$
5	$(d_1^*, d_2^*) = \left(y + \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}, y - \frac{1}{2}\frac{3\zeta+1}{\theta(2\zeta+1)}\right)$	$d_1^* \leq y \leq d_2^*,\ d_1^* + d_2^* \geq 2y \ \Leftrightarrow\ \theta < 0$
6	$(d_1^*, d_2^*) = \left(y - \frac{1}{2}\frac{3\zeta+1}{\theta(2\zeta+1)}, y + \frac{1}{2}\frac{\zeta+1}{\theta(2\zeta+1)}\right)$	$d_2^* \leq y \leq d_1^*,\ d_1^* + d_2^* \geq 2y \ \Leftrightarrow\ \theta < 0$

Table B.2: Condition on θ so that the points of contact of biquadratic with each facet occurs in the required range.

optimal θ^* is given by:

$$\theta^* = -\frac{(\zeta+1)}{2\,(2\zeta+1)\,\sqrt{(y-m)^2+\gamma\sigma^2}}. \tag{B.20}$$

The corresponding dual value is $\frac{1}{4(2\zeta+1)}(a - 2\sqrt{bc})$, where a, b, c are according to the equations (B.17)–(B.19), which simplifies to:

$$(\zeta+1)\left(y + m - \sqrt{(y-m)^2+\gamma\sigma^2}\right) \leq M(y), \tag{B.21}$$

where the inequality follows from weak duality.

All that is left to prove the proposition is to show that, if $\gamma(1+\nu^2) \geq 2$ where $\nu := \frac{3\zeta+1}{\zeta+1} \in (1,3)$, then there exists a six-point distribution $f_y^* \in \mathcal{F}^{m\sigma\rho}$ whose objective value (B.1) is

$$\mathbb{E}_{f_y^*}\left[\zeta \min(\tilde{d}_1 + \tilde{d}_2, 2y) + \sum_{j=1,2} \min(\tilde{d}_j, y)\right] = (\zeta+1)\left(y + m - \sqrt{(y-m)^2+\gamma\sigma^2}\right). \tag{B.22}$$

To construct the distribution, we use the contact points of the biquadratic to each facet as the support points. Define $z_y := (y-m)/\sigma$ and $\Phi(z_y) := \sqrt{z_y^2+\gamma}$, where we note that $\Phi(z_y) > z_y$. If we use the optimal θ^*, defined in (B.20), to find the associated contact points in Table B.2, where we use the fact that $\theta^*\left(2\zeta+1\right) = \frac{-1}{2\Phi(z_y)\sigma}(\zeta+1)$, we get the following six support points of f_y^*:

$$
\mathbf{D}^{(1)} = \begin{bmatrix} m + (z_y - \Phi(z_y))\sigma \\ m + (z_y - \Phi(z_y))\sigma \end{bmatrix},\ \mathbf{D}^{(2)} = \begin{bmatrix} m + (z_y - \nu\Phi(z_y))\sigma \\ m + (z_y + \Phi(z_y))\sigma \end{bmatrix},\ \mathbf{D}^{(3)} = \begin{bmatrix} m + (z_y + \Phi(z_y))\sigma \\ m + (z_y - \nu\Phi(z_y))\sigma \end{bmatrix}
$$

$$
\mathbf{D}^{(4)} = \begin{bmatrix} m + (z_y + \Phi(z_y))\sigma \\ m + (z_y + \Phi(z_y))\sigma \end{bmatrix},\ \mathbf{D}^{(5)} = \begin{bmatrix} m + (z_y - \Phi(z_y))\sigma \\ m + (z_y + \nu\Phi(z_y) \end{bmatrix},\ \mathbf{D}^{(6)} = \begin{bmatrix} m + (z_y + \nu\Phi(z_y)\sigma \\ m + (z_y - \Phi(z_y))\sigma \end{bmatrix}
$$

We next construct probabilities for the distribution f_y^* to ensure that it is a feasible distribution in $\mathcal{F}^{m\sigma\rho}$. In particular, we find the probabilities $\pi_1, \pi_2, \ldots, \pi_6$ such that the following relationships are true:

$$
\sum_{i=1}^{6} \pi_i = 1 \tag{B.23}
$$

$$
\sum_{i=1}^{6} \pi_i \mathbf{D}^{(i)} = \begin{bmatrix} m \\ m \end{bmatrix} \tag{B.24}
$$

$$
\sum_{i=1}^{6} \pi_i \mathbf{D}^{(i)} \odot \mathbf{D}^{(i)} = \begin{bmatrix} m^2 + \sigma^2 \\ m^2 + \sigma^2 \end{bmatrix} \tag{B.25}
$$

$$
\sum_{i=1}^{6} \pi_i d_1^{(i)} d_2^{(i)} = m^2 + \rho\sigma^2, \tag{B.26}
$$

where $\mathbf{a} \odot \mathbf{b} = (a_i b_i)$ denotes element-wise multiplication of vectors $\mathbf{a}, \mathbf{b}$.

From the equalities (B.23)–(B.26), we have the following system of linear equations: (where for notational brevity, we drop the subscript on z_y and drop the dependence of Φ on z_y)

$$
\begin{array}{llllllr}
\pi_1 & +\pi_2 & +\pi_3 & +\pi_4 & +\pi_5 & +\pi_6 & = 1 \\
(z-\Phi)\pi_1 & +(z-\nu\Phi)\pi_2 & +(z+\Phi)\pi_3 & +(z+\Phi)\pi_4 & (z-\Phi)\pi_5 & +(z+\nu\Phi)\pi_6 & = 0 \\
(z-\Phi)\pi_1 & +(z+\Phi)\pi_2 & +(z-\nu\Phi)\pi_3 & +(z+\Phi)\pi_4 & +(z+\nu\Phi)\pi_5 & +(z-\Phi)\pi_6 & = 0 \\
(z-\Phi)^2\pi_1 & +(z-\nu\Phi)^2\pi_2 & +(z+\Phi)^2\pi_3 & +(z+\Phi)^2\pi_4 & (z-\Phi)^2\pi_5 & +(z+\nu\Phi)^2\pi_6 & = 1 \\
(z-\Phi)^2\pi_1 & +(z+\Phi)^2\pi_2 & +(z-\nu\Phi)^2\pi_3 & +(z+\Phi)^2\pi_4 & +(z+\nu\Phi)^2\pi_5 & +(z-\Phi)^2\pi_6 & = 1 \\
(z-\Phi)^2\pi_1 & +(z-\nu\Phi)(z+\Phi)\pi_2 & +(z-\nu\Phi)(z+\Phi)\pi_3 & +(z+\Phi)^2\pi_4 & +(z+\nu\Phi)(z-\Phi)\pi_5 & +(z+\nu\Phi)(z-\Phi)\pi_6 & = \rho
\end{array}
$$

By simple row operations, we can show that the last equation is linearly dependent on the others. Additionally, it is easy to see that if we interchange π_2 and π_3 as well as π_5 and π_6, the equations remain unaltered, thus $\pi_2 = \pi_3$ and $\pi_5 = \pi_6$. Thus, the new

Facet i	**Support Point**	**Probability**
1	$(d_1^*, d_2^*) = (m + (z_y - \Phi(z_y))\sigma,\ m + (z_y - \Phi(z_y))\sigma)$	$\frac{1}{2} + \frac{2\zeta\pi}{\zeta+1} + \frac{z_y}{2\sqrt{z_y^2+\gamma}} - \frac{(1-\gamma)}{2(\nu-1)(z_y^2+\gamma)}$
2	$(d_1^*, d_2^*) = (m + (z_y - \nu\Phi(z_y))\sigma,\ m + (z_y + \Phi(z_y))\sigma)$	$-\pi + \frac{(1-\gamma)}{(\nu^2-1)(z_y^2+\gamma)}$
3	$(d_1^*, d_2^*) = (m + (z_y + \Phi(z_y))\sigma,\ m + (z_y - \nu\Phi(z_y))\sigma)$	$-\pi + \frac{(1-\gamma)}{(\nu^2-1)(z_y^2+\gamma)}$
4	$(d_1^*, d_2^*) = (m + (z_y + \Phi(z_y))\sigma,\ m + (z_y + \Phi(z_y))\sigma)$	$\frac{1}{2} - \frac{2\zeta\pi}{\zeta+1} - \frac{z_y}{2\sqrt{z_y^2+\gamma}} - \frac{(1-\gamma)(3-\nu)}{2(\nu^2-1)(z_y^2+\gamma)}$
5	$(d_1^*, d_2^*) = (m + (z_y - \Phi(z_y))\sigma,\ m + (z_y + \nu\Phi(z_y))$	π
6	$(d_1^*, d_2^*) = (m + (z_y + \nu\Phi(z_y))\sigma,\ m + (z_y - \Phi(z_y))\sigma)$	π

Table B.3: The support points and the corresponding probabilities in a worst-case probability distribution $f^*_{y,\pi}$, where $\max(0, \alpha_1) \leq \pi \leq \min(\beta_1, \beta_2)$.

system of equations are:

$$\begin{aligned}
&\pi_1 &&+ 2\pi_2 &&+\pi_4 &&+ 2\pi_5 &&= 1\\
&(z-\Phi)\pi_1 &&+ (2z + (1-\nu)\Phi)\pi_2 &&+(z+\Phi)\pi_4 &&+ (2z - (1-\nu)\Phi)\pi_5 &&= 0\\
&(z-\Phi)^2\pi_1 &&+ ((z-\nu\Phi)^2 + (z+\Phi)^2)\pi_2 &&+(z+\Phi)^2\pi_4 &&((z-\Phi)^2 + (z+\nu\Phi)^2)\pi_5 &&= 1
\end{aligned}$$

Since we have three equations and four unknowns, we use parameter $\pi_5 = \pi$. Then, the solution to the set of equations are as follows:

$$\begin{aligned}
\pi_1 &= \frac{1}{2} + \frac{2\zeta\pi}{\zeta+1} + \frac{z}{2\sqrt{z^2+\gamma}} - \frac{(1-\gamma)}{2(\nu-1)(z^2+\gamma)}\\
\pi_2 = \pi_3 &= -\pi + \frac{(1-\gamma)}{(\nu^2-1)(z^2+\gamma)}\\
\pi_4 &= \frac{1}{2} - \frac{2\zeta\pi}{\zeta+1} - \frac{z}{2\sqrt{z^2+\gamma}} - \frac{(1-\gamma)(3-\nu)}{2(\nu^2-1)(z^2+\gamma)}\\
\pi_5 = \pi_6 &= \pi
\end{aligned}$$

We need to ensure that the probabilities lie in $[0, 1]$ (they already sum up to one because of (B.23)), which can be accomplished by putting restrictions on the value of π. Defining:

$$\begin{aligned}
\alpha_1(z) &:= \left(\frac{\zeta+1}{2\zeta}\right)\left(-\frac{1}{2} - \frac{z}{2\sqrt{z^2+\gamma}} + \frac{(1-\gamma)(\nu+1)}{2(\nu^2-1)(z^2+\gamma)}\right),\\
\beta_1(z) &:= \frac{1-\gamma}{(\nu^2-1)(z^2+\gamma)},\\
\beta_2(z) &:= \left(\frac{\zeta+1}{2\zeta}\right)\left(\frac{1}{2} - \frac{z}{2\sqrt{z^2+\gamma}} + \frac{(1-\gamma)(\nu-3)}{2(\nu^2-1)(z^2+\gamma)}\right),
\end{aligned}$$

we have that the probabilities are nonnegative for $\max(0,\alpha_1) \le \pi \le \min(\beta_1,\beta_2)$. Note that $\beta_1 \ge 0$. If $\beta_2 < 0$, $\beta_1 < \alpha_1$, or $\beta_2 < \alpha_2$, then the set of feasible values for π is empty. However, according to the following lemma, if $\gamma(\nu^2+1) \ge 2$ then this set is non-empty for all values of z.

Lemma B.1.1 *If $\gamma(\nu^2+1) \ge 2$, then $\beta_2(z) \ge 0$, $\alpha_1(z) \le \beta_1(z)$, and $\alpha_1(z) \le \beta_2(z)$ for all $z \in \Re$.*

PROOF. Since $\nu - 1 = \frac{2\zeta}{\zeta+1}$, we can rewrite the following:

$$\begin{aligned}
\alpha_1 &:= \left(\frac{1}{2(\nu-1)}\right)\left(-1 - \frac{z}{\Phi(z)} + \frac{(1-\gamma)(\nu+1)}{(\nu^2-1)\Phi^2(z)}\right),\\
\beta_1 &:= \frac{1-\gamma}{(\nu^2-1)\Phi^2(z)},\\
\beta_2 &:= \left(\frac{1}{2(\nu-1)}\right)\left(1 - \frac{z}{\Phi(z)} + \frac{(1-\gamma)(\nu-3)}{(\nu^2-1)\Phi^2(z)}\right),
\end{aligned}$$

Note that $\beta_2(z) \ge 0$ if and only if: $\Phi(z)(\Phi(z)-z) \ge \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$. Let $w(z) = \Phi(z)(1-z)$. Then $w'(z) = 2z - \frac{z^2}{\sqrt{z^2+\gamma}} - \sqrt{z^2+\gamma}$, and $w''(z) = 2 - \frac{3z}{\sqrt{z^2+\gamma}} + \frac{z^3}{(z^2+\gamma)^{\frac{3}{2}}}$. Note that $w''(z)$ can be shown to be non-negative (we can prove: $-2 \le -\frac{3z}{\sqrt{z^2+\gamma}} + \frac{z^3}{(z^2+\gamma)^{\frac{3}{2}}} \le 2$), implying that $w(z)$ is a convex function minimized at $z=0$ (from equating $w'(z)=0$). Thus, whenever $\gamma \ge \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$, we have $\beta_2(z) \ge 0$ for all z. The sufficient condition translates to: $\gamma \ge \frac{3-\nu}{\nu^2-\nu+2}$.

$\beta_2(z) \ge \alpha_1(z)$ if and only if: $\frac{2(1-\gamma)}{(\nu^2-1)\Phi^2(z)} \le 1$, which simplifies to: $z^2 \ge \frac{2-\gamma(\nu^2+1)}{\nu^2-1}$. Thus, a sufficient condition is given by: $0 \ge \frac{2-\gamma(\nu^2+1)}{\nu^2-1}$, which translates to: $\gamma(\nu^2+1) \ge 2$. Note that the condition $\gamma \ge \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$ is implied by $\gamma(\nu^2+1) \ge 2$.

$\beta_1(z) \ge \alpha_1(z)$ if and only if: $\Phi(z)(\Phi(z)+z) \ge \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$. The left hand side can be shown to be a convex function minimized at $z=0$ from the same argument in the case $\beta_2(z) \ge 0$. Thus, the sufficient condition is the same as the case $\beta_2(z) \ge 0$. ■

Let us define $f^*_{y,\pi}$ as the six-point distribution that is summarized in Table B.3 for some valid π. Note that the probabilities of $f^*_{y,\pi}$ only ensure that the distribution has the appropriate moments to belong in $\mathcal{F}^{m\sigma\rho}$. We also need to ensure that the strong duality condition (B.22) is true. The left-hand side of (B.22) evaluates to

$$
\begin{aligned}
P(y) = &\left(\frac{1}{2} + \frac{2\zeta\pi}{\zeta+1} + \frac{z}{2\sqrt{z^2+\gamma}} - \frac{(1-\gamma)}{2(\nu-1)(z^2+\gamma)}\right)\Big(2\left(\zeta+1\right)\left(m+\left(z-\Phi(z)\right)\sigma\right)\Big) \\
&+ 2\left(-\pi + \frac{(1-\gamma)}{(\nu^2-1)(z^2+\gamma)}\right)\Big((\zeta+1)(2m+\sigma(2z+(1-\nu)\Phi(z))) - \sigma\Phi(z)\Big) \\
&+ \left(\frac{1}{2} - \frac{2\zeta\pi}{\zeta+1} - \frac{z}{2\sqrt{z^2+\gamma}} + \frac{(1-\gamma)(\nu-3)}{2(\nu^2-1)(z^2+\gamma)}\right)\Big(2(\zeta+1)y\Big) \\
&+ 2\left(\pi\right)\Big((\zeta+1)(2y) - \sigma\Phi(z)\Big)
\end{aligned} \tag{B.27}
$$

The coefficient of π in (B.27) is given by:

$$
\begin{aligned}
&= \frac{2\zeta}{\zeta+1}\Big(2(\zeta+1)\left[m+\left(z-\Phi(z)\right)\sigma - y\right]\Big) + 2(\zeta+1)\left[2y-2m-\sigma(2z+(1-\nu)\Phi(z))\right] \\
&= 4\zeta(-\sigma\Phi(z)) + 2(\zeta+1)(-\sigma(1-\nu)\Phi(z)) \qquad\qquad (y = m + z\sigma) \\
&= 0 \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (1-\nu = -2\zeta/(\zeta+1))
\end{aligned}
$$

Hence, for any π, the left-hand side of (B.22) with $f_y^* = f_{y,\pi}^*$ simplifies to

$$
\begin{aligned}
P(y) &= (\zeta+1)(2y-(z+\Phi(z))\sigma) = (\zeta+1)(y+m-\Phi(z)\sigma) \\
&= (\zeta+1)(y+m-\sqrt{(y-m)^2+\gamma\sigma^2})
\end{aligned}
$$

which is equal to the right-hand side of (B.22). This completes our proof. ■

B.1.3 Proof of Proposition 3.4.2

Let $\mathbf{y}^* = (y_1^*, y_2^*)$ be the optimal solution of the distributionally robust problem (3.5). Since the locations are identical, we have that $y_1^* = y_2^* = y^*$ for some y^*. Hence, we need only consider the subset of inventory levels $\mathbf{y} = (y, y)$, for which we derive an analytic expression of the worst-case cost as $\bar{C}(y)$ in (3.10). Thus, the distributionally robust problem (3.5) is equivalent to $\min_y \bar{C}(y)$. The first two derivatives of $\bar{C}(y)$ are

$$
\begin{aligned}
\bar{C}'(y) &= -(p-h-s_0) + \frac{(p+h-s_0)(y-m)}{\sqrt{(y-m)^2+\gamma\sigma^2}} \\
\bar{C}''(y) &= \frac{(p+h-s_0)\gamma\sigma^2}{((y-m)^2+\gamma\sigma^2)^{\frac{3}{2}}}
\end{aligned}
$$

Since $\gamma > 0$, $\bar{C}(y)$ is convex in y, and the optimal solution is given by the first-order condition $\bar{C}'(y^*) = 0$, which gives y^* as the right-hand side of (3.11). ■

B.1.4 Proof of Lemma 3.5.1

We define

$$
\begin{aligned}
M(\mathbf{y}) = &\max_f \ \mathbb{E}_f \left[\sum_{\ell=0}^{L-1} \eta_\ell^\top \left(\mathbf{E}_\ell \mathbf{D} - \mathbf{E}_\ell \mathbf{y} \right)^+ \right] \\
&\text{s.t. } \mathbb{E}_f(1) = 1 \\
&\qquad \mathbb{E}_f \left(\tilde{\mathbf{D}} \right) = \mathbf{m} \\
&\qquad \mathbb{E}_f \left(\tilde{\mathbf{D}} \tilde{\mathbf{D}}^\top \right) = \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \\
&\qquad f(\mathbf{D}) \geq 0, \quad \forall \mathbf{D} \in \Re^n.
\end{aligned}
$$

which is equal to the left-hand side of (3.13).

Since $\Sigma \succ 0$, then the moments $(\mathbf{m}, \mathbf{\Sigma})$ are strictly in the interior of the feasible moment cone. Hence, strong duality of moment problems holds (Smith, 1995). The dual of the moment problem is

$$
\begin{aligned}
M(\mathbf{y}) = \quad &\min_{t, \mathbf{r}, \mathbf{Y}} \quad t + \mathbf{r}^\top \mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle \\
&\text{s.t} \quad t + \mathbf{r}^\top \mathbf{x} + \mathbf{x}^\top \mathbf{Y} \mathbf{x} \geq \sum_{\ell=0}^{L-1} \eta_\ell^\top \left(\mathbf{E}_\ell \mathbf{x} - \mathbf{E}_\ell \mathbf{y} \right)^+, \quad \forall \mathbf{x} \in \Re^n
\end{aligned}
$$

We can reformulate the dual as the following semi infinite linear program:

$$
\begin{aligned}
M(\mathbf{y}) = \quad &\min_{t, \mathbf{r}, \mathbf{Y}} \quad t + \mathbf{r}^\top \mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle \\
&\text{s.t} \quad t + \mathbf{r}^\top \mathbf{x} + \mathbf{x}^\top \mathbf{Y} \mathbf{x} \geq \sum_{\ell=0}^{L-1} \left(\eta_\ell \odot \mathbf{e}_{A_\ell} \right)^\top \left(\mathbf{E}_\ell \mathbf{x} - \mathbf{E}_\ell \mathbf{y} \right), \quad \forall \mathbf{x} \in \Re^n \\
&\qquad\qquad \forall (A_0, A_1, \cdots, A_{L-1}) \in 2^{[n_0]} \times 2^{[n_1]} \times 2^{[n_{L-1}]}
\end{aligned}
$$

where $\odot$ is the element-wise product operator, and $\mathbf{e}_{A_\ell}$ is an n_ℓ-dimensional binary vector whose k^{th} element is 1 if and only if $k \in A_\ell$. For simplicity, we can write the right-hand side as $\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top \mathbf{y}$ for $k \in [2^N]$, where $N = \sum_{\ell=0}^{L-1} n_\ell$. The constraint now becomes: $\mathbf{x}^\top \mathbf{Y}\mathbf{x} + (\mathbf{r}-\mathbf{a_k})^\top \mathbf{x} + t - \mathbf{b}_k^\top \mathbf{y} \geq 0, \forall \mathbf{x}$. This is true if and only if

$$\begin{bmatrix} \mathbf{Y} & \frac{1}{2}(\mathbf{r}-\mathbf{a_k}) \\ \frac{1}{2}(\mathbf{r}-\mathbf{a_k})^\top & t - \mathbf{b}_k^\top \mathbf{y} \end{bmatrix} \succeq 0, \quad \forall k. \quad \blacksquare$$

$\square$

B.2 Optimal Inventory Solutions for Two-Locations Systems

We have four cases for which the distributionally robust solution needs to be calculated: pooling/no pooling (P/NP), and known/unknown correlation (C/NC). Note that we restrict the search to identical solutions of the form (y, y).

1. **No pooling, ρ unknown:** This is the same setting as Scarf (1958), and the optimal inventory and worst-case cost are given by:

$$\begin{aligned} y^{NP,NC} &= m + \frac{p-h-s_0}{2\sqrt{h(p-s_0)}} \cdot \sigma \\ C^{NP,NC} &= 2m(s_0-h) + 2hy + (p+h-s_0)(m-y+\sqrt{\sigma^2+(m-y)^2}) \end{aligned}$$

2. **No pooling, ρ known:** This is the same setting as Natarajan and Teo (2017),

and the solutions are given through an SDP.

$$
\begin{aligned}
C^{NP,C} := \min_{t_0,r,u,v,y} \quad & 2(s_0-h)m + 2hy + (p+h-s_0)\big(t_0 + 2rm \\
& \qquad + 2u(m^2+\sigma^2) + v(m^2+\rho\sigma^2)\big) \\
\text{s.t.} \quad & \begin{pmatrix} t_0+2y & \frac{1}{2}(r-1) & \frac{1}{2}(r-1) \\ \frac{1}{2}(r-1) & u & \frac{1}{2}v \\ \frac{1}{2}(r-1) & \frac{1}{2}v & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+y & \frac{1}{2}(r-1) & \frac{1}{2}r \\ \frac{1}{2}(r-1) & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+y & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0 & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0 \\
& y \geq 0
\end{aligned}
$$

3. **With pooling, ρ unknown:** This is simply an extension of our setting where only marginal information (m, σ) is known, and cross-moment information (ρ) is

unknown. The solutions are given through an SDP.

$$
\begin{aligned}
C^{P,NC} := \min_{t_0,r,u,y} \quad & h(2y-2m)+2s_0m+t_0+2rm+2u(m^2+\sigma^2) \\
\text{s.t.} \quad & \begin{pmatrix} t_0 & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+(s-s_0)y & \frac{1}{2}r & \frac{1}{2}(r-s+s_0) \\ \frac{1}{2}r & u & 0 \\ \frac{1}{2}(r-s+s_0) & 0 & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+(s-s_0)y & \frac{1}{2}(r-s+s_0) & \frac{1}{2}r \\ \frac{1}{2}(r-s+s_0) & u & 0 \\ \frac{1}{2}r & 0 & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+2(p+h-s_0)y & \frac{1}{2}(r-p-h+s_0) & \frac{1}{2}(r-p-h+s_0) \\ \frac{1}{2}(r-p-h+s_0) & u & 0 \\ \frac{1}{2}(r-p-h+s_0) & 0 & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+y(s-s_0)+2y(p+h-s) & \frac{1}{2}(r-p-h+s) & \frac{1}{2}(r-p-h+s_0) \\ \frac{1}{2}(r-p-h+s) & u & 0 \\ \frac{1}{2}(r-p-h+s_0) & 0 & u \end{pmatrix} \succeq 0 \\
& \begin{pmatrix} t_0+y(s-s_0)+2y(p+h-s) & \frac{1}{2}(r-p-h+s_0) & \frac{1}{2}(r-p-h+s) \\ \frac{1}{2}(r-p-h+s_0) & u & 0 \\ \frac{1}{2}(r-p-h+s) & 0 & u \end{pmatrix} \succeq 0 \\
& y \geq 0
\end{aligned}
$$

4. **With pooling, ρ known:** This is the setting considered by our paper, and the solutions $y^{P,C}, C^{P,C}$ are given in closed-form in Proposition 3.4.2.

B.3 Example for Generating Nested Fulfillment Structure with $L < n$

Example 2 *Consider Figure B.1, where a nested fulfillment structure with $L = 4$ is created from the dendrogram in Figure 3.7a. Here, the range of distances are partitioned into three quantiles by the two lines drawn on the dendrogram. In Figure B.1a, the lower line gives rise to three connected components: $\{\{3\}, \{1\}, \{4, 5, 2\}\}$. The nodes in each connected component are considered to be a single cluster in level $l = 1$, and the*

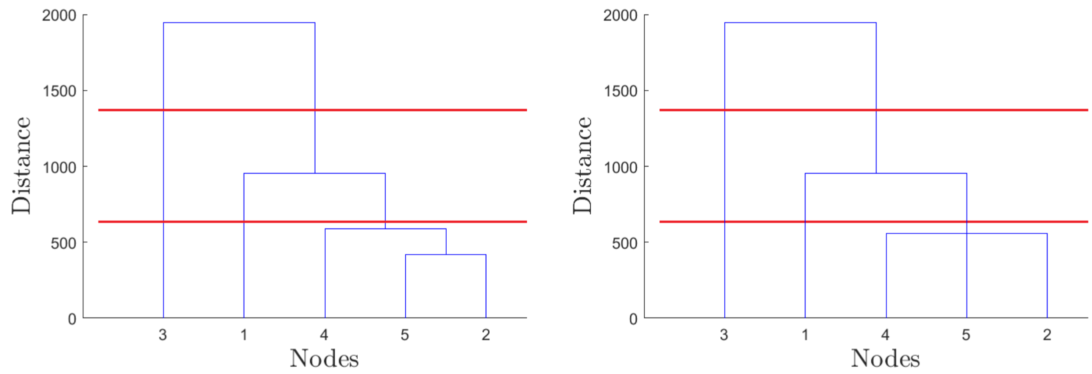

(a) (b)

Figure B.1: Creating a nested fulfillment structure with $L = 4$ from a dendrogram.

UPGMA distances are recalculated for the new clusters. The upper line gives rise to two connected components: $\{\{3\}, \{1, 4, 5, 2\}\}$*, which form the two components at level* $l = 2$*, resulting in Figure B.1b.*

B.4 Details for Numerical Experiments

B.4.1 Constant Fulfillment Heuristic

For $n = 5$, the marginal distribution parameters for the four distributions (Normal, Exponential, BetaPrime and Student-t) in the following way:

1. Normal: the means are identical with $m = 300$, and the standard deviation is chosen at random from $[100, 800]$.

2. Exponential: the mean of the exponential distribution is chosen at random from [100,500]. The standard deviation is equal to the mean.

3. BetaPrime: the mean is fixed at $m = 2$. The parameters α and *beta* are chosen as follows. β is chosen at random from $[2, 3]$, and $\alpha = m \cdot (\beta - 1)$.

4. Student-t: the parameter ν is chosen at random from $[2, 3]$.

We generate 50 such instances of marginal distribution parameters. We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012). Then, using the method of Gaussian copula, we generate 5000 correlated random demand samples for each distribution, and report the sample average approximations.

B.4.2 Nested Fulfillment Heuristic

The mean and covariance matrices are calculated based on the populations for each fulfillment center. Given mean m and variance v for a demand distribution, we calculate the marginal distribution parameters for four distributions (Normal, Exponential, BetaPrime and Pareto) as follows:

1. Normal: Mean $\mu = m$, Variance $\sigma^2 = v$
2. Exponential: Mean $\frac{1}{\lambda} = m = \sqrt{v}$
3. BetaPrime: $\beta = 2 + \frac{m\cdot(m+1)}{v}$, $\alpha = m \cdot (\beta - 1)$
4. Generalized Pareto: $k = -\frac{v}{m^2} + \sqrt{\frac{v^2}{m^4} + \frac{v}{m^2}}$, $\sigma = m \cdot k \cdot (1-k)$, $\theta = \frac{\sigma}{k}$.

We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012) such that the correlation coefficients do not exceed .4 in magnitude. We then use the Gaussian copula to generate 10^3 training samples of correlated random vectors. The stochastic solutions are calculated based on a sample average approximation linear program using these training samples, and the robust solution is calculated based on the partitioned statistics estimated from the training samples. The inventory solutions are then evaluated through simulations based on 10^3 test samples generated in a similar fashion to the training samples.

B.5 Asymmetry Information

Based on Natarajan et al. (2017), we incorporate into our robust models the partitioned statistics information. Specifically, the mean and covariance of random vector $(\tilde{\mathbf{D}}^+, \tilde{\mathbf{D}}^-)$ whose i^{th} elements are $(\tilde{d}_i - m_i)^+$ and $(m_i - \tilde{d}_i)^+$, respectively, are defined to be:

$$\mathbb{E}\left[\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}\right] =: \bar{\mathbf{m}} \qquad \mathbb{E}\left[\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}^\intercal\right] =: \bar{\mathbf{Q}} \tag{B.28}$$

The set of distributions that the random demand can take is defined as $\bar{\mathcal{F}}_{\geq 0}$, which specifies that the random demand has non-negative support, with mean $\mathbf{m}$, and with mean and covariance of the partitioned statistics given in (B.28). We follow the same approach as in Theorem 4.3 in Natarajan et al. (2017) to derive the following upper bound including the partitioned statistics information. We omit the proof to avoid repetition.

Proposition B.5.1 *For the n-location newsvendor problem under inventory risk pooling with a L-level nested fulfillment cost structure, we have* $\sup_{f \in \bar{\mathcal{F}}_{\geq 0}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] \leq \bar{\bar{C}}_L(\mathbf{y})$ *for any* $\mathbf{y} \in \Re^n$*, where*

$$
\begin{aligned}
\bar{\bar{C}}_L(\mathbf{y}) := \min_{\substack{t_0, \mathbf{t}, \mathbf{Y}, \mathbf{u}, \\ \mathbf{B}, \mathbf{W}, \mathbf{U}, \mathbf{V}}} \; & h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{m}) + \mathbf{s}_0^\top \mathbf{m} + t_0 + \mathbf{t}^\top \bar{\mathbf{m}} + \langle \mathbf{Y}, \bar{\mathbf{Q}} \rangle + \mathbf{e}^\top \mathbf{B} \mathbf{e} \\
s.t. \; & \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top \\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top \\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U} \end{pmatrix} \succeq 0 \\
& \mathbf{u} = -\mathbf{W}\mathbf{e} + \left(\mathbf{B} + \mathbf{B}^\top\right)\mathbf{e} + \mathbf{P}(\mathbf{y} - \mathbf{m}) \\
& \mathbf{V} \geq \bar{\mathbf{P}} \\
& \mathbf{U} \leq \mathbf{W} - \mathbf{B} \\
& \mathbf{W}, \mathbf{B} \geq 0 \\
& t_0 \in \Re, \; \mathbf{t} \in \Re^{2n}, \; \mathbf{u} \in \Re^N, \; \mathbf{Y} \in \Re^{2n \times 2n}, \; \mathbf{B}, \mathbf{W}, \mathbf{U} \in \Re^{N \times N}, \; \mathbf{V} \in \Re^{N \times 2n},
\end{aligned}
$$

with $\mathbf{P} := \Big(\mathbf{E}_{L-1}^\top \mathrm{diag}(\eta_{L-1}) \quad \mathbf{E}_{L-2}^\top \mathrm{diag}(\eta_{L-2}) \quad \cdots \quad \mathbf{E}_0^\top \mathrm{diag}(\eta_0)\Big)^\top \in \Re^{N \times n}$*, and* $\bar{\mathbf{P}} = \Big[\mathbf{P} \quad -\mathbf{P}\Big] \in \Re^{N \times 2n}$.

The heuristic solution can be similarly obtained by setting $\mathbf{y}$ as a decision variable, constrained by $\mathbf{y} \geq 0$.

B.6 Multiple Demand Channels

To simplify our discussion, we consider a two-level nested fulfillment cost structure for the online demand (i.e., where cross-location fulfillment cost is constant), though the technique can be generalized to an L-level structure. Let p_b and p_o be the penalty cost of unmet brick-and-mortar store demand and online demand, respectively. The per-unit overage cost is h. We normalize the cost for meeting store demand to zero. As before, the cost of in-location fulfillment of online demand is s_0, and the cost of cross-location fulfillment is s, where $s > s_0$. For a customer region $j \in [n]$, let $\tilde{d}_j^o$ and $\tilde{d}_j^b$ be the stochastic online demand and the stochastic store demand, respectively. We denote the vector of online demands as $\tilde{\mathbf{D}}^o = (\tilde{d}_j^o)$ and the vector of store demands as $\tilde{\mathbf{D}}^b = (\tilde{d}_j^b)$. We let $\tilde{\mathbf{D}} = (\tilde{\mathbf{D}}^b, \tilde{\mathbf{D}}^o)$ as the vector of all demands with a mean vector $\mathbf{m} = \left(\mathbf{m}^b, \; \mathbf{m}^o\right)$ and covariance matrix $\mathbf{\Sigma}$.

Store demand can only be met with inventory from the same location. However, online demand can be fulfilled from inventory from any location. We assume that

$p_o + h > s$, that $p_b + h > s$, and that $p_b + s_0 > p_o$. Given our assumptions on the cost parameters, it is optimal for each local store to first meet the store demands to the maximum extent possible, then for excess inventory to be used to fulfill in-location online demand, before cross-location fulfillment is used. To see why, note that since $p_b + s_0 > p_o$, then it is cheaper to use an inventory unit to meet store demand than to fulfill a local online demand. Moreover, the assumptions imply that $p_b + h + s_0 > s$, so it is cheaper to use cross-location fulfillment on an online demand than to use in-location fulfillment and not meet a store demand. Therefore, we can write the cost as

$$
\begin{aligned}
C(\mathbf{y}, \mathbf{D}) = h \cdot &\left(\sum_{j\in[n]} \left(y_j - d_j^b\right)^+ - \sum_{j\in[n]} d_j^o \right)^+ + p_o \cdot \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(y_j - d_j^b\right)^+ \right)^+ \\
&+ p_b \cdot \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ + s_0 \cdot \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(d_j^o - \left(y_j - d_j^b\right)^+\right)^+ \right) \\
&+ s \cdot \left(\sum_{j\in[n]} \left(d_j^o - \left(y_j - d_j^b\right)^+\right)^+ - \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(y_j - d_j^b\right)^+ \right)^+ \right)
\end{aligned}
$$

We observe that, due to the presence of store demand which is prioritized due to its lower cost of fulfillment, the cost structure is more complicated than before. In particular, the last term in the cost function has a composition of a function $f(x) = (a - x)^+$ and $g(\mathbf{x}) = \sum_j x_j^+$. This requires a careful treatment in developing the tractable SDP heuristic. We first simplify the cost function by reducing the number of such terms using the relationship that if $a \geq 0$, then $(a-(b-c)^+) = (a+c-b)^+ - (c-b)^+$. Also using the fact that $(c-b)^+ = b - c + (c-b)^+$, we can simplify the cost function to

$$
\begin{aligned}
C(\mathbf{y}, \mathbf{D}) = h \cdot \mathbf{e}^\top &\left(\mathbf{y} - \mathbf{D}^o - \mathbf{D}^b\right) + s_0 \cdot \mathbf{e}^\top \mathbf{D}^o \\
&+ (h + p_b + s_o - s) \cdot \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ + (s - s_0) \cdot \sum_{j\in[n]} \left(d_j^o + d_j^b - y_j\right)^+ \\
&+ (h + p_o - s) \cdot \left(\sum_{j\in[n]} \left(d_j^o + d_j^b - y_j\right) - \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ \right)^+
\end{aligned}
$$

We define the constants $\gamma := h + p_b + s_o - s$, $\eta_0 := s - s_0$, and $\eta_1 := h + p_o - s$.

Hence, the minmax expected cost under the omni-channel demand is equivalent to

$$C_{\mathrm{o}}^* := \min_{\mathbf{y}} \left((s_0 - h) \cdot \mathbf{e}^\top \mathbf{m}_o + h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{m}_s) + M_{\mathrm{o}}(\mathbf{y}) \right)$$

where $M_{\mathrm{o}}(\mathbf{y})$ is the optimal value of the moment problem

$$\begin{aligned} M_{\mathrm{o}}(\mathbf{y}) := \max_{f \in \mathcal{F}} \mathbb{E}_f \Bigg[& \gamma \cdot \sum_{j \in [n]} \left(\tilde{d}_j^b - y_j \right)^+ + \eta_0 \cdot \sum_{j \in [n]} \left(\tilde{d}_j^o + \tilde{d}_j^b - y_j \right)^+ \\ & + \eta_1 \cdot \left(\sum_{j \in [n]} \left(\tilde{d}_j^o + \tilde{d}_j^b - y_j \right) - \sum_{j \in [n]} \left(\tilde{d}_j^b - y_j \right)^+ \right)^+ \Bigg]. \end{aligned}$$

We can write the moment problem as

$$\begin{aligned} M_{\mathrm{o}}(\mathbf{y}) := \max_{f \in \mathcal{F}_{\geq 0}} \Bigg[& \max_{\mathbf{x}^{(0)} \in \{0,1\}^n, x^{(1)} \in \{0,1\}, \mathbf{z} \in \{0,1\}^n} \gamma \cdot \mathbf{z}^\top \left(\tilde{\mathbf{D}}^b - \mathbf{y} \right) + \eta_0 \cdot \mathbf{x}^{(0)\top} \left(\tilde{\mathbf{D}}^o + \tilde{\mathbf{D}}^b - \mathbf{y} \right) \\ & + \eta_1 \cdot x^{(1)} \cdot \left(\mathbf{e}^\top \left(\tilde{\mathbf{D}}^o + \tilde{\mathbf{D}}^b - \mathbf{y} \right) - \mathbf{z}^\top \left(\tilde{\mathbf{D}}^b - \mathbf{y} \right) \right) \Bigg] \end{aligned}$$

To see why, note that the coefficient of z_j is equal to $\left(\gamma - \eta_1 x^{(1)}\right) \cdot \left(\tilde{d}_j^b - y_j\right)$. Based on our assumptions on the cost parameters, we have that $\gamma = h + p_b + s_o - s > 0$, and $\gamma - \eta_1 = p_b - p_o + s_0 > 0$. Therefore, z_j is equal to 1 if and only if $d_j^b - y_j \geq 0$. Note that unlike in the previous section where the newly introduced variables only interact with other constants or the random demand, we have cross interactions between the new variables from the term $x^{(1)} \cdot \mathbf{z}$. Hence, we introduce a new n-dimensional vector $\mathbf{w} = x^{(1)} \cdot \mathbf{z}$.

Consider the $(3n+1)$-dimensional random vector $\tilde{\mathbf{x}} := \begin{pmatrix} \tilde{x}^{(1)\top} & \tilde{\mathbf{x}}^{(0)\top} & \tilde{\mathbf{z}}^\top & \tilde{\mathbf{w}} \end{pmatrix}^\top$, which collects all the new binary variables into a single vector. We again have the following transformation

$$\begin{aligned} \mathbf{x} &:= \mathbb{E}_f \left(\tilde{\mathbf{x}} \right) \in \Re^{3n+1}, \\ \mathbf{Q} &:= \mathbb{E}_f \left(\tilde{\mathbf{x}} \tilde{\mathbf{D}}^\top \right) \in \Re^{(3n+1) \times (2n)}, \\ \mathbf{R} &:= \mathbb{E}_f \left(\tilde{\mathbf{x}} \tilde{\mathbf{x}}^\top \right) \in \Re^{(3n+1) \times (3n+1)}. \end{aligned}$$

Therefore, we have linearized the objective to

$$\gamma \cdot \sum_{j\in[n]} (Q_{1+n+j,j} - x_{1+n+j} \cdot y_j) + \eta_0 \cdot \sum_{j\in[n]} (Q_{1+j,j} + Q_{1+j,n+j} - x_{1+j} \cdot y_j) \\ + \eta_1 \cdot \sum_{j\in[n]} (Q_{1,j} + Q_{1,n+j} - x_1 \cdot y_j - Q_{2n+1+j,j} + x_{2n+1+j} \cdot y_j)$$

The constraints are the same as before, but with the addition of a few other constraints that follow from the fact that $\tilde{w}_j = \tilde{x}^{(1)} \cdot \tilde{z}_j$ for all $j \in [n]$. In particular, note that

$$\begin{aligned} R_{1,n+1+j} &= x_{1+2n+j}, & \forall j \in [n], \\ R_{1+n+i,1+2n+j} &= R_{1+2n+1,1+2n+j}, & \forall i \in [n], j \in [n]. \end{aligned}$$

The first constraint follows since the left-hand side is by definition equal to $\mathbb{E}_f\left(\tilde{x}^{(1)} \cdot \tilde{z}_j\right)$, and the right-hand side is $\mathbb{E}_f\left(\tilde{w}_j\right)$. In the second constraint, the left-hand side is equal to $\mathbb{E}_f\left(\tilde{z}_i \cdot \tilde{w}_j\right) = \mathbb{E}_f\left(\tilde{x}^{(1)}\tilde{z}_i\tilde{z}_j\right)$. The right-hand side is equal to $\mathbb{E}_f\left(\tilde{w}_i \cdot \tilde{w}_j\right) = \mathbb{E}_f\left(\tilde{x}^{(1)}\tilde{x}_i\tilde{x}_j\right)$ since $(\tilde{x}^{(1)})^2 = \tilde{x}^{(1)}$. Due to the nonnegativity of demand, aside from the constraint that $\mathbf{Q} \geq 0$, we also have that $\mathbf{z} \leq \mathbf{x}^{(0)}$. This is because $d_j^b - y_j \geq 0$ impies that $d_j^b + d_j^o - y_j \geq 0$, which is equivalent to the condition that $\tilde{z}_j \leq \tilde{x}_j^{(0)}$.

Proposition B.6.1 *For the n-location newsvendor problem under inventory risk pooling with online and store demand in each location, if the cross-location fulfillment costs of online demand are all equal to s, then $\sup_{f\in\mathcal{F}_{\geq 0}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] \leq \bar{C}_o(\mathbf{y})$ for any $\mathbf{y} \in \Re^n$,* *where*

$$
\begin{aligned}
&\bar{C}_o(\mathbf{y}) := \\
&\min_{\substack{t_0,\mathbf{t},\mathbf{Y},\mathbf{u},\\ \mathbf{B},\mathbf{W},\mathbf{U},\mathbf{V},\\ \mathbf{g},\mathbf{h},\mathbf{H}}} \ h\cdot\mathbf{e}^\top\left(\mathbf{y}-\mathbf{m}^o-\mathbf{m}^b\right)+s_0\cdot\mathbf{e}^\top\mathbf{m}^o+t_0+\mathbf{t}^\top\mathbf{m}+\left\langle\mathbf{Y},\boldsymbol{\Sigma}+\mathbf{m}\mathbf{m}^\top\right\rangle+\mathbf{e}^\top\mathbf{B}\mathbf{e} \\
&s.t. \ \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top \\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top \\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U} \end{pmatrix} \succeq 0 \\
&\qquad \mathbf{u} = -\mathbf{W}\mathbf{e}+\left(\mathbf{B}+\mathbf{B}^\top\right)\mathbf{e}+\begin{pmatrix}\mathbf{0}_{1,n}\\ -\mathbf{I}_n\\ \mathbf{I}_n\\ \mathbf{0}_{n,n}\end{pmatrix}\mathbf{g}+\begin{pmatrix}\mathbf{0}_{1,n}\\ \mathbf{0}_{n,n}\\ \mathbf{0}_{n,n}\\ -\mathbf{I}_n\end{pmatrix}\mathbf{h}+\begin{pmatrix}\eta_1\cdot\mathbf{e}_n^\top\\ \eta_0\cdot\mathbf{I}_n\\ \gamma\cdot\mathbf{I}_n\\ -\eta_0\cdot\mathbf{I}_n\end{pmatrix}\mathbf{y} \\
&\qquad \mathbf{V} \geq \begin{pmatrix}\eta_1\cdot\mathbf{e}_n^\top & \eta_1\cdot\mathbf{e}_n^\top\\ \eta_0\cdot\mathbf{I}_n & \eta_0\cdot\mathbf{I}_n\\ \gamma\cdot\mathbf{I}_n & \mathbf{0}_{n,n}\\ -\eta_0\cdot\mathbf{I}_n & \mathbf{0}_{n,n}\end{pmatrix} \\
&\qquad \mathbf{U} \leq \mathbf{W}-\mathbf{B}+\begin{pmatrix}\mathbf{0}_{1,n+1} & \mathbf{h}^\top & \mathbf{0}_n^\top\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & \mathbf{0}_{n,n}\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & \mathbf{H}\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & -\mathbf{H}\end{pmatrix} \\
&\qquad \mathbf{g},\mathbf{W},\mathbf{B}\geq 0 \\
&\qquad t_0\in\Re,\ \mathbf{g},\mathbf{h}\in\Re^n,\ \mathbf{t}\in\Re^{2n},\ \mathbf{u}\in\Re^{3n+1},\ \mathbf{H}\in\Re^{n\times n} \\
&\qquad \mathbf{Y}\in\Re^{2n\times 2n},\ \mathbf{B},\mathbf{W},\mathbf{U}\in\Re^{(3n+1)\times(3n+1)},\ \mathbf{V}\in\Re^{(3n+1)\times 2n}.
\end{aligned}
$$

PROOF. Suppose that $(z(\mathbf{D}),\mathbf{x}(\mathbf{D}),\mathbf{w}(\mathbf{D}))$ are the optimal recourse variables for demand realization $\mathbf{D}$. Let us define the following variables

$$\begin{aligned}
\begin{pmatrix} z \\ \mathbf{x} \\ \mathbf{w} \end{pmatrix} &= \mathbb{E}_f \left(\begin{pmatrix} z(\tilde{\mathbf{D}}) \\ \mathbf{x}(\tilde{\mathbf{D}}) \\ \mathbf{w}(\tilde{\mathbf{D}}) \end{pmatrix} \right) \\
\begin{pmatrix} \mathbf{y}_{zs}^\top & \mathbf{y}_{zo}^\top \\ \mathbf{Y}_{xs} & \mathbf{Y}_{xo} \\ \mathbf{Y}_{ws} & \mathbf{Y}_{wo} \end{pmatrix} &= \mathbb{E}_f \left(\begin{pmatrix} z(\tilde{\mathbf{D}}) \\ \mathbf{x}(\tilde{\mathbf{D}}) \\ \mathbf{w}(\tilde{\mathbf{D}}) \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{D}}_s \\ \tilde{\mathbf{D}}_o \end{pmatrix}^\top \right) \\
\bar{\mathbf{X}} &= \mathbb{E}_f \left(\mathbf{x}(\mathbf{D}) \mathbf{x}(\mathbf{D})^\top \right) \\
\hat{\mathbf{X}} &= \mathbb{E}_f \left(z(\mathbf{D}) \mathbf{x}(\mathbf{D}) \mathbf{x}(\mathbf{D})^\top \right).
\end{aligned}$$

Also define the constants

$$\boldsymbol{\Sigma} + \mathbf{m}\mathbf{m}^\top = \begin{pmatrix} \mathbf{Q}_{ss} & \mathbf{Q}_{so}^\top \\ \mathbf{Q}_{so} & \mathbf{Q}_{oo} \end{pmatrix}.$$

Note that

$$\begin{pmatrix} 1 \\ \mathbf{D}_s \\ \mathbf{D}_o \\ z(\mathbf{D}) \\ \mathbf{x}(\mathbf{D}) \\ \mathbf{w}(\mathbf{D}) \end{pmatrix} \begin{pmatrix} 1 \\ \mathbf{D}_s \\ \mathbf{D}_o \\ z(\mathbf{D}) \\ \mathbf{x}(\mathbf{D}) \\ \mathbf{w}(\mathbf{D}) \end{pmatrix}^\top = \begin{pmatrix} 1 & \mathbf{D}_s^\top & \mathbf{D}_o^\top & z(\mathbf{D}) & \mathbf{x}(\mathbf{D})^\top & \mathbf{w}(\mathbf{D})^\top \\ \mathbf{D}_s & \mathbf{D}_s\mathbf{D}_s^\top & \mathbf{D}_s\mathbf{D}_o^\top & \mathbf{D}_s z(\mathbf{D}) & \mathbf{D}_s\mathbf{x}(\mathbf{D})^\top & \mathbf{D}_s\mathbf{w}(\mathbf{D})^\top \\ \mathbf{D}_o & \mathbf{D}_o\mathbf{D}_s^\top & \mathbf{D}_o\mathbf{D}_o^\top & \mathbf{D}_o z(\mathbf{D}) & \mathbf{D}_o\mathbf{x}(\mathbf{D})^\top & \mathbf{D}_o\mathbf{w}(\mathbf{D})^\top \\ z(\mathbf{D}) & z(\mathbf{D})\mathbf{D}_s^\top & z(\mathbf{D})\mathbf{D}_o^\top & z(\mathbf{D}) & \mathbf{w}(\mathbf{D})^\top & \mathbf{w}(\mathbf{D})^\top \\ \mathbf{x}(\mathbf{D}) & \mathbf{x}(\mathbf{D})\mathbf{D}_s^\top & \mathbf{x}(\mathbf{D})\mathbf{D}_o^\top & \mathbf{w}(\mathbf{D}) & \mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top \\ \mathbf{w}(\mathbf{D}) & \mathbf{w}(\mathbf{D})\mathbf{D}_s^\top & \mathbf{w}(\mathbf{D})\mathbf{D}_o^\top & \mathbf{w}(\mathbf{D}) & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top \end{pmatrix}$$

where we use: $z(\mathbf{D})^2 = z(\mathbf{D})$, $\mathbf{x}(\mathbf{D})\mathbf{w}(\mathbf{D})^\top = z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top$, $z(\mathbf{D})\mathbf{x}(\mathbf{D}) = \mathbf{w}(\mathbf{D})$, $\mathbf{z}(\mathbf{D})\mathbf{w}(\mathbf{D}) = z(\mathbf{D})^2\mathbf{x}(\mathbf{D}) = z(\mathbf{D})\mathbf{x}(\mathbf{D}) = \mathbf{w}(\mathbf{D})$, and finally, $\mathbf{w}(\mathbf{D})\mathbf{w}(\mathbf{D})^\top = z(\mathbf{D})^2\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top = z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top$. Taking the expectation on both sides, we have that

$$\begin{pmatrix} 1 & \mathbf{m}_s^\top & \mathbf{m}_o^\top & z & \mathbf{x}^\top & \mathbf{w}^\top \\ \mathbf{m}_s & \mathbf{Q}_{ss} & \mathbf{Q}_{so}^\top & \mathbf{y}_{zs} & \mathbf{Y}_{xs}^\top & \mathbf{Y}_{ws}^\top \\ \mathbf{m}_o & \mathbf{Q}_{so} & \mathbf{Q}_{oo} & \mathbf{y}_{zo} & \mathbf{Y}_{xo}^\top & \mathbf{Y}_{wo}^\top \\ z & \mathbf{y}_{zs}^\top & \mathbf{y}_{zo}^\top & z & \mathbf{w}^\top & \mathbf{w}^\top \\ \mathbf{x} & \mathbf{Y}_{xs} & \mathbf{Y}_{xo} & \mathbf{w} & \bar{\mathbf{X}} & \hat{\mathbf{X}} \\ \mathbf{w} & \mathbf{Y}_{ws} & \mathbf{Y}_{wo} & \mathbf{w} & \hat{\mathbf{X}} & \hat{\mathbf{X}} \end{pmatrix} \succeq 0,$$

and that

$$\begin{pmatrix} 1 & z & \mathbf{x}^\top & \mathbf{w}^\top \\ z & z & \mathbf{w}^\top & \mathbf{w}^\top \\ \mathbf{x} & \mathbf{w} & \bar{\mathbf{X}} & \hat{\mathbf{X}} \\ \mathbf{w} & \mathbf{w} & \hat{\mathbf{X}} & \hat{\mathbf{X}} \end{pmatrix} \in \mathrm{BQP}.$$

Note that the a linear relaxation of the BQP constraints is the following:

$$\begin{aligned} \mathbf{w} &\le \mathbf{x}, \\ \mathbf{w} &\le z \cdot \mathbf{e}, \\ -\mathbf{w} + \mathbf{x} + z \cdot \mathbf{e} &\le 1, \\ \bar{X}_{ii} &= x_i, \\ \bar{X}_{ij} &\le x_i, \\ -\bar{X}_{ij} + x_i + x_j &\le 1, \\ \hat{X}_{ii} &= w_i, \\ \hat{X}_{ij} &\le w_i, \\ -\hat{X}_{ij} + w_i + w_j &\le 1, \\ \hat{X}_{ij} &\le x_i, \\ -\hat{X}_{ij} + x_i + w_j &\le 1. \end{aligned}$$

Removing redundant constraints and taking the dual of this SDP gives the Lemma.
□

APPENDIX C

Appendix to Chapter 4

C.1 Proofs

C.1.1 Proof of Proposition 4.4.1

The customer of type θ will choose to buy if and only if their expected utility from buying is non-negative. The uncertainty involved in the customer's decision is the valuation uncertainty ϵ, which is assumed to be independent of their type θ. Also, the value that the return window $T^{\pi}(\theta)$ takes is either T_S or T_L which is independent of θ.

Let K denote the return disutility imposed on the customer. Thus, a customer will buy if she observes a disutility that satisfies:

$$U(K) := \mathbb{E}_{\epsilon}\left[\max\left(V + \epsilon - p,\ -K\right)\right] \geq 0 \tag{C.1}$$

We see that $U(K)$ is decreasing in K, and is independent of θ. Thus, there exists a threshold $\bar{K}$ such the customer observing a return disutility of K will buy if and only if $K \leq \bar{K}$.

Note that $\bar{K}$ can be found as the highest value of K that satisfies Equation C.1. Also, $\bar{K}$ is finite, as if $\bar{K} = \infty$, we have:

$$U(\bar{K}) := \mathbb{E}_{\epsilon}\left[\max\left(V + \epsilon - p,\ -\infty\right)\right] = \mathbb{E}_{\epsilon} \max\left(V + \epsilon - p\right) = V - p < 0$$

Thus, $\bar{K}$ can be found by solving:

$$U(\bar{K}) = \mathbb{E}_{\epsilon}\left[\max\left(V + \epsilon - p,\ -\bar{K}\right)\right] = 0 \tag{C.2}$$

□

C.1.2 Proof of Propostion 4.4.2

The profit function is given by:

$$\Pi(y,\pi) = (s-c)\cdot y + (p-\tilde{p}^{\pi}-s)\cdot \mathbb{E}\min\left(y,\xi^{\pi}D\right) \tag{C.3}$$
$$= (p-\tilde{p}^{\pi}-c)\cdot y - (p-\tilde{p}^{\pi}-s)\cdot \mathbb{E}\left(y-\xi^{\pi}D\right)^{+} \tag{C.4}$$
$$= (p-\tilde{p}^{\pi}-c)\cdot y - (p-\tilde{p}^{\pi}-s)\cdot \int\limits_{0}^{\frac{y}{\xi^{\pi}}} \left(y-\xi^{\pi}D\right) f(D)dD \tag{C.5}$$

Given π, the first and second order differentials with respect to y are:

$$\frac{d\Pi(y,\pi)}{dy} = p-\tilde{p}^{\pi}-c-(p-\tilde{p}-s)\cdot F\left(\frac{y}{\xi^{\pi}}\right)$$
$$\frac{d^2\Pi(y,\pi)}{dy^2} = -(p-\tilde{p}-s)\cdot f\left(\frac{y}{\xi^{\pi}}\right)$$

It is clear that $\frac{d^2\Pi(y,\pi)}{dy^2}\le 0$, and hence $\Pi(y,\pi)$ is concave in y, and the first order conditions yield the optimal $y^*(\pi)$ given in the Proposition. $\square$

C.1.3 Proof of Proposition 4.5.1

Proof of i): Consider the case where $s'=s$. For any value of r_{SB}, we can choose $r_{LB} = \bar{K}\cdot\left(1-\frac{T_S}{T_L}\right)+r_{SB}\cdot\frac{T_S}{T_L}$ such that from Propostion 4.5.3, we have $\Omega_S^{SB}=\Omega_L^{LB}$, and hence, $\xi^{SB}=\xi^{LB}$. Also, for any $\theta\in\Omega_S^{SB}$,

$$\begin{aligned}\frac{\theta}{T_L}+r_{LB}-\left(\frac{\theta}{T_S}+r_{SB}\right) &= \theta\cdot\left(\frac{1}{T_L}-\frac{1}{T_S}\right)+r_{LB}-r_{SB}\\ &= \theta\cdot\left(\frac{1}{T_L}-\frac{1}{T_S}\right)+\left(\bar{K}-r_{SB}\right)\cdot\left(1-\frac{T_S}{T_L}\right)\\ &= \left(\frac{1}{T_L}-\frac{1}{T_S}\right)\cdot\left(\left(\bar{K}-r_{SB}\right)T_S-\theta\right)\\ &\ge 0\end{aligned}$$

where the second equality follows by substituting the value of r_{LB}, and the final inequality follows from Proposition 4.5.3. Thus, we have:

$$\begin{aligned}\psi^{LB} &= \mathbb{E}_{\theta}\left[\mathcal{G}\left(p-V-\frac{\theta}{T_L}-r_{LB}\right)\mid\theta\in\Omega_L^{LB}\right]\\ &\le \mathbb{E}_{\theta}\left[\mathcal{G}\left(p-V-\frac{\theta}{T_S}-r_{SB}\right)\mid\theta\in\Omega_S^{SB}\right]\\ &= \psi^{SB}\end{aligned}$$

We have shown that for any chosen r_{SB}, we can find a value of r_{LB} such that $\Pi^{SB}(y,\pi^{SB}) \leq \Pi^{LB}(y,\pi^{LB})$. Thus, $\max\limits_{y,\pi^{SB}} \Pi^{SB}(y,\pi^{SB}) \leq \Pi^{LB}(\bar{y},\bar{\pi}^{LB})$ for some $\bar{y}$ and $\bar{\pi}^{LB}$, and we also have: $\Pi^{LB}(\bar{y},\bar{\pi}^{LB}) \leq \max\limits_{y,\pi^{LB}} \Pi^{LB}(y,\pi^{LB})$.

Hence, we have $\Delta\Pi \leq 0$ when $s' = s$, and hence the firm should offer the short-blanket policy. By the continuity of $\Delta\Pi$, this is also true for values of s' close to s. □

Proof of ii): The difference in the optimal profits between the short-blanket and long-blanket policies is given by:

$$\Delta\Pi = \max_{y,\pi^{SB}} \Pi^{SB}(y,\pi^{SB}) - \max_{y,\pi^{LB}} \Pi^{LB}(y,\pi^{LB}) \tag{C.6}$$

It is easy to see from Equation 4.19 that $\tilde{p}^{SB}$ is continuous and decreasing in s', and hence $\Pi^{SB}(y,\pi^{SB})$ is increasing in s'. Since Π^{LB} is independent of s', this implies that $\Delta\Pi$ is increasing in s'. □

C.1.4 Proof of Proposition 4.5.2

We have:

$$y^*(r) = \xi(r) \cdot F^{-1}\left(1 - \frac{c-s}{p-s-\tilde{p}(r)}\right) \tag{C.7}$$

where

$$\begin{aligned}
\xi(r) &= \frac{(\bar{K}-r)T}{\bar{\theta}} \\
\tilde{p}(r) &= (p-s'-r) \cdot \left(p - V - \underline{\epsilon} - \frac{\bar{K}}{2} - \frac{r}{2}\right)
\end{aligned}$$

Dropping the dependence on r for convenience, we have: $y^*(r) := \xi \cdot \alpha$, where we define $\alpha = F^{-1}\left(1 - \frac{c-s}{p-s-\tilde{p}(r)}\right)$. We have:

$$\begin{aligned}
\frac{d\xi}{dr} &= -\frac{T}{\bar{\theta}} \\
\frac{d\tilde{p}}{dr} &= -\frac{1}{\bar{\epsilon}-\underline{\epsilon}} \cdot \left(\frac{1}{2}\cdot(p-s'-r) + p - V - \underline{\epsilon} - \frac{r}{2} - \frac{\bar{K}}{2}\right)
\end{aligned}$$

We also have:

$$\begin{aligned}
F(\alpha) &= 1 - \frac{c-s}{p-\tilde{p}-s} \\
\frac{dF(\alpha)}{d\alpha}\cdot\frac{d\alpha}{dr} &= -\frac{(c-s)\cdot\frac{d\tilde{p}}{dr}}{(p-s-\tilde{p})^2} \\
\frac{d\alpha}{dr} &= -\frac{(c-s)\cdot\frac{d\tilde{p}}{dr}}{(p-s-\tilde{p})^2\cdot f(\alpha)}
\end{aligned}$$

Given a return window T, the expected profit function only in terms of r is given by:

$$\Pi(r) = (s-c)\cdot(\xi\alpha) + (p-s-\tilde{p})\cdot\mathbb{E}\min(\xi\alpha,\xi D) \tag{C.8}$$
$$= (p-c-\tilde{p})\cdot\xi\cdot\alpha - (p-s-\tilde{p})\cdot\xi\cdot\mathbb{E}(\alpha-D)^{+} \tag{C.9}$$

Differentiating the profit function in Equation C.9, we have:

$$\begin{aligned}
\frac{d\Pi}{dr} &= \left[(p-c-\tilde{p})\left(-\frac{T}{\bar{\theta}}\right) + \xi\cdot\left(-\frac{d\tilde{p}}{dr}\right)\right]\cdot\alpha \; + (p-c-\tilde{p})\cdot\xi\cdot\frac{d\alpha}{dr} \\
&\quad - \left[(p-s-\tilde{p})\left(-\frac{T}{\bar{\theta}}\right) + \xi\cdot\left(-\frac{d\tilde{p}}{dr}\right)\right]\cdot\mathbb{E}(\alpha-D)^{+} \; - (p-s-\tilde{p})\cdot\xi\cdot\frac{d\alpha}{dr}\cdot F(\alpha) \\
&= -\left[(p-c-\tilde{p})\left(\frac{T}{\bar{\theta}}\right) + \xi\cdot\left(\frac{d\tilde{p}}{dr}\right)\right]\cdot\alpha + \left[(p-s-\tilde{p})\left(\frac{T}{\bar{\theta}}\right) + \xi\cdot\left(\frac{d\tilde{p}}{dr}\right)\right]\cdot\mathbb{E}(\alpha-D)^{+} \\
&\quad + \frac{d\alpha}{dr}\cdot\xi\cdot(p-c-\tilde{p}-(p-s-\tilde{p})\cdot F(\alpha)) \\
&= -\frac{\Pi(r)}{(\bar{K}-r)} - \left[\xi\cdot\left(\frac{d\tilde{p}}{dr}\right)\right]\cdot\alpha + \left[\xi\cdot\left(\frac{d\tilde{p}}{dr}\right)\right]\cdot\mathbb{E}(\alpha-D)^{+} \\
&\quad + \frac{d\alpha}{dr}\cdot\xi\cdot\left(p-c-\tilde{p}-(p-s-\tilde{p})\cdot\left(\frac{p-c-\tilde{p}}{p-s-\tilde{p}}\right)\right)
\end{aligned}$$

Simplifying, we get:

$$\frac{d\Pi}{dr} = -\frac{\Pi(r)}{(\bar{K}-r)} + \left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D) \tag{C.10}$$

Now, the second order differential of the profit function is given by differentiating Equation C.10:

$$\begin{aligned}
\frac{d^2\Pi}{dr^2} &= -\frac{(\bar{K}-r)\frac{d\Pi}{dr} - \Pi(r)(-1)}{(\bar{K}-r)^2} + \frac{d}{dr}\left(\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D)\right) &\text{(C.11)} \\
&= -\frac{(\bar{K}-r)\cdot\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D)}{(\bar{K}-r)^2} + \frac{d}{dr}\left(\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D)\right) \\
&= \frac{1}{\bar{K}-r}\cdot\left[-\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D) + (\bar{K}-r)\cdot\frac{d}{dr}\left(\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D)\right)\right] \\
\frac{d^2\Pi}{dr^2} &= \frac{1}{\bar{K}-r}\cdot\frac{d}{dr}\left[(\bar{K}-r)\cdot\left(-\frac{d\tilde{p}}{dr}\right)\cdot\mathbb{E}\min(\xi\alpha,\xi D)\right] &\text{(C.12)}
\end{aligned}$$

The second equality follows from Equation C.10. Note that $\Pi(\bar{K}) = 0$, as there is no demand. Thus, we will assume that $\Pi(r)$ cannot be monotone increasing, as it is then optimal to not sell the product. Thus, $\Pi(r)$ is either monotone decreasing, or there exists at least one interior point r^* where $\frac{d\Pi}{dr} = 0$. For any such r^*, we have:

$$
\begin{aligned}
&-\frac{\Pi(r^*)}{(\bar{K}-r^*)} - \xi \cdot \left(\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*} \right) \cdot \mathbb{E}\min(\alpha, D) = 0 \\
&-\frac{\Pi(r^*)}{(\bar{K}-r^*)} - \xi \cdot \left(\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*} \right) \cdot \left(\frac{\frac{\Pi(r^*)}{\xi} - (s-c)\alpha}{p-s-\tilde{p}} \right) = 0 \\
&\Pi(r^*) = \frac{(s-c)\xi\alpha \cdot (\bar{K}-r^*) \cdot \left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}}{p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}} \\
&(s-c)\cdot(\cancel{\xi}\alpha) + (p-s-\tilde{p}) \cdot \mathbb{E}\min(\cancel{\xi}\alpha, \cancel{\xi}D) = \frac{(s-c)\cancel{\xi}\alpha \cdot (\bar{K}-r^*) \cdot \left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}}{p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}} \\
&(p-s-\tilde{p}) \cdot \mathbb{E}\min(\alpha, D) = (s-c)\alpha \cdot \left(\frac{(\bar{K}-r^*) \cdot \left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}}{p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}} - 1 \right) \\
&\cancel{(p-s-\tilde{p})} \cdot \mathbb{E}\min(\alpha, D) = -(s-c)\alpha \cdot \left(\frac{\cancel{p-s-\tilde{p}}}{p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*}} \right) \\
&\left(p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*} \right) \cdot \mathbb{E}\min(\alpha, D) = -(s-c)\alpha \\
&\left(p-s-\tilde{p} + (\bar{K}-r)\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*} \right) \cdot \mathbb{E}\min(\xi\alpha, \xi D) = -(s-c)\xi\alpha \qquad \text{(C.13)} \\
&\Pi(r^*) = (\bar{K}-r)\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=r^*} \right) \cdot \mathbb{E}\min(\xi\alpha, \xi D) \qquad \text{(C.14)}
\end{aligned}
$$

Note that from Equation C.14, $\Pi(r^*) \geq 0$. For any such r^*, if we show that $\left.\frac{d^2\Pi}{dr^2}\right|_{r^*} \leq 0$, then $\Pi(r)$ is unimodal. From Equation C.12, it is clear that if $\mathbb{E}\min(\xi\alpha, \xi D)$ is decreasing in r at $r = r^*$, we will have $\left.\frac{d^2\Pi}{dr^2}\right|_{r^*} \leq 0$.

We have:

$$
\begin{aligned}
&\frac{d\left(\mathbb{E}\min\left(\xi\alpha,\xi D\right)\right)}{dr}\\
&=\frac{d}{dr}\left(\frac{\Pi(r)-(s-c)\cdot(\xi\alpha)}{p-s-\tilde{p}}\right)\\
&=\frac{1}{p-s-\tilde{p}}\cdot\left(\frac{d\Pi}{dr}+(c-s)\frac{d(\xi\alpha)}{dr}\right)+\left(\Pi(r)-(s-c)\cdot(\xi\alpha)\right)\cdot\left(-\frac{1}{(p-s-\tilde{p})^2}\right)\cdot\left(-\frac{d\tilde{p}}{dr}\right)\\
&=\frac{1}{p-s-\tilde{p}}\cdot\left[\frac{d\Pi}{dr}+(c-s)\frac{d(\xi\alpha)}{dr}+\mathbb{E}\min\left(\xi\alpha,\xi D\right)\cdot\left(\frac{d\tilde{p}}{dr}\right)\right]\\
&=\frac{1}{p-s-\tilde{p}}\cdot\left[-\frac{\Pi(r)}{(\bar{K}-r)}+(c-s)\frac{d(\xi\alpha)}{dr}\right]
\end{aligned}
$$

The first equality follows from Equation C.8, the second equality follows from differentiation, the third equality again follows from C.8, and the fourth equality follows from Equation C.10. Since $\Pi(r)\geq 0$ and $r\leq\bar{K}$, if we have $\left.\frac{d(\xi\alpha)}{dr}\right|_{r^*}\leq 0$, we are done. □

Now, if $\Pi(r)$ is either monotone decreasing or unimodal, then $r=0$ is optimal whenever $\left.\frac{d\Pi}{dr}\right|_{r=0}\leq 0$. let ξ_0, α_0 denote the values of ξ and α when $r=0$ respectively.

From (C.10), we have:

$$
\begin{aligned}
\left.\frac{d\Pi}{dr}\right|_{r=0}&=-\frac{\Pi(0)}{\bar{K}}+\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\cdot\mathbb{E}\min\left(\xi_0\alpha_0,\xi_0 D\right) &\text{(C.15)}\\
&=\frac{1}{\bar{K}}\cdot\left((c-s)\xi_0\alpha_0-\left(p-s-\tilde{p}(0)-\bar{K}\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\right)\cdot\mathbb{E}\min\left(\xi_0\alpha_0,\xi_0 D\right)\right) &\text{(C.16)}
\end{aligned}
$$

Whenever $c-s\geq\left(p-s-\tilde{p}(0)-\bar{K}\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\right)$, we have:

$$
\begin{aligned}
\left.\frac{d\Pi}{dr}\right|_{r=0}&\geq\frac{c-s}{\bar{K}}\cdot\left(\xi_0\alpha_0-\mathbb{E}\min\left(\xi_0\alpha_0,\xi_0 D\right)\right)\\
&\geq\frac{c-s}{\bar{K}}\cdot\left(\xi_0\alpha_0-\xi_0\alpha_0\right)\\
&=0
\end{aligned}
$$

Thus, whenever $c\geq\left(p-\tilde{p}(0)-\bar{K}\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\right)$, the function $\Pi(r)$ is unimodal with an interior solution $r^*>0$.

Consider the case where $c\to s$. In this case, $\alpha_0\to\infty$, hence:

$$
\left.\frac{d\Pi}{dr}\right|_{r=0}=-\frac{\xi_0\cdot\mathbb{E}(D)}{\bar{K}}\cdot\left(p-s-\tilde{p}(0)-\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\right) \tag{C.17}
$$

Thus, only when $p-s-\tilde{p}(0)-\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)\geq 0$, or $\tilde{p}(0)\leq p-s-\left(-\left.\frac{d\tilde{p}}{dr}\right|_{r=0}\right)$, we have $\left.\frac{d\Pi}{dr}\right|_{r=0}\leq 0$, and the optimal solution is $r=0$. Otherwise, we will have $\left.\frac{d\Pi}{dr}\right|_{r=0}\geq 0$, and an interior point $r^*>0$ will be optimal. □

C.1.5 Proof of Proposition 4.5.3

Any customer of type θ will choose the option that yields the maximum expected utility. By Proposition 4.4.1, we know that the customer will choose to buy only if $\theta \leq max\left(\left(\bar{K}-r_S\right)T_S, \left(\bar{K}-r_L\right)T_L\right) = \left(\bar{K}-r_L\right)T_L = \theta_{M,L}$. Thus, for $\theta \leq \theta_{M,L}$, the expected utility is:

$$\max\left(\mathbb{E}_\epsilon \max\left[V+\epsilon-p, -\frac{\theta}{T_S}-r_S\right],\ \mathbb{E}_\epsilon \max\left[V+\epsilon-p,\ -\frac{\theta}{T_L}-r_L\right]\right)$$

Clearly, the choice between (T_S, r_S) and (T_L, r_L) only depends on the return disutility imposed on the customer by these policies. Thus, customer of type $\theta \leq \theta_{M,L}$ will choose (T_S, r_S) if and only if:

$$\begin{aligned} & \frac{\theta}{T_S}+r_S \leq \frac{\theta}{T_L}+r_S \\ \Rightarrow \qquad & \theta \leq \frac{r_L-r_S}{\frac{1}{T_S}-\frac{1}{T_L}} = \theta_{M,SL} \end{aligned}$$

and the customer chooses (T_L, r_L) for $\theta_{M,SL} < \theta \leq \theta_{M,L}$. □

C.1.6 Proof of Proposition 4.5.4

Let:

$$\begin{aligned} r^*_{SB} &= \arg\max_{r_S\in[0,\bar{K}]} \Pi(y, \pi^{SB}) \\ r^*_{LB} &= \arg\max_{r_L\in[0,\bar{K}]} \Pi(y, \pi^{LB}) \\ (r_S^{*,M}, r_L^{*,M}) &= \arg\max_{\substack{r_S\in[0,\bar{K}],\\ r_L\in\left[r_S, r_S\cdot\frac{T_S}{T_L}+\bar{K}\cdot\left(1-\frac{T_S}{T_L}\right)\right]}} \Pi(y, \pi^M) \end{aligned}$$

Let $\pi^{*,SB}$, $\pi^{*,LB}$ and $\pi^{*,M}$ denote the optimal policies respectively. The proof follows by showing that both r^*_{SB} and r^*_{LB} can be implemented under π^M, and hence are feasible solutions to the menu of policies.

Consider a feasible policy π_1^M under the menu of policies such that $r_S = r^*_{SB}$ and $r_L = r_{SB}\cdot\frac{T_S}{T_L}+\bar{K}\cdot\left(1-\frac{T_S}{T_L}\right)$. We thus have: $\Omega_L^{\pi_1^M} = \emptyset$ since $\theta_{M,S} = \theta_{M,SL} = \theta_{M,L}$. Also, $\Omega_S^{\pi_1^M} = \Omega_S^{\pi^{*,SB}}$, and each buying customer in both policies is offered the same return window and return fee (T_S, r^*_{SB}). Hence, $\Pi^{SB}(y) = \Pi(y, \pi_1^M) \leq \Pi^M(y)$.

Similarly, consider another feasible policy π_2^M under the menu of policies such that $r_S = r_L = r^*_{LB}$. In this case, we have: $\Omega_S^{\pi_2^M} = \emptyset$ since $\theta_{M,SL} = 0$, and we also have $\Omega_L^{\pi_2^M} = \Omega_L^{\pi^{*,LB}}$. Since each buying customer in both policies is offered the same return window and fee (T_L, r^*_{LB}), we have $\Pi^{LB}(y) = \Pi(y, \pi_2^M) \leq \Pi^M(y)$.

Thus, we have: $\Pi^M(y) \geq \max\left(\Pi^{SB}(y), \Pi^{LB}(y)\right).$ $\square$

C.1.7 Proof of Proposition 4.6.1

We begin the proof by noticing that, keeping ξ_S^π and ξ_L^π constant, the expected profit $\Pi(y,\pi)$ is decreasing in ψ_S^π and ψ_L^π. We have:

$$\xi_S^\pi = \frac{1}{\bar{\theta}} \cdot \int\limits_{\theta \in \Omega_S^\pi} d\theta$$

$$\xi_L^\pi = \frac{1}{\bar{\theta}} \cdot \int\limits_{\theta \in \Omega_L^\pi} d\theta$$

$$\xi_S^\pi \psi_S^\pi = \frac{1}{\bar{\theta}(\bar{\epsilon}-\underline{\epsilon})} \cdot \int\limits_{\theta \in \Omega_S^\pi} \left(p - V - \underline{\epsilon} - \frac{\theta}{T_S}\right) d\theta$$

$$\xi_L^\pi \psi_L^\pi = \frac{1}{\bar{\theta}(\bar{\epsilon}-\underline{\epsilon})} \cdot \int\limits_{\theta \in \Omega_L^\pi} \left(p - V - \underline{\epsilon} - \frac{\theta}{T_L}\right) d\theta$$

$$\xi_S^\pi \psi_S^\pi + \xi_L^\pi \psi_L^\pi = \frac{1}{\bar{\theta}(\bar{\epsilon}-\underline{\epsilon})} \cdot \left[\int\limits_{\theta \in \Omega_S^\pi \cup \Omega_L^\pi} \left(p - V - \underline{\epsilon} - \frac{\theta}{T_S}\right) d\theta + \int\limits_{\theta \in \Omega_L^\pi} \theta \left(\frac{1}{T_S} - \frac{1}{T_L}\right) d\theta\right]$$

Proof of i): For any $\theta \leq \bar{K}T_S$, if $T^{*,F}(\theta) = T_\emptyset$, then for any other $\theta' < \theta$, $T^{*,F}(\theta') = T_\emptyset$.

Consider an optimal policy $\pi^{*,F}$ where $T^{*,F}(\theta) = T_\emptyset$ for some $\theta \leq \bar{K}T_S$, and $T^{*,F}(\theta') = T_S$ for some $\theta' < \theta$. The proof for the case where $T^{*,F}(\theta') = T_L$ follows similarly.

Consider an alternative policy $\hat{\pi}$ that is identical to $\pi^{*,F}$ except that $\hat{T}(\theta') = T_\emptyset$, and $\hat{T}(\theta) = T_S$. Thus, $\xi_S^{\pi^{*,F}} = \xi_S^{\hat{\pi}}$, and $\xi^{\pi^{*,F}} = \xi^{\hat{\pi}}$. However, $\xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} > \xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}}$ since $\theta' < \theta$. Thus, we have $\psi_S^{\pi^{*,F}} > \psi_S^{\hat{\pi}}$, implying $\Pi(y, \pi^{*,F}) \leq \Pi(y, \hat{\pi})$, and thus $\hat{\pi}$ is also optimal.

Proof of ii): For any $\theta \leq \bar{K}T_L$, if $T^{*,F}(\theta) = T_S$, then for any other $\theta' < \theta$, $T^{*,F}(\theta') = T_S$ or $T^{*,F}(\theta') = T_\emptyset$.

Consider an optimal policy $\pi^{*,F}$ where $T^{*,F}(\theta) = T_S$ for some $\theta \leq \bar{K}T_L$, and $T^{*,F}(\theta') = T_L$ for some $\theta' < \theta$. There are three cases that ensue depending on the values of θ' and θ.

Case 1: $\theta' < \theta \leq \bar{K}T_S$: Consider an alternative policy $\hat{\pi}$ that is identical to $\pi^{*,F}$ except that $\hat{T}(\theta') = T_S$. This ensures that $\xi^{\pi^{*,F}} = \xi^{\hat{\pi}}$. Since $\theta' < \theta$, we also have

$$\xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} < \xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}}$$
$$\xi_L^{\pi^{*,F}} \psi_L^{\pi^{*,F}} > \xi_L^{\hat{\pi}} \psi_L^{\hat{\pi}}$$
$$\xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} + \xi_L^{\pi^{*,F}} \psi_L^{\pi^{*,F}} > \xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}} + \xi_L^{\hat{\pi}} \psi_L^{\hat{\pi}}$$

Let $\xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}} - \xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} = \Delta_1$, and $\xi_L^{\pi^{*,F}} \psi_L^{\pi^{*,F}} - \xi_L^{\hat{\pi}} \psi_L^{\hat{\pi}} = \Delta_2$. Thus, $0 < \Delta_1 < \Delta_2$. Hence, we have:

$$
\begin{aligned}
\tilde{p}^{\pi^{*,F}} - \tilde{p}^{\hat{\pi}} &= \frac{1}{\xi^{\hat{\pi}}} \cdot \big[(p - s')(-\Delta_1) + (p - s)(\Delta_2)\big] \\
&\geq \frac{p - s'}{\xi^{\hat{\pi}}} \cdot [\Delta_2 - \Delta_1] \\
&\geq 0
\end{aligned}
$$

Hence, $\Pi(y, \pi^{*,F}) \leq \Pi(y, \hat{\pi})$, and thus $\hat{\pi}$ is also optimal.

Case 2: $\theta' \leq \bar{K}T_S < \theta$. Consider an alternative policy $\hat{\pi}$ that is identical to $\pi^{*,F}$ except that $\hat{T}(\theta') = T_\emptyset$, and $\hat{T}(\theta) = T_L$. This ensures that $\xi_S^{\pi^{*,F}} = \xi_S^{\hat{\pi}}$, and $\xi^{\pi^{*,F}} = \xi^{\hat{\pi}}$. It is also clear that since $\theta' < \theta$, we have $\psi_L^{\pi^{*,F}} > \psi_L^{\hat{\pi}}$, with $\psi_S^{\pi^{*,F}} > \psi_S^{\hat{\pi}}$. Hence, $\Pi(y, \pi^{*,F}) \leq \Pi(y, \hat{\pi})$, and thus $\hat{\pi}$ is also optimal.

Case 3: $\bar{K}T_S < \theta' < \theta$. Consider an alternative policy $\hat{\pi}$ that is identical to $\pi^{*,F}$ except that $\hat{T}(\theta') = T_S$, and $\hat{T}(\theta) = T_L$. This ensures that $\xi_S^{\pi^{*,F}} = \xi_S^{\hat{\pi}}$, and $\xi^{\pi^{*,F}} = \xi^{\hat{\pi}}$. It is also clear that since $\theta' < \theta$, we have $\psi_L^{\pi^{*,F}} > \psi_L^{\hat{\pi}}$, with $\psi_S^{\pi^{*,F}} > \psi_S^{\hat{\pi}}$. Hence, $\Pi(y, \pi^{*,F}) \leq \Pi(y, \hat{\pi})$, and thus $\hat{\pi}$ is also optimal. $\square$

Proof of iii): If $\theta \leq \bar{K}T_S$, $T^{*,F}(\theta) \neq T_L$.

Let there exist an optimal policy $\pi^{*,F}$ where $T^{*,F}(\theta) = T_L$ for some $\theta \leq \bar{K}T_S$. Consider an alternative policy $\hat{\pi}$ that is identical to $\pi^{*,F}$ except that $\hat{T}(\theta) = T_S$. Then, we have:

$$
\begin{aligned}
\xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} &< \xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}} \\
\xi_L^{\pi^{*,F}} \psi_L^{\pi^{*,F}} &> \xi_L^{\hat{\pi}} \psi_L^{\hat{\pi}} \\
\xi_S^{\pi^{*,F}} \psi_S^{\pi^{*,F}} + \xi_L^{\pi^{*,F}} \psi_L^{\pi^{*,F}} &> \xi_S^{\hat{\pi}} \psi_S^{\hat{\pi}} + \xi_L^{\hat{\pi}} \psi_L^{\hat{\pi}}
\end{aligned}
$$

The rest of the proof follow the same steps as the proof of Case 1 of Part ii). $\square$

C.1.8 Proof of Lemma 4.6.1

The expected return fees are:

$$
\begin{aligned}
R_S &= \mathbb{E}\left[r^\pi(\theta) \;\middle|\; \theta \in \Omega_S,\ V + \epsilon - p < -\frac{\theta}{T_S} - r^\pi(\theta) \right] \\
&= \frac{1}{\psi_S} \cdot \mathbb{E}\left[r^\pi(\theta) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^\pi(\theta) \right) \;\middle|\; \theta \in \Omega_S \right] \qquad \text{(C.18)} \\
R_L &= \mathbb{E}\left[r^\pi(\theta) \;\middle|\; \theta \in \Omega_L,\ V + \epsilon - p < -\frac{\theta}{T_L} - r^\pi(\theta) \right] \\
&= \frac{1}{\psi_L} \cdot \mathbb{E}\left[r^\pi(\theta) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^\pi(\theta) \right) \;\middle|\; \theta \in \Omega_L \right] \qquad \text{(C.19)}
\end{aligned}
$$

We have $\tilde{p}^{\pi} = \frac{1}{\xi^{\pi}} \cdot \left[\xi_S \left((p - s')\psi_S - R_S\psi_S\right) + \xi^L \left((p - s)\psi_L - R_L\psi_L\right)\right]$. Considering each term, we have:

$$
\begin{aligned}
&(p - s')\psi_S - R_S\psi_S \\
=\quad & (p - s') \cdot \mathbb{E}\left[\mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right) \mid \theta \in \Omega_S\right] \\
-\quad & \mathbb{E}\left[r^{\pi}(\theta) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right) \mid \theta \in \Omega_S\right] \\
=\quad & \mathbb{E}\left[\left(p - s' - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right) \middle| \, \theta \in \Omega_S\right] && \text{(C.20)} \\
&(p - s)\psi_L - R_L\psi_L \\
=\quad & (p - s) \cdot \mathbb{E}\left[\mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right) \mid \theta \in \Omega_L\right] \\
-\quad & \mathbb{E}\left[r^{\pi}(\theta) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right) \mid \theta \in \Omega_L\right] \\
=\quad & \mathbb{E}\left[\left(p - s - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right) \middle| \, \theta \in \Omega_L\right] && \text{(C.21)}
\end{aligned}
$$

It is easy to see that for any θ, both $(p - s - r^{\pi}(\theta)) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right)$ as well as $(p - s' - r^{\pi}(\theta)) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right)$ are decreasing in $r^{\pi}(\theta)$. Since $\Omega^{S,\pi_1} = \Omega^{S,\pi_2}$ and $\Omega^{L,\pi_1} = \Omega^{L,\pi_2}$, this implies that $\tilde{p}^{\pi_1} \leq \tilde{p}^{\pi_2}$, which completes the proof. $\square$

C.1.9 Proof of Proposition 4.6.2

The proofs follow by constructing an optimal policy that possesses the properties in Proposition 4.6.2 from another optimal policy that does not possess these properties. *Proof of i):* Let π^* be an optimal policy where there exists $\bar{\Omega} = \left\{\theta \,:\, -\frac{\theta}{T^*(\theta)} - r^*(\theta) \neq -\bar{K}\right\} \neq \emptyset$ such that $\bar{\Omega} \subset \Omega$. Consider policy π_1 which is identical to π^* except that $r^*(\theta) = \bar{K} - \frac{\theta}{T^*(\theta)}$ for all $\theta \in \bar{\Omega}$. By Lemma 4.6.1, it follows that $\tilde{p}^{\pi^*} \geq \tilde{p}^{\pi_1}$ and consequently $\Pi(y, \pi^*) \leq \Pi(y, \pi_1)$, which implies that π_1 is also optimal.

Proof of ii): From (i), we can update the definition of Ω_S and Ω_L as:

$$\Omega_S^{\pi^*} = \left\{\theta \,:\, T^*(\theta) = T_S,\, -\frac{\theta}{T_S} - r^*(\theta) = -\bar{K}\right\} \tag{C.22}$$

$$\Omega_L^{\pi^*} = \left\{\theta \,:\, T^*(\theta) = T_L,\, -\frac{\theta}{T_L} - r^*(\theta) = -\bar{K}\right\} \tag{C.23}$$

From Equations C.20 and C.21, for $\theta \sim U[0, \bar{\theta}]$, we have:

$$\begin{aligned}
\tilde{p}^{\pi} &= \frac{1}{\xi^{\pi}} \cdot \left(\mathbb{E}\left[\left(p - s' - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right) \mathbb{1}\left(\theta \in \Omega_S^{\pi}\right) \right] \right. \\
&\quad \left. + \mathbb{E}\left[\left(p - s - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right) \mathbb{1}\left(\theta \in \Omega_L^{\pi}\right) \right] \right) \\
&= \frac{1}{\bar{\theta} \cdot \xi^{\pi}} \cdot \left(\int\limits_{\theta \in \Omega_S^{\pi}} \left(p - s' - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r^{\pi}(\theta)\right) d\theta \right. \\
&\quad \left. + \int\limits_{\theta \in \Omega_L^{\pi}} \left(p - s - r^{\pi}(\theta)\right) \cdot \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r^{\pi}(\theta)\right) d\theta \right) \qquad \text{(C.24)}
\end{aligned}$$

By (i), for any optimal policy π^*, we have:

$$\begin{aligned}
\tilde{p}^{\pi^*} &= \frac{1}{\bar{\theta} \cdot \xi^{\pi^*}} \cdot \left(\int\limits_{\theta \in \Omega_S^{\pi^*}} \left(p - s' - r^*(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \right. \\
&\quad \left. + \int\limits_{\theta \in \Omega_L^{\pi^*}} \left(p - s - r^*(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \right) \\
&= \frac{\mathcal{G}\left(p - V - \bar{K}\right)}{\bar{\theta} \cdot \xi^{\pi^*}} \cdot \left(\int\limits_{\theta \in \Omega_S^{\pi^*}} \left(p - s' - r^*(\theta)\right) d\theta + \int\limits_{\theta \in \Omega_L^{\pi^*}} \left(p - s - r^*(\theta)\right) d\theta \right) \qquad \text{(C.25)}
\end{aligned}$$

It is clear from (i) that for any $\theta \leq \bar{K} \cdot T_S$ and $\theta \in \Omega^{\pi}$, two policies are possible: $\left(T_S, r_S = \bar{K} - \frac{\theta}{T_S}\right)$ and $\left(T_L, r_L = \bar{K} - \frac{\theta}{T_L}\right)$. Since the rate of return is constant under both policies ($= \mathcal{G}(p - V - \bar{K})$), the contribution to $\tilde{p}^{\pi^*}$ of both these policies are $p - s' - r_S$ and $p - s - r_L$ respectively. Due to Lemma 4.6.1, the policy that is chosen for θ is one that leads to the lowest $\tilde{p}^{\pi^*}$. Thus, (T_S, r_S) is chosen if and only if:

$$\begin{aligned}
p - s' - r_S &\leq p - s - r_L \\
-s' - \bar{K} + \frac{\theta}{T_S} &\leq -s - \bar{K} + \frac{\theta}{T_L} \\
\theta &\leq \frac{s' - s}{\frac{1}{T_S} - \frac{1}{T_L}}
\end{aligned}$$

□

Proof of iii): Note that for any $\theta > \bar{K} \cdot T_S$ and $\theta \in \Omega^{\pi}$, it must be the case that $T^*(\theta) = T_L$. Thus, a corollary of (ii) is that there is an optimal policy in which there exists a threshold

θ^{SL} such that $T^*(\theta) = T_S$ only if $\theta \leq \theta^{SL}$, where

$$\theta^{SL} = \min\left(\frac{s' - s}{\frac{1}{T_S} - \frac{1}{T_L}}, \bar{K} \cdot T_S\right)$$

Thus, it is sufficient to show that if for any $\theta_2 > \theta^{SL}$, $\theta_2 \in \Omega_L$, then any other θ_1 such that $\theta^{SL} < \theta_1 < \theta_2$ should also be in Ω_L.

Let $\theta_1 \notin \Omega_L$. This implies that $\theta_1 \notin \Omega$, as $\theta_1 > \theta^{SL}$. Consider an alternate strategy π' which is identical to π^* except that $\theta_2 \notin \Omega_L$, and $\theta_1 \in \Omega_L$. That is, $-\frac{\theta_1}{T_L} - r'(\theta_1) = -\bar{K} = -\frac{\theta_2}{T_L} - r^*(\theta_2)$. Hence, we have $r'(\theta_1) > r^*(\theta_2)$ and $\xi^{\pi'} = \xi^{\pi^*} := \xi$.

Let $\hat{\Omega}^L = \Omega^{L,\pi^*} - \{\theta_2\}$, and $\Omega^{L,\pi'} = \hat{\Omega}^L \cup \{\theta_1\}$. Thus,

$$\begin{aligned}\tilde{p}^{\pi^*} = \frac{1}{\bar{\theta} \cdot \xi} \cdot \Bigg(&\int\limits_{\theta \in \Omega_S} \left(p - s' - r^*(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \\ &+ \int\limits_{\theta \in \hat{\Omega}^L \cup \{\theta_2\}} \left(p - s - r^*(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \Bigg)\end{aligned}$$

$$\begin{aligned}\tilde{p}^{\pi'} = \frac{1}{\bar{\theta} \cdot \xi} \cdot \Bigg(&\int\limits_{\theta \in \Omega_S} \left(p - s' - r'(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \\ &+ \int\limits_{\theta \in \hat{\Omega}^L \cup \{\theta_1\}} \left(p - s - r'(\theta)\right) \cdot \mathcal{G}\left(p - V - \bar{K}\right) d\theta \Bigg)\end{aligned}$$

Since $p - s' - r'(\theta_1) < p - s' - r^*(\theta_2)$, we have $\tilde{p}^{\pi'} < \tilde{p}^{\pi^*}$, and hence $\Pi(y, \pi') \geq \Pi(y, \pi^*)$, which implies that π' is also optimal. $\square$

C.2 Customer Parameters under Return Policies

C.2.1 General Setting

$$
\begin{aligned}
\xi_S^\pi &= \frac{1}{\bar{\theta}} \cdot \int\limits_{\theta \in \Omega_S^\pi} d\theta \\
\xi_L^\pi &= \frac{1}{\bar{\theta}} \cdot \int\limits_{\theta \in \Omega_L^\pi} d\theta \\
\psi_S^\pi &= \frac{\int\limits_{\theta \in \Omega_S^\pi} \left(p - V - \underline{\epsilon} - \frac{\theta}{T_S} - r^\pi(\theta)\right) d\theta}{\xi_S^\pi \bar{\theta}(\bar{\epsilon} - \underline{\epsilon})} \\
\psi_L^\pi &= \frac{\int\limits_{\theta \in \Omega_L^\pi} \left(p - V - \underline{\epsilon} - \frac{\theta}{T_L} - r^\pi(\theta)\right) d\theta}{\xi_L^\pi \bar{\theta}(\bar{\epsilon} - \underline{\epsilon})} \\
R_S^\pi &= \frac{1}{\xi_S^\pi \psi_S^\pi \bar{\theta}(\bar{\epsilon} - \underline{\epsilon})} \cdot \int\limits_{\theta \in \Omega_S^\pi} r^\pi(\theta) \left(p - V - \underline{\epsilon} - \frac{\theta}{T_S} - r^\pi(\theta)\right) d\theta \\
R_L^\pi &= \frac{1}{\xi_L^\pi \psi_L^\pi \bar{\theta}(\bar{\epsilon} - \underline{\epsilon})} \cdot \int\limits_{\theta \in \Omega_L^\pi} r^\pi(\theta) \left(p - V - \underline{\epsilon} - \frac{\theta}{T_L} - r^\pi(\theta)\right) d\theta
\end{aligned}
$$

C.2.2 Blanket Policies

For a return window T and flat return fee r, we have:

$$
\begin{aligned}
\xi &= \frac{(\bar{K} - r)T}{\bar{\theta}} \\
\psi &= \left(p - V - \underline{\epsilon} - \frac{\bar{K}}{2} - \frac{r}{2}\right) \\
R &= r
\end{aligned}
$$

C.2.3 Menu of Policies

$$\xi_S = \mathcal{H}(\theta_{M,SL}) = \frac{\theta_{M,SL}}{\bar{\theta}} = \frac{r_L - r_S}{\bar{\theta}\left(\frac{1}{T_S} - \frac{1}{T_L}\right)}$$

$$\xi_L = \mathcal{H}(\theta_{M,L}) - \mathcal{H}(\theta_{M,SL}) = \frac{\theta_{M,L} - \theta_{M,SL}}{\bar{\theta}} = \frac{\bar{K}\left(\frac{T_L}{T_S} - 1\right) + r_S - r_L \cdot \frac{T_L}{T_S}}{\bar{\theta}\left(\frac{1}{T_S} - \frac{1}{T_L}\right)}$$

$$\begin{aligned}
\psi_S &= \frac{\int\limits_{\theta \le \theta_{M,SL}} \mathcal{G}\left(p - V - \frac{\theta}{T_S} - r_S\right) h(\theta) d\theta}{\mathcal{H}(\theta_{M,SL})} \\
&= \frac{1}{\bar{\epsilon} - \underline{\epsilon}} \cdot \left[p - V - \underline{\epsilon} - r_S - \frac{\theta_{M,SL}}{2T_S}\right] \\
\psi_L &= \frac{\int\limits_{\theta_{M,SL}}^{\theta_{M,L}} \mathcal{G}\left(p - V - \frac{\theta}{T_L} - r_L\right) h(\theta) d\theta}{\mathcal{H}(\theta_{M,L}) - \mathcal{H}(\theta_{M,SL})} \\
&= \frac{1}{\bar{\epsilon} - \underline{\epsilon}} \cdot \left[p - V - \underline{\epsilon} - r_L - \frac{\theta_{M,SL} + \theta_{M,L}}{2T_L}\right]
\end{aligned}$$

$$\begin{aligned}
R_S &= r_S \\
R_L &= r_L
\end{aligned}$$

C.2.4 Personalized Policies with General Refunds

The fraction of customers buying the product are:

$$\begin{aligned}
\xi_S^* &= H\left(\min(\theta^*, \theta_1^*)\right) = \frac{\min(\theta^*, \theta_1^*)}{\bar{\theta}} \\
\xi_L^* &= H(\theta^*) - H\left(\min(\theta^*, \theta_1^*)\right) = \frac{\theta^* - \min(\theta^*, \theta_1^*)}{\bar{\theta}} \\
\xi^{\pi^*} &= H(\theta^*) = \frac{\theta^*}{\bar{\theta}}
\end{aligned}$$

The conditional probability of return (probability of return given purchase) for the customers in Ω:

$$\begin{aligned}
\psi_S^* &= \mathcal{G}\left(p - V - \bar{K}\right) \\
\psi_L^* &= \mathcal{G}\left(p - V - \bar{K}\right)
\end{aligned}$$

Note that the rate of return is independent of θ, as all customers obtain the same utility upon return $(-\bar{K})$. The expected return fees collected by the firm R_S and R_L are given by:

$$
\begin{aligned}
R_S^* &= \bar{K} - \frac{1}{T_S} \cdot \mathbb{E}\left[\theta \mid \theta \le \min(\theta^*, \theta_1^*)\right] \\
&= \bar{K} - \frac{1}{2T_S} \cdot \min(\theta^*, \theta_1^*) \\
R_L^* &= \bar{K} - \frac{1}{T_L} \cdot \mathbb{E}\left[\theta \mid \min(\theta^*, \theta_1^*) < \theta \le \theta^*\right] \\
&= \bar{K} - \frac{1}{2T_L} \cdot \left(\theta^* + \min(\theta^*, \theta_1^*)\right)
\end{aligned}
$$

The effective revenue loss per unit is thus:

$$
\begin{aligned}
\tilde{p}^{\pi^*} &= \frac{1}{\xi^\pi} \cdot \left[\left(p - s' - R_S\right) \xi_S \psi_S + \left(p - s - R_L\right) \xi_L \psi_L\right] \\
&= \frac{\mathcal{G}(p - V - \bar{K})}{H(\theta^*)} \cdot \Bigg[\left(p - s' - \bar{K}\right) \xi_S + \left(p - s - \bar{K}\right) \xi_L \\
&\qquad + \frac{1}{T_S} \mathbb{E}\left[\theta \mid \theta \le \min(\theta^*, \theta_1^*)\right] \xi_S + \frac{1}{T_L} \mathbb{E}\left[\theta \mid \min(\theta^*, \theta_1^*) < \theta \le \theta^*\right] \xi_L \Bigg] \\
&= \mathcal{G}(p - V - \bar{K}) \cdot \left[p - s - \bar{K} - (s' - s) \frac{H\left(\min(\theta^*, \theta_1^*)\right)}{H\left(\theta^*\right)}\right] \\
&\quad + \frac{\mathcal{G}(p - V - \bar{K})}{H(\theta^*)} \cdot \left[\frac{1}{T_S} \mathbb{E}\left[\theta \cdot \mathbb{1}\left(\theta \le \min(\theta^*, \theta_1^*)\right)\right] + \frac{1}{T_L} \mathbb{E}\left[\theta \cdot \mathbb{1}\left(\min(\theta^*, \theta_1^*) < \theta \le \theta^*\right)\right]\right] \\
&= \mathcal{G}(p - V - \bar{K}) \cdot \left[p - s - \bar{K} - (s' - s) \frac{H\left(\min(\theta^*, \theta_1^*)\right)}{H\left(\theta^*\right)}\right] \\
&\quad + \frac{\mathcal{G}(p - V - \bar{K})}{H(\theta^*)} \cdot \left[\left(\frac{1}{T_S} - \frac{1}{T_L}\right) \cdot \mathbb{E}\left[\theta \cdot \mathbb{1}\left(\theta \le \min(\theta^*, \theta_1^*)\right)\right] + \frac{1}{T_L} \mathbb{E}\left[\theta \cdot \mathbb{1}\left(\theta \le \theta^*\right)\right]\right] \\
&= \mathcal{G}(p - V - \bar{K}) \cdot \left[p - s - \bar{K} - (s' - s) \frac{\min(\theta^*, \theta_1^*)}{\theta^*}\right] \\
&\quad + \frac{\mathcal{G}(p - V - \bar{K})}{\theta^*} \cdot \left[\left(\frac{1}{T_S} - \frac{1}{T_L}\right) \cdot \frac{\left(\min(\theta^*, \theta_1^*)\right)^2}{2} + \frac{1}{T_L} \frac{\left(\theta^*\right)^2}{2}\right]
\end{aligned}
$$

Hence,

$$
\tilde{p}^{\pi^*} = \begin{cases} \mathcal{G}(p - V - \bar{K}) \cdot \left(p - s' - \bar{K} + \frac{\theta^*}{2T_S}\right), & \text{if } \theta \le \theta_1 \\ \mathcal{G}(p - V - \bar{K}) \cdot \left(p - s - \bar{K} + \frac{\theta^*}{2T_L} - \frac{\theta_1^2}{2\theta^*} \cdot \left(\frac{1}{T_S} - \frac{1}{T_L}\right)\right), & \text{o/w} \end{cases}
$$

BIBLIOGRAPHY

James Abbey, Michael Ketzenberg, and Richard Metters. A more profitable approach to product returns. *MIT Sloan Management Rev.*, July 23, 2018. https://sloanreview.mit.edu/article/a-more-profitable-approach-to-product-returns/.

Jason Acimovic and Stephen C Graves. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing Service Oper. Management*, 17(1): 34–51, 2014.

Jason Acimovic and Stephen C Graves. Mitigating spillover in online retailing via replenishment. *Manufacturing Service Oper. Management*, 19(3):419–436, 2017.

Narendra Agrawal and Stephen A Smith. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.*, 43(6): 839–861, 1996.

Yalçın Akçay, Tamer Boyacı, and Dan Zhang. Selling with money-back guarantees: The impact on prices, quantities, and retail profitability. *Production Oper. Management*, 22(4):777–791, 2013.

Aydın Alptekinoğlu and Alex Grasas. When to carry eccentric products? Optimal retail assortment under consumer returns. *Production Oper. Management*, 23(5): 877–892, 2014.

Mehmet Sekip Altug and Tolga Aydinliyim. Counteracting strategic purchase deferrals: The impact of online retailers' return policy decisions. *Manufacturing Service Oper. Management*, 18(3):376–392, 2016.

Asim Ansari, Carl F Mela, and Scott A Neslin. Customer channel migration. *J. Marketing Res.*, 45(1):60–76, 2008.

Appriss Retail. Customer returns in the retail industry. 2018. Accessed August 18, 2019, https://appriss.com/retail/wp-content/uploads/sites/4/2018/12/AR3018_2018-Customer-Returns-in-the-Retail-Industry_Digital.pdf.

Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Proc. 30th Annual ACM Sympos. Theory Comput.* ACM, New York, 1998.

David R Bell, Santiago Gallino, and Antonio Moreno. Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Sci.*, 64(4):1629–1651, 2017.

Aharon Ben-Tal, Alexander Goryashko, Elana Guslitzer, and Arkadi Nemirovski. Adjustable robust solutions of uncertain linear programs. *Math. Programming*, 99(2): 351–376, 2004.

Oded Berman, Dmitry Krass, and M Mahdi Tajbakhsh. On the benefits of risk pooling in inventory management. *Production Oper. Management*, 20(1):57–71, 2011.

Marco Bijvank and Iris FA Vis. Lost-sales inventory theory: A review. *Eur. J. Oper. Res.*, 215(1):1–13, 2011.

Kostas Bimpikis and Mihalis G Markakis. Inventory pooling under heavy-tailed demand. *Management Sci.*, 62(6):1800–1813, 2015.

Alexander Bleier and Maik Eisenbeiss. Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Sci.*, 34(5):669–688, 2015.

Brightpearl. Banned? A returning problem. 2018. Accessed August 18, 2019, `https://info.brightpearl.com/banned`.

Erik Brynjolfsson, Yu Jeffrey Hu, and Mohammad S Rahman. Competing in the age of omnichannel retailing. *MIT Sloan Management Rev.*, 54(4):23, 2013.

Yeon-Koo Che. Customer return policies for experience goods. *J. Indust. Econom.*, pages 17–24, 1996.

Annie I-An Chen. *Large-scale optimization in online-retail inventory management.* PhD thesis, Massachusetts Institute of Technology, 2017.

Jian Chen, Youhua Chen, Mahmut Parlar, and Yongbo Xiao. Optimal inventory and admission policies for drop-shipping retailers serving in-store and online customers. *IIE Trans.*, 43(5):332–347, 2011.

Chicago Consulting. Ten best warehouse networks. 2013. Accessed August 18, 2019, `http://www.chicago-consulting.com/ten-best-warehouse-networks/`.

Chongwoo Choe, Stephen King, and Noriaki Matsushima. Pricing with cookies: Behavior-based price discrimination and spatial competition. *Management Sci.*, 64 (12):5669–5687, 2017.

Sunil Chopra and Peter Meindl. Supply Chain Management. Strategy, Planning and Operations. In *Das Summa Summarum des Management*, pages 265–275. Springer, 2007.

Mabel Chou, Melvyn Sim, and R So. A robust optimization framework for analyzing distribution systems with transshipment. Working Paper, National University of Singapore, Singapore, 2006.

Wujin Chu, Eitan Gerstner, and James D Hess. Managing dissatisfaction: How to decrease customer opportunism by partial refunds. *J. Service Res.*, 1(2):140–155, 1998.

Charles J Corbett and Kumar Rajaram. A generalization of the inventory pooling effect to nonnormal dependent demand. *Manufacturing Service Oper. Management*, 8(4):351–358, 2006.

Keith J Crocker and Paolo Letizia. Optimal policies for recovering the value of consumer returns. *Production Oper. Management*, 23(10):1667–1680, 2014.

Scott Davis, Eitan Gerstner, and Michael Hagerty. Money back guarantees in retailing: Matching products to consumer tastes. *J. Retailing*, 71(1):7–22, 1995.

Scott Davis, Michael Hagerty, and Eitan Gerstner. Return policies and the optimal level of "hassle". *J. Econom. Bus.*, 50(5):445–460, 1998.

Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.*, 58(3):595–612, 2010.

Levi DeValve, Yehua Wei, Di Wu, and Rong Yuan. Understanding the value of fulfillment flexibility in an online retailing environment. 2018. Working Paper, Duke University.

Lingxiu Dong and Nils Rudi. Who benefits from transshipment? exogenous vs. endogenous wholesale prices. *Management Sci.*, 50(5):645–657, 2004.

Gary D Eppen. Note – Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.*, 25(5):498–501, 1979.

Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Computer System Sci.*, 69(3):485–497, 2004.

Mark Ferguson, V Daniel R Guide Jr, and Gilvan C Souza. Supply chain coordination for false failure returns. *Manufacturing Service Oper. Management*, 8(4):376–393, 2006.

Forrester. Customer Desires Vs. Retailer Capabilities: Minding The OmniChannel Commerce Gap. 2014. Accessed August 18, 2019, `http://global.sap.com/asia/campaigns/2014-07/hybris-accenture-forrester-tlp-omni-channel.pdf`.

FTI Consulting. The U.S. online retail forecast. 2015. Accessed August 18, 2019, `http://www.fticonsulting.com/~/media/Files/us-files/insights/featured-perspectives/fti-2015onlineretailforecast.pdf`.

Drew Fudenberg and J Miguel Villas-Boas. Behavior-based price discrimination and customer recognition. *Handbook Econom. Inform. Systems*, 1:377–436, 2006.

Guillermo Gallego and Ilkyeong Moon. The distribution free newsboy problem: review and extensions. *J. Oper. Res. Society*, 44(8):825–834, 1993.

Guillermo Gallego, Anran Li, Van-Anh Truong, and Xinshang Wang. Online personalized resource allocation with customer choice. 2016. Working paper, Hong Kong University of Science and Technology, Hong Kong.

Santiago Gallino and Antonio Moreno. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Sci.*, 60(6):1434–1451, 2014.

Santiago Gallino, Antonio Moreno, and Ioannis Stamatopoulos. Channel integration, sales dispersion, and inventory management. *Management Sci.*, 63(9):2813–2831, 2017.

Fei Gao and Xuanming Su. Online and offline information for omnichannel retailing. *Manufacturing Service Oper. Management*, 19(1):84–98, 2016.

Fei Gao and Xuanming Su. Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Sci.*, 63(8):2478–2492, 2017.

Nicole Giannopoulos. The macy's advantage: Secrets to same-day delivery, omnichannel. *RIS News*, 2014. Accessed August 18, 2019, `https://risnews.com/macys-advantage-secrets-same-day-delivery-omnichannel`.

Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-time optimization of personalized assortments. *Management Sci.*, 60(6):1532–1551, 2014.

Aravind Govindarajan, Amitabh Sinha, and Joline Uichanco. Joint inventory and fulfillment decisions for omnichannel retail networks. Working Paper, University of Michigan, 2018.

Grani A Hanasusanto, Daniel Kuhn, Stein W Wallace, and Steve Zymler. Distributionally robust multi-item newsvendor problems with multimodal demand distributions. *Math. Programming*, 152(1-2):1–32, 2015.

Mary Hanbury. Costco's generous return policy is ripe for abuse. *Business Insider*, August 17, 2018. `https://www.businessinsider.com/costco-membership-canceled-return-policy-2018-8`.

J Michael Harrison and Jan A van Mieghem. Multi-resource investment strategies: Operational hedging under demand uncertainty. *Eur. J. Oper. Res.*, 113(1):17–29, 1999.

James D Hess, Wujin Chu, and Eitan Gerstner. Controlling product returns in direct marketing. *Marketing Letters*, 7(4):307–317, 1996.

Daniel P Heyman and Matthew J Sobel. *Stochastic models in operations research: Stochastic optimization*, volume 2. Courier Corporation, 2003.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *J. American Statistical Association*, 58(301):13–30, 1963.

Lu Hsiao and Ying-Ju Chen. Returns policy and quality risk in e-business. *Production Oper. Management*, 21(3):489–503, 2012.

Lu Hsiao and Ying-Ju Chen. Return policy: Hassle-free or your money-back guarantee? *Naval Res. Logist.*, 61(5):403–417, 2014.

Lu Hsiao and Ying-Ju Chen. Retailer's rationale to refuse consumer returns in supply chains. *Naval Res. Logist.*, 62(8):686–701, 2015.

Woonghee Tim Huh, Ganesh Janakiraman, John A Muckstadt, and Paat Rusmevichientong. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Sci.*, 55(3):404–420, 2009.

Elnaz Jalilipour Alishah, Kamran Moinzadeh, and Yong-Pin Zhou. Inventory fulfillment strategies for an omni-channel retailer. Working Paper, University of Washington, 2015.

Narayan Janakiraman and Holly Syrdal. How to design a return policy. *Harvard Business Rev.*, September 9, 2016. `https://hbr.org/2016/08/how-to-design-a-return-policy`.

Stefanus Jasin and Amitabh Sinha. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Oper. Res.*, 63(6):1336–1351, 2015.

Stephen C Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3):241–254, 1967.

Michael E Ketzenberg and Rob A Zuidwijk. Optimal pricing, ordering, and return policies for consumer goods. *Production Oper. Management*, 18(3):344–360, 2009.

Eugene Kim. More than half of online sales growth in the US came from Amazon last year. *Business Insider*, February 2, 2017. `http://www.businessinsider.com/amazon-drives-more-than-half-us-ecommerce-growth-2016-2017-2`.

Marcel J Kleijn and Rommert Dekker. An overview of inventory systems with several demand classes. *Lecture Notes in Economics and Mathematical Systems*, pages 253–266, 1999.

Kurt Salmon. Building a solid omnichannel foundation with effective inventory management. 2016. Accessed December 12, 2016, `www.kurtsalmon.com/en-us/Retail/vertical-insight/1478/Building-a-Solid-Omnichannel-Foundation-with-Effective-Inventory-Management`.

Yanzhe Lei, Stefanus Jasin, and Amitabh Sinha. Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing Service Oper. Management*, 20 (2):269–284, 2018.

Jordy Leiser. Think tank: Why an omnichannel approach can help retail escape the Amazon. *WWD*, 2016. Accessed August 18, 2019, wwd.com/retail-news/forecasts-analysis/blazing-the-amazon-and-why-an-omnichannel-approach-can-help-retail,-10312172/.

Sam Lewis. Macy's grows order fulfillment centers to 500. *Retail Solutions Online*, 2013. Accessed August 18, 2019, https://www.retailitinsights.com/doc/macy-s-grows-order-fulfillment-centers-to-0001.

Matt Lindner. Target now ships online orders from more than 1,000 stores. *InternetRetailer*, 2016. Accessed August 18, 2019, https://www.internetretailer.com/2016/11/16/target-now-ships-online-orders-more-1000-stores.

Matt Lindner. E-commerce is expected to grow to 17% of US retail sales by 2022. *InternetRetailer*, August 2017. Retrieved from: https://www.digitalcommerce360.com/2017/08/09/e-commerce-grow-17-us-retail-sales-2022/.

Yunchuan Liu and Z John Zhang. Research note – The benefits of personalized pricing in a channel. *Marketing Sci.*, 25(1):97–105, 2006.

Guodong Lyu, Wang Chi Cheung, Mabel Chou, Chung-Piaw Teo, Zhichao Zheng, and Yuanguang Zhong. Capacity allocation in flexible production networks: Theory and applications. *Management Sci.*, 2017. Forthcoming.

Sridhar Moorthy and Kannan Srinivasan. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Sci.*, 14(4):442–466, 1995.

Julien Mostard, Rene De Koster, and Ruud Teunter. The distribution-free newsboy problem with resalable returns. *Internat. J. Production Econom.*, 97(3):329–342, 2005.

Leela Nageswaran, Soo-Haeng Cho, and Alan Andrew Scheller-Wolf. Consumer return policies in omnichannel operations. 2017. Working Paper, Carnegie Mellon University.

Sajjad Najafi and Izak Duenyas. Price markdowns to induce customers to opt out of free returns. 2018. Working Paper, University of Michigan.

Kim Nash. Walmart build supply chain to meet e-commerce demands. *The Wall Street Journal*, May 7, 2015. http://www.wsj.com/articles/wal-mart-builds-supply-chain-to-meet-e-commerce-demands-1431016708.

Karthik Natarajan and Chung-Piaw Teo. On reduced semidefinite programs for second order moment bounds with applications. *Math. Programming*, 161(1-2):487–518, 2017.

Karthik Natarajan, Chung Piaw Teo, and Zhichao Zheng. Mixed 0-1 linear programs under objective uncertainty: A completely positive representation. *Oper. Res.*, 59 (3):713–728, 2011.

Karthik Natarajan, Melvyn Sim, and Joline Uichanco. Asymmetry and ambiguity in newsvendor models. *Management Sci.*, 64(7):3146–3167, 2017.

Kawee Numpacharoen and Amporn Atsawarungruangkit. Generating correlation matrices based on the boundaries of their coefficients. *PLoS One*, 7(11):e48902, 2012.

Elie Ofek, Zsolt Katona, and Miklos Sarvary. "Bricks and clicks": The impact of product returns on the strategies of multichannel retailers. *Marketing Sci.*, 30(1): 42–60, 2011.

Manfred Padberg. The boolean quadric polytope: some characteristics, facets and relatives. *Math. Programming*, 45(1-3):139–172, 1989.

Colin Paterson, Gudrun Kiesmüller, Ruud Teunter, and Kevin Glazebrook. Inventory models with lateral transshipments: A review. *Eur. J. Oper. Res.*, 210(2):125–136, 2011.

J Andrew Petersen and V Kumar. Are product returns a necessary evil? Antecedents and consequences. *J. Marketing*, 73(3):35–51, 2009.

Erica E. Phillips. What stores do with $90 billion in merchandise returns. *The Wall Street Journal*, February 16, 2018. `https://www.wsj.com/articles/what-stores-do-with-90-billion-in-merchandise-returns-1518777000`.

Darrell Rigby. The future of shopping. *Harvard Business Rev.*, 89(12):65–76, 2011.

Lawrence W Robinson. Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Oper. Res.*, 38(2):278–295, 1990.

Donna Rosato. Guide to returning gifts: Retailers with the best and worst policies. *Consumer Reports*, December 10, 2018. `https://www.consumerreports.org/returns-refunds-exchanges/guide-to-returning-gifts/`.

Khadeeja Safdar. How your returns are used against you at best buy, other retailers. *The Wall Street Journal*, March 13, 2018. `https://www.wsj.com/articles/how-your-returns-are-used-against-you-at-best-buy-other-retailers-1520933400`.

Khadeeja Safdar and Laura Stevens. Banned from amazon: The shoppers who make too many returns. *The Wall Street Journal*, May 22, 2018. `https://www.wsj.com/articles/banned-from-amazon-the-shoppers-who-make-too-many-returns-1526981401`.

Khalid Saleh. E-commerce product return rate – statistics and trends. 2016. Accessed August 18, 2019, `https://www.invespcro.com/blog/ecommerce-product-return-rate-statistics/`.

Herbert Scarf. A min max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209, 1958.

Ralf W Seifert, Ulrich W Thonemann, and Marcel A Sieke. Relaxing channel separation: Integrating a virtual store into the supply chain via transshipments. *IIE Trans.*, 38(11):917–931, 2006.

Guangzhi Shang, Bikram P Ghosh, and Michael R Galbreth. Optimal retail return policies with wardrobing. *Production Oper. Management*, 26(7):1315–1332, 2017.

Hanif D Sherali and Amine Alameddine. A new reformulation-linearization technique for bilinear programming problems. *J. Global Optim.*, 2(4):379–410, 1992.

Jeffrey D Shulman, Anne T Coughlan, and R Canan Savaskan. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. *Manufacturing Service Oper. Management*, 11(4):577–594, 2009.

Jeffrey D Shulman, Anne T Coughlan, and R Canan Savaskan. Optimal reverse channel structure for consumer product returns. *Marketing Sci.*, 29(6):1071–1085, 2010.

James E Smith. Generalized chebychev inequalities: Theory and applications in decision analysis. *Oper. Res.*, 43(5):807–825, 1995.

Statista. Costs of return deliveries in the United States from 2016 to 2020 (in billion U.S. dollars). *Statista.com*, 2018. Accessed July 31, 2019, `https://www.statista.com/statistics/871365/reverse-logistics-cost-united-states/`.

Xuanming Su. Consumer returns policies and supply chain performance. *Manufacturing Service Oper. Management*, 11(4):595–612, 2009.

George Tagaras. Effects of pooling on the optimization and service levels of two-location inventory systems. *IIE Trans.*, 21(3):250–257, 1989.

George Tagaras and Morris A Cohen. Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Sci.*, 38(8):1067–1083, 1992.

M Ali Ülkü and Ülkü Gürler. The impact of abusing return policies: A newsvendor model with opportunistic consumers. *Internat. J. Production Econom.*, 203:124–133, 2018.

M Ali Ülkü, Lynn C Dailey, and H Müge Yayla-Küllü. Serving fraudulent consumers? The impact of return policies on retailer's profitability. *Service Sci.*, 5(4):296–309, 2013.

UPS. 2017 UPS rate and service guide. 2017. `https://www.ups.com/media/en/daily_rates.pdf`.

UPS Compass. Ship from store: A smart competitive strategy for retailers. 2014. Accessed August 18, 2019, `https://compass.ups.com/ship-from-store-benefits/`.

U.S. Census Bureau. Quarterly retail e-commerce sales. Accessed on December 12, 2016, `https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf`, 2016.

Jan A van Mieghem. Commissioned paper: Capacity management, investment, and hedging: Review and recent developments. *Manufacturing Service Oper. Management*, 5(4):269–302, 2003.

Jan A van Mieghem and Nils Rudi. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management*, 4(4):313–335, 2002.

Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM Rev.*, 38 (1):49–95, 1996.

Walker Sands. The future of retail 2018. 2018. Accessed August 18, 2019, `https://www.walkersands.com/resources/the-future-of-retail-2018/`.

Lai Wei, Stefanus Jasin, and Roman Kapuscinski. Shipping consolidation with delivery deadline and expedited shipment options. Working Paper, University of Michigan, 2017.

Wikipedia. List of United States cities by population. 2016. Accessed 12 December, 2016, `https://en.wikipedia.org/wiki/List_of_United_States_cities_by_population`.

Ping Josephine Xu, Russell Allgor, and Stephen C Graves. Benefits of reevaluating real-time order fulfillment decisions. *Manufacturing Service Oper. Management*, 11 (2):340–355, 2009.

Zhenzhen Yan, Sarah Yini Gao, and Chung Piaw Teo. On the design of sparse but efficient structures in operations. *Management Sci.*, 64(7):2973—-3468, 2018.

Hongsuk Yang and Linus Schrage. Conditions that cause risk pooling to increase inventory. *Eur. J. Oper. Res.*, 192(3):837–851, 2009.

Jian Yang and Zhaoqiong Qin. Capacitated production control with virtual lateral transshipments. *Oper. Res.*, 55(6):1104–1119, 2007.

David D Yao, Sean X Zhou, and Weifen Zhuang. Joint initial stocking and transshipment–Asymptotics and bounds. *Production Oper. Management*, 25(2):273–289, 2016.

Stefany Zaroban. U.S. e-commerce sales grow 16.0% in 2017. *InternetRetailer*, 2018. Accessed August 18, 2019, `https://www.digitalcommerce360.com/article/us-ecommerce-sales/`.