Joint Inventory and Fulfillment Decisions for Omnichannel Retail Networks

Aravind Govindarajan
Target Corporation, Sunnyvale, CA 94086, aravind.govindarajan@target.com

Amitabh Sinha
Amazon.com, Seattle, WA 98109, amitabsi@amazon.com

Joline Uichanco
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI 48109, jolineu@umich.edu

An omnichannel retailer with a network of physical stores and online 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 with available inventory. We consider inventory decisions at the start of the selling horizon for a seasonal product, with online fulfillment decisions made multiple times over the horizon. To address the intractability in considering inventory and fulfillment decisions together, we relax the problem using a hindsight-optimal bound, for which the inventory decision can be made independent of the optimal fulfillment decisions, while still incorporating virtual pooling of online demands across locations. We develop a computationally fast and scalable inventory heuristic for the multi-location problem based on the two-store analysis. The inventory heuristic directly informs dynamic fulfillment decisions that guide online demand fulfillment from stores. Using a numerical study based on a fictitious network embedded in the USA, we show that our heuristic significantly outperforms traditional strategies. The value of centralized inventory planning is highest when there is a moderate mix of online and in-store demands leading to synergies between pooling within and across locations, and this value increases with the size of the network. The inventory-aware fulfillment heuristic considerably outperforms myopic policies seen in practice, and is found to be near-optimal under a wide range of problem parameters.

Key words: omnichannel; e-commerce; inventory management; fulfillment; heuristic; asymptotic analysis

1. Introduction

In 2019, e-commerce sales accounted for around 11% of the total retail sales in the United States (U.S. Census Bureau 2019). Although this is a small portion of the total sales, online sales have been increasing at a rapid year-to-year growth rate around 18% in 2019 (Young 2019), and projected to account for 22% of all retail sales within the next five years (Chaffey 2019). With customers increasingly favoring the online channel, traditional brick-and-mortar (B&M) retail firms are equipping themselves with the ability to fulfill online orders from multiple inventory nodes (stores, fulfillment centers, etc.).

Omnichannel refers to this 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: 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). There is an industry-wide shift to omnichannel retailing, with one-time B&M firms like Macy's and Walmart integrating an online channel by leveraging their existing network of retail stores (Nash 2015); the e-commerce giant Amazon has also added a network of physical stores through the purchase of Whole Foods Market.

A key aspect of omnichannel operations is *store fulfillment*, which is the use of physical stores to fulfill online orders. An online customer who chooses ship-to-customer fulfillment may receive a package originating from, in theory, any inventory node in the network. Usually, the closest node is used since the shipping rate increases with the distance. However, when the closest node stocks out, 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 risk pooling across inventory nodes is reminiscent of transshipments between B&M stores; a distinction is that e-commerce fulfillment allows for this flexibility without prepositioning the inventory in a customer location.

Allowing demand spillovers essentially pools the geographically separate inventories. Since traditional inventory models do not account for demand spillovers, they overestimate the inventory required across the network. Indeed, this is reflected in the fact that, from 2010 to 2014, even as retail and online sales increased, inventory turnover decreased (Samuel 2017). Therefore, to reduce the burden of carrying too much inventory, omnichannel inventory planning must use network-based models that capture fulfillment flexibility. In this paper, we develop *network-based strategies* (that take into account online demand spillover) to *optimize inventory levels and fulfillment decisions* for an omnichannel firm. In particular, this firm has a network of physical stores and online fulfillment centers facing online (ship-to-customer) demands and in-store demands.

We consider the problem for a single, seasonal product with long lead times, such that inventory decisions are only made once for the start of the selling horizon. We assume that the selling horizon is divided into multiple periods, with the following dynamic in each period: in-store demands are fulfilled as they arrive; online fulfillment decisions (assigning online orders to fulfillment locations) are made at the end of the period with the available inventory; and unmet demands are lost. This assumption can be used to approximate the continuous-time problem, as the length of these periods can be arbitrarily small such that at most one unit of demand arrives in any period. This model leads to a multi-stage stochastic programming problem, and is intractable due to the size of the state space as well as complexities in the action space.

Our main contribution is a joint heuristic that co-ordinates network-based inventory and fulfillment decisions for omnichannel retailers. We derive the inventory heuristic based on a two-stage approximation that allows optimization of the inventory levels for the hindsight-optimal fulfillment policy. The fulfillment heuristic provides location-specific, time-varying inventory thresholds (calculated with the help of the inventory heuristic) which dictate the rationing of store inventory between in-store and online demands. We show by means of numerical studies on realistic inventory networks embedded in the USA, that by virtue of taking into account demand spillovers, our heuristic solution outperforms traditional decentralized strategies. The inventory-aware fulfillment heuristic is shown to be near-optimal under a variety of problem parameters, and provides significant cost savings compared to the myopic policy commonly seen in practice. Our solutions are highly scalable and easy to understand, which are of utmost importance in practice for retail networks.

We show that the value of such centralized planning (taking demand spillovers into account) is highest when there is a moderate mix of online and in-store demands that takes advantage of synergies in pooling across locations (demand spillovers) and within each location (in-store and online). Based on our analysis, for the current state of the industry (with around 10% of sales occurring online and growing quickly), we emphasize that retailers can significantly benefit by switching to centralized planning strategies that take into account online demand spillover.

We organize the paper as follows. We discuss relevant literature in Section 2, and introduce the general problem in Section 3. In Section 4, we develop the replenishment policy for the multi-location problem based on the two-stage hindsight-optimal approximation. In Section 5, we develop fulfillment thresholds informed by the inventory solution to guide dynamic fulfillment of online orders from stores. In Section 6, using realistic retail networks, we numerically test our heuristics against the benchmark solution and the hindsight optimal bound. Finally, we conclude with Section 7 discussing extensions and future research.

2. 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. There have been a number of papers that have focussed on customer behavior – 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. Nageswaran et al. (2020) consider optimal return policies for omnichannel firms when customers can return items in physical stores that were bought online.

Optimal fulfillment decisions for e-commerce demand 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); Acimovic and Graves (2014) focus on fulfillment strategies to minimize outbound shipping costs; 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; and Bayram and Cesaret (2020) consider the optimal dynamic fulfillment decisions for an omnichannel retail network when the initial inventory levels are given, by modeling shipping costs from each location as independent random variables.

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, and we consider the problem of deciding inventory levels across the network in the presence of e-commerce fulfillment flexibility. To the best of our knowledge, the study closest to ours is Jalilipour Alishah et al. (2017), who 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. They propose (without theoretical guarantees) using these rationing policies when there are multiple stores and a single online fulfillment center, with fulfillment costs not dependent on shipping distances. In contrast, we analyze the problem for general networks with multiple stores and fulfillment centers, with realistic fulfillment costs that depend shipping distances, which introduces complexity due to an additional rationing decision - online orders from other regions.

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). 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 regions (Tagaras 1989) or identical shipping costs across regions (Dong and Rudi 2004). Tagaras and Cohen (1992) show that when there is positive replenishment leadtime, the problem becomes intractable even for two regions due to the interdependence of optimal decisions 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. Their analysis is limited to two stores, as key mathematical properties like submodularity do not extend to multiple regions. Lim et al. (2020) consider a robust optimization approach to the joint allocation-fulfillment problem for e-commerce networks. They consider a two-step approach to optimization, wherein the periods in which replenishments

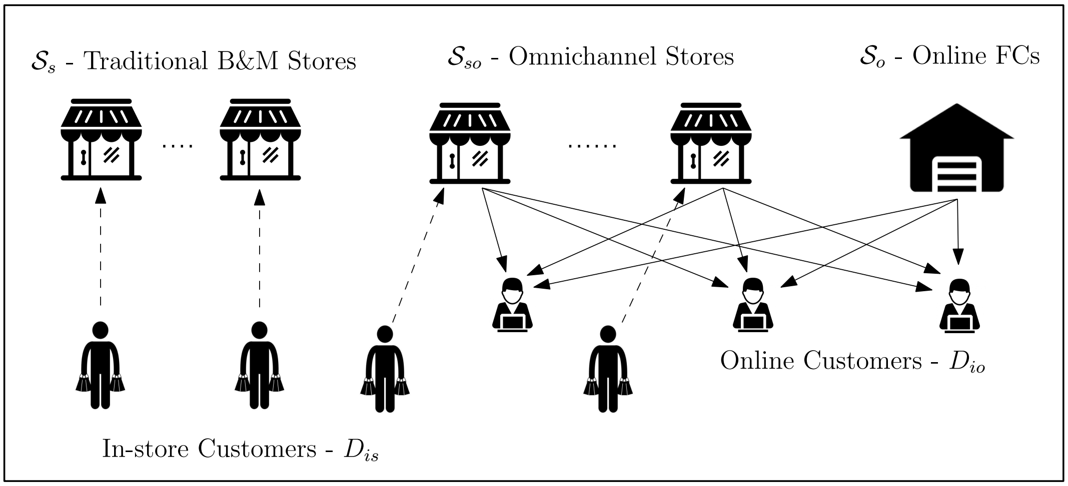

Figure 1 **Three types of facilities in an omnichannel fulfillment network a) Traditional brick-and-mortar stores ($\mathcal{S}_s$), b) Omnichannel stores ($\mathcal{S}_{so}$), and c) Online Fulfillment Centers ($\mathcal{S}_o$).**

arrive are optimized first, followed by the allocation and fulfillment decisions. To the best of our knowledge, our paper is the first to consider the multi-location omnichannel inventory and fulfillment problem, where we provide provably scalable and easy-to-implement policies that coordinate inventory and fulfillment strategies.

3. The General Problem - Model and Assumptions

Consider a firm selling a single product to multiple customer regions. We assume that a facility (inventory node) is located in each customer region. Customer regions without a facility in it can also be incorporated as zero-inventory facilities. We consider the problem of optimizing inventory and fulfillment decisions for this 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 customer region i, modeled by non-negative and continuous random variables:

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, and
2. the *online (ship-to-customer) demand* (D_{io}), consisting of customers ordering through the website or mobile app, expecting items to be delivered directly to their homes.

The demands are exogenous and are temporally independent, but can have any general channel or location correlation structure. A typical retail fulfillment network is shown in

Figure 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 regions with such stores from our analysis. We are only interested in regions where the facility is involved in online fulfillment, namely omnichannel stores and online fulfillment centers. Let $\mathcal{S} = \mathcal{S}_o \cup \mathcal{S}_{so}$ be the set of regions with omnichannel facilities. We denote $N = |\mathcal{S}|$ as its cardinality. We refer to the regions in $\mathcal{S}$ as $R_1, \ldots, R_N$.

There are two important features to be noted in the omnichannel problem. The first feature is that unfulfilled in-store demand at one region cannot be fulfilled by facilities in other regions. The second feature is that a facility in $\mathcal{S}$ with available inventory can fulfill online orders from any customer region. Hence, there is risk pooling of online demands across regions, as well as risk pooling of in-store and online demands within each region.

3.1. Sequence of Events

We consider a seasonal product for which the replenishment lead time is longer than the selling season, and hence there is only one chance at the start of the selling season to decide the inventory levels $y_1, \ldots, y_N$ across the network. We divide the selling season into T fulfillment periods. In each period $t \in [T]$, in-store demand (D_{is}^t) is fulfilled as it arrives, whereas fulfillment decisions to satisfy the online demands (D_{io}^t) are made at the end of the period with the available inventory in the network, and unmet demands are lost.

We note that the length of these periods can be arbitrarily small such that at most one unit of demand arrives in any period, and hence we can model the case where fulfillment decisions need to be made as soon as online orders arrive. In some cases, batch fulfillment of online orders also makes practical sense – 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.

3.2. Cost Parameters

We consider a per-unit fulfillment 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 $s_{ij} = s_{ji}$ and $s_{ij} \geq s_{ii}$ for any $i, j \in [N]$ since it is costlier to ship an item over longer distances. We will refer to s_{ii} (within the same region) as *in-location shipping* costs, and s_{ij} (across regions) as *cross-shipping* costs.

We assume that the in-location shipping costs are identical across regions: $s_{ii} = s$ for all $i \in [N]$. This assumption makes sense from a practical perspective, as the fixed component of the fulfillment cost (pick-pack-and-label) dominates for shipping over small distances, as this is usually done through human labor. At the end of a fulfillment period, each unit of unfulfilled in-store and online demands incurs 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 is costlier to lose, and cross-shipping always leads to a myopic reduction in cost: $s_{ij} < p_o$ for all $i, j \in [N]$. At the end of the selling horizon, each unit of unsold inventory incurs an overage cost h. We ignore the purchasing cost, but this can be added easily through linear terms. We summarize our assumptions on cost parameters in the set Ψ:

$$\Psi = \Big\{ p_s > p_o - s_{ij} > 0, \;\; \forall i, j \in [N]; \quad s_{ij} \geq s, \;\; \forall i, j \in [N] \text{ with } j \neq i \Big\}. \tag{1}$$

3.3. Stochastic Programming Formulation

We next develop the joint inventory planning and fulfillment problem of the seller as a stochastic program. In fulfillment period $t \in [T]$, let the starting inventory levels be denoted by $\mathbf{x^t} = (x_i^t)_i$, and $\tilde{\mathbf{D}}^\mathbf{t} = (D_{is}^t, D_{io}^t)_i$ denote the demands. From a facility in region 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 R_j. If region R_i has an OFC facility, then this region has no in-store demand, hence $z_i^t = D_{is}^t = 0$. We denote the period t fulfillment decisions in vector form as $\mathbf{z^t}, \mathbf{Z^t}$ respectively.

Suppose the seller starts the season with inventory levels $y_1, \ldots, y_N$. The optimal fulfillment decisions can be found by solving a $(T+1)$-stage stochastic program which can be solved through dynamic programming. Note that physical stores always prioritize in-store demand (since $p_s > p_o - s$) and online fulfillment is decided at the end of the period. Therefore, without loss of generality, we can assume that $\mathbf{z^t}, \mathbf{Z^t}$ is decided with knowledge of the demand realization $\tilde{\mathbf{D}}^\mathbf{t}$. Hence, $(\mathbf{x^t}, \tilde{\mathbf{D}}^\mathbf{t})$ is the state of the dynamic program. The optimal cost-to-go function in period t can be written as:

$$C_t(\mathbf{x^t}, \tilde{\mathbf{D}}^\mathbf{t}) = \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta^t} \left[P(\mathbf{x^t}, \tilde{\mathbf{D}}^\mathbf{t}, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E} C_{t+1} \left(\Big(x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t \Big)_i, \tilde{\mathbf{D}}^\mathbf{t+1} \right) \right] \tag{2}$$

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

$$P(\mathbf{x^t}, \tilde{\mathbf{D}}^\mathbf{t}, \mathbf{z^t}, \mathbf{Z^t}) = \sum_{i=1}^{N} p_s(D_{is}^t - z_i^t) + \sum_{j=1}^{N} p_o \left(D_{jo}^t - \sum_{i=1}^{N} Z_{ij}^t \right) + \sum_{i=1}^{N} sZ_{ii}^t + \sum_{i=1}^{N} \sum_{j=1, j\neq i}^{N} s_{ij} Z_{ij}^t \tag{3}$$

and Δ^t is the set of feasible fulfillment decisions in period t, described by the following set of constraints:

$$\begin{aligned} z_i^t + \sum_{j=1}^{N} Z_{ij}^t &\leq x_i^t, & \forall i \in [N] \\ z_i^t &\leq D_{is}^t, & \forall i \in [N] \\ \sum_{i=1}^{N} Z_{ij}^t &\leq D_{jo}^t, & \forall j \in [N] \\ \mathbf{z^t}, \mathbf{Z^t} &\geq 0. \end{aligned} \tag{4}$$

The first inequality in Δ^t represents the supply constraint (where $x_i^0 = y_i$ for all $i \in [N]$), 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 (4). At the end of the horizon, leftover inventory incurs a per-unit overage cost h, and hence, we have the boundary condition:

$$C_{T+1}(\mathbf{x^{T+1}}, \tilde{\mathbf{D}}^{T+1}) = h \cdot x^{T+1}, \quad \text{for any } \tilde{\mathbf{D}}^{T+1} \tag{5}$$

Hence, given the initial stocking level $\mathbf{y} = (y_i)_i$ at stage 0, the total expected cost over the T fulfillment periods is $C(\mathbf{y}) := \mathbb{E}C_1(\mathbf{y}, \tilde{\mathbf{D}}^\mathbf{1})$. The optimal initial stocking level minimizes this total expected cost. This initial inventory problem is shown to be a convex optimization problem in the following Proposition.

PROPOSITION 1. *$C(\mathbf{y}) := \mathbb{E}C_1(\mathbf{y}, \tilde{\mathbf{D}}^\mathbf{1})$ is jointly convex in the inventory levels $\mathbf{y} = (y_i)$.*

All proofs are relegated to the Appendix. We now discuss the intractability in the joint inventory and fulfillment problem. First, due to the assumption that $p_s > p_o - s$, in the optimal solution, we will have $z_i^t = \min(x_i^t, D_{is}^t)$, i.e. in-store demand is fulfilled first. However, the online fulfillment decisions are not straightforward – myopic policies that try to fulfill all online demand in the current period may be sub-optimal, as future in-store demands have a higher penalty cost. In addition, cross-shipping items to fulfill online demands at other regions may be sub-optimal as well, as it may be more profitable to reserve these items for cheaper future in-location store or online demands.

These two rationing elements render the multi-stage stochastic program intractable. In fact, a similar issue arises for optimal transshipment decisions with non-negligible lead times. Tagaras and Cohen (1992) show that for the two-store transshipment problem, even if the optimal policy is threshold-based, the optimization becomes intractable due to the complexity of the decision space in the dynamic programming formulation. The intractability arises solely due to the fact that in any period, it may be *ex-post* optimal to reserve inventory for future, less costly demand. Overlaying multiple regions, an additional class of demands, and inventory optimization, we cannot hope to solve this problem optimally.

4. Determining the Static Replenishment Policy

Our primary goal is to develop inventory solutions that take into account virtual pooling of online demand across regions. To this end, we relax the intractable joint inventory and fulfillment problem through a hindsight-optimal bound, and develop the inventory heuristic based on this bound. This inventory heuristic directly informs a simple, threshold-based fulfillment policy described later in Section 5, and we show that our joint inventory and fulfillment policy performs significantly better than traditional decentralized strategies.

4.1. The Two-Stage Approximation

First, we introduce a lower bound to $C(\mathbf{y})$. We do this by introducing the following hindsight-optimal T-period cost:

$$\begin{aligned}
\underline{C}(\mathbf{y}, \tilde{\mathbf{D}}) := \min_{\mathbf{x}^t, \mathbf{z}^t, \mathbf{Z}^t} \quad & \sum_{t=1}^{T} \left[\sum_{i=1}^{N} p_s (D_{is}^t - z_i^t) + \sum_{j=1}^{N} p_o \left(D_{jo}^t - \sum_{i=1}^{N} Z_{ij}^t \right) \right. \\
& \left. + \sum_{i=1}^{N} \left(s Z_{ii}^t + \sum_{j=1, j \neq i}^{N} s_{ij} Z_{ij}^t \right) \right] + \sum_{i=1}^{N} h \left(x_i^T - z_i^T - \sum_{j=1}^{N} Z_{ij}^T \right) \\
\text{s.t.} \quad & z_i^t + \sum_{j=1}^{N} Z_{ij}^t \leq x_i^t, \qquad \forall i \in [N], \forall t \in [T], \\
& z_i^t \leq D_{is}^t, \qquad \forall i \in [N], \forall t \in [T], \\
& \sum_{i=1}^{N} Z_{ij}^t \leq D_{jo}^t, \qquad \forall j \in [N], \forall t \in [T], \\
& \mathbf{z^t}, \mathbf{Z^t} \geq 0, \qquad \forall t \in [T], \\
& x_i^{t+1} = x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t, \qquad \forall i \in [N], \forall t \in [T], \\
& \mathbf{x^1} = \mathbf{y}.
\end{aligned} \tag{6}$$

Given the initial inventory $\mathbf{y}$ and after the realization of the demands $\tilde{\mathbf{D}} = (D^t_{is}, D^t_{io})_{i,t}$ in all regions and in all periods, the two-stage model (6) finds the T-period fulfillment decision $(\mathbf{z}^t, \mathbf{Z}^t)_t$ that minimizes the total cost. Since this sequence of fulfillment decisions are optimized after *all* demands are realized for the entire T periods, any non-anticipating fulfillment policy cannot achieve a total cost lower than $\underline{C}(\mathbf{y}, \tilde{\mathbf{D}})$ on a given sample path $\tilde{\mathbf{D}}$. Therefore, $\mathbb{E}(\underline{C}(\mathbf{y}, \tilde{\mathbf{D}}))$ is a lower bound on the expected total cost $C(\mathbf{y})$ of the optimal non-anticipating fulfillment policy. Hence, we refer to $\mathbb{E}(\underline{C}(\mathbf{y}, \tilde{\mathbf{D}}))$ as the *hindsight-optimal* bound for the multi-stage stochastic program.

The hindsight-optimal bound is useful since it simplifies the fulfillment decisions. The following Lemma provides a closed-form expression for the optimal fulfillment in (6).

LEMMA 1. *An optimal solution* $(\mathbf{z}^{t^*}, \mathbf{Z}^{t^*})$ *to (6) satisfies the following:*

1. $\sum_{t=1}^{T} z_i^{t^*} = \min\left(y_i\,,\, \sum_{t=1}^{T} D^t_{is}\right), \quad \forall i \in [N]$
2. $\sum_{t=1}^{T} Z_{ii}^{t^*} = \min\left(\left(y_i - \sum_{t=1}^{T} D^t_{is}\right)^+, \sum_{t=1}^{T} D^t_{io}\right), \quad \forall i \in [N]$
3. $\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1}^{N} Z_{ij}^{t^*} = \min\left(\sum_{i=1}^{N}\left(y_i - \sum_{t=1}^{T} D^t_{is}\right)^+, \sum_{i=1}^{N}\sum_{t=1}^{T} D^t_{io}\right)$

Here, $x^+ = \max(x, 0)$ for any $x \in \Re$. The implication of this lemma is that the hindsight-optimal solution can be determined sequentially. First, all in-store demands are fulfilled to the maximum extent. Then, with the leftover inventory, in-location online demands in each region are fulfilled. Finally, cross-fulfillment takes place to fulfill unmet online demands with leftover inventory in the network. This characterization of optimal fulfillment decisions is helpful, as we can rewrite the second-stage cost function in an interpretable manner.

Proposition 2. $\underline{C}(\mathbf{y},\tilde{\mathbf{D}})$ *is equivalent to the following linear program:*

$$\begin{aligned}
\underline{C}(\mathbf{y},\tilde{\mathbf{D}}) = \min_{\mathbf{z},\mathbf{Z}} \quad & \sum_{i=1}^{N} p_s(D_{is}-z_i) + \sum_{j=1}^{N} p_o\left(D_{jo} - \sum_{i=1}^{N} Z_{ij}\right) \\
& + \sum_{i=1}^{N}\left(sZ_{ii} + \sum_{j=1, j\neq i}^{N} s_{ij} Z_{ij}\right) + \sum_{i=1}^{N} h\left(y_i - z_i - \sum_{j=1}^{N} Z_{ij}\right) \\
\text{s.t.} \quad & z_i + \sum_{j=1}^{N} Z_{ij} \leq y_i, \quad \forall i \in [N], \\
& z_i \leq D_{is}, \quad \forall i \in [N], \\
& \sum_{i=1}^{N} Z_{ij} \leq D_{jo}, \quad \forall j \in [N], \\
& \mathbf{z}, \mathbf{Z} \geq 0
\end{aligned} \tag{7}$$

where $D_{is} \leftarrow \sum_{t=1}^{T} D_{is}^t$*, and* $D_{io} \leftarrow \sum_{t=1}^{T} D_{io}^t$.

We observe that the difference between (6) and (7) is that the former model determines the fulfillment quantities in each period, while the latter model determines the fulfillment quantities aggregated over the T periods.

The expected hindsight-optimal cost is defined to be

$$C^{IIP}(\mathbf{y}) := \mathbb{E}\left[\underline{C}(\mathbf{y},\tilde{\mathbf{D}})\right] \tag{8}$$

where, by Proposition 2, $\underline{C}(\mathbf{y},\tilde{\mathbf{D}})$ is equivalent to a single-fulfillment-period problem. Recall that $C^{IIP}(\mathbf{y}) \leq C(\mathbf{y})$. We can then interpret $C^{IIP}(\mathbf{y})$ as an approximation of the optimal cost of the stochastic program $C(\mathbf{y})$ which simplifies the fulfillment problem. Hence, by using cost function $C^{IIP}(\mathbf{y})$ to optimize the inventory levels, we are able to decouple the inventory problem from the intractable dynamic fulfillment problem, while still capturing the effect of virtual pooling of online demands across locations. We refer to the problem of optimizing $C^{IIP}(\mathbf{y})$ as the Integrated Inventory Planning (IIP) problem, since it jointly optimizes the inventory in the entire network. We denote C^{IIP} as the optimal IIP expected cost, i.e., $C^{IIP} := \min_{\mathbf{y}\geq 0} C^{IIP}(\mathbf{y})$.

Before we analyze the IIP problem, it is worthwhile discussing the simpler case where cross-fulfillment is not taken into account in inventory planning: this corresponds to the solution where $Z_{ij}=0$ for all $i,j\in[N]$ where $i\neq j$. When cross-fulfillment is not allowed, (7) yields the following optimal fulfillment quantities:

$$\begin{aligned} z_i &= \min\left(y_i, D_{is}\right) && \forall i\in[N] \\ Z_{ii} &= \min\left(\left(y_i-D_{is}\right)^+, D_{io}\right) && \forall i\in[N] \end{aligned} \tag{9}$$

When cross-fulfillment is ignored, the problem of determining the optimal initial inventory to minimize the total expected cost is referred to as the Decentralized Inventory Planning (DIP) problem, where the expected cost can be decomposed by region as follows:

$$\begin{aligned} \underline{C}^{DIP}(\mathbf{y}) = \sum_{i=1}^{N}\mathbb{E}\Big[& h\left(\left(y_i-D_{is}\right)^+-D_{io}\right)^+ + p_s\left(D_{is}-y_i\right)^+ \\ & + p_o\left(D_{io}-\left(y_i-D_{is}\right)^+\right)^+ + s\min\left(\left(y_i-D_{is}\right)^+, D_{io}\right)\Big]. \end{aligned} \tag{10}$$

Most contemporary omnichannel firms plan for inventory in this decentralized fashion for each store, a relic from traditional brick-and-mortar inventory models.

We can further simplify the decentralized expected cost. To do so, we use the identities $\min(x,y)=y-(y-x)^+$, and $(D_{is}-y_i)^+ + \left(D_{io}-(y_i-D_{is})^+\right)^+ = (D_i-y_i)^+$. Hence, (10) can be expressed in terms of the total demands $D_i=D_{is}+D_{io}$ as follows:

$$\underline{C}^{DIP}(\mathbf{y}) = \sum_{i=1}^{N}\left(s\mu_{io} + \mathbb{E}\Big[h\left(y_i-D_i\right)^+ + \left(p_o-s\right)\left(D_i-y_i\right)^+ + \left(p_s-p_o+s\right)\left(D_{is}-y_i\right)^+\Big]\right) \tag{11}$$

where $\mu_{io}=\mathbb{E}[D_{io}]$. It is evident from (11) that $\underline{C}^{DIP}(\mathbf{y})$ is a jointly convex function in $\mathbf{y}$. Therefore, the following first-order conditions are necessary and sufficient for the optimal DIP inventory solution:

$$\left(h+p_o-s\right)F_{D_i}\left(y_i\right) + \left(p_s-p_o+s\right)F_{D_{is}}\left(y_i\right) = p_s, \qquad \forall i\in[N] \tag{12}$$

where F_D denotes the cumulative distribution function of demand D. Note that the left-hand side in (12) is non-decreasing in y_i, while the right-hand side is constant. Therefore, a simple bisection method yields the unique optimum $\mathbf{y^{DIP}} = \left(y_i^{DIP}\right)_{i\in[N]}$.

Thus, we see that the inventory solution is simple when cross-fulfillment is ignored in optimizing the inventory levels. However, solving the IIP problem (8) where cross-fulfillment is allowed, is not straightforward. This is because the total fulfillment costs can only be characterized through a linear program. Even if there are only two locations, where fulfillment costs can be characterized in closed-form, certain complications arise. We first analyze the two-store setting to exhibit the complicated nature of the inventory problem alone. The insights derived in this case inform our analysis of a generalized multi-location case, which includes a network of omnichannel stores and online fulfillment centers.

4.2. The Two-Store Integrated Inventory Planning (IIP) Strategy

We consider the single period problem described in (7), with $N=2$, and the facilities in each region are omnichannel stores. Using the property in Lemma 1, the quantity cross-shipped to region R_j by the store in R_i is the minimum of the inventory remaining at R_i and the unfulfilled online demand at R_j, after both stores have fulfilled their own demands. Hence, the expected cost function is:

$$\begin{aligned} C^{IIP}(\mathbf{y}) := \mathbb{E}\Bigg[\sum_{i=1,2} \Bigg(& h\left((y_i - D_{is})^+ - D_{io}\right)^+ + p_s(D_{is} - y_i)^+ \\ & + p_o\left(D_{io} - (y_i - D_{is})^+\right)^+ + s\min\left((y_i - D_{is})^+, D_{io}\right)\Bigg) \\ & + (s_{12} - h - p_o)\min\left(\left((y_1 - D_{1s})^+ - D_{1o}\right)^+, \left(D_{2o} - (y_2 - D_{2s})^+\right)^+\right) \\ & + (s_{21} - h - p_o)\min\left(\left((y_2 - D_{2s})^+ - D_{2o}\right)^+, \left(D_{1o} - (y_1 - D_{1s})^+\right)^+\right)\Bigg] \end{aligned} \tag{13}$$

The last two terms in Equation 13 represent the value of cross-shipping: the total savings by cross-shipping a unit from R_i to R_j, $h + p_o - s_{ij}$, times the total quantity cross-shipped from R_i to R_j (and vice versa). The total cross-shipped quantity can be expressed as:

$$\sum_{i=1,2}\left(D_{io} - (y_i - D_{is})^+\right)^+ - \left(\sum_{i=1,2} D_{io} - \sum_{i=1,2}(y_i - D_{is})^+\right)^+ \tag{14}$$

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. Since $s_{12} = s_{21}$, we can simplify Equation 13 as follows:

$$\begin{aligned} C^{IIP}(\mathbf{y}) = s\sum_{i=1,2}\mu_{io} + \sum_{i=1,2}\mathbb{E}\Big[& h\,(y_i - D_i)^+ + (p_s - p_o + s)(D_{is} - y_i)^+ + (p_o - s)\,(D_i - y_i)^+\Big] \\ + (s_{12} - h - p_o)\Bigg[& \sum_{i=1,2}(D_i - y_i)^+ - \sum_{i=1,2}(D_{is} - y_i)^+ - \left(\sum_{i=1,2} D_{io} - \sum_{i=1,2}(y_i - D_{is})^+\right)^+\Bigg] \end{aligned} \tag{15}$$

Note that (15) is a special case of $C(\mathbf{y})$ where $N=2$ and $T=1$. Hence, from Proposition 1 we know that $C^{IIP}(\mathbf{y})$ is a convex function whose optimizer can be found efficiently using gradient descent methods. However, the nested term $\left(\sum_i D_{io} - \sum_i (y_i - D_{is})^+\right)^+$ complicates the calculation of the gradient. Note that the presence of this nested piece-wise linear term is due to the fact that in-store demand is prioritized. (This type of term does not arise for the case of traditional transshipments which can be seen by setting $D_{is}=0$ for all $i=1,2$.) By noting structural similarities of the IIP problem 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 (7).

Proposition 3 **(van Mieghem and Rudi 2002)**. *Under the conditions on cost parameters in Ψ, with $N=2$, there exists a partition $(\Omega_k(y_1,y_2))_{k=1}^{20}$ of the demand space such that in region k of the partition, the dual-price vector of the inventory constraints is equal to $\lambda^k = (\lambda_1^k, \lambda_2^k)$. Hence, the gradient of the IIP cost function can be written as*

$$\nabla C^{IIP}(\mathbf{y}) = (h,h)^\intercal - \sum_{k=1}^{20} \lambda^k \cdot \mathbb{P}\left(\tilde{\mathbf{D}} \in \Omega_k(y_1,y_2)\right). \tag{16}$$

The four-dimensional demand vector is separable into 20 independent regions based on the values of y_1 and y_2, within which the dual price vector of the inventory constraints is constant (refer to Appendix B for a detailed discussion). This enables formulating the gradient as in Equation 16. The optimal solution $\mathbf{y^{IIP}} = (y_1^{IIP}, y_2^{IIP})$ can thus be obtained by gradient descent, where in each iterative step, the probability of realization of every demand region has to be recalculated.

However, this gradient-based approach does not extend to more than two stores, as the number of regions in the partition increases exponentially, and the regions cannot be enumerated tractably. This is due to the fact that cross-shipment quantities are now set by a linear program, as compared to explicit expressions in the two-store case. Hence we develop a tractable lower bound based on a relaxation motivated by practice, yielding a heuristic solution for the two-store case, which we then extend to multiple regions.

4.2.1. 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 the nested piecewise linear terms in the cost function. We relax this by treating unfulfilled in-store demand as online demand which can be fulfilled by cross-shipping. This is commonly seen in practice, where if an in-store customer is unable to find an item on the shelf, store personnel are equipped with the ability to place an online order for the item to be delivered directly to the customer's home.

Mathematically, we make the following replacement:

$$\underbrace{\sum_{i=1,2}(D_{is}-y_i)^+ + \left(\sum_{i=1,2} D_{io} - \sum_{i=1,2}(y_i - D_{is})^+\right)^+}_{\text{unfulfilled demand with in-store demand not pooled}} \overset{(\geq)}{\longrightarrow} \underbrace{\left(\sum_{i=1,2} D_i - \sum_{i=1,2} y_i\right)^+}_{\text{unfulfilled demand when all demands are pooled}}. \tag{17}$$

Here, recall that $D_i = D_{is} + D_{io}$ is the aggregate demand in region i. Using the relaxation (17), Proposition 4 formally establishes a lower bound to the IIP expected cost:

PROPOSITION 4. *For any* $\mathbf{y} \geq 0$*, we have* $C^{LB}(\mathbf{y}) \leq C^{IIP}(\mathbf{y})$*, where*

$$\begin{aligned} C^{LB}(\mathbf{y}) := s(\mu_{1o}+\mu_{2o}) + \mathbb{E}\Big[& h\,(y_1+y_2-D)^+ + (p_o - s_{12})\,(D-y_1-y_2)^+ \\ & + (s_{12}-s)\,(D_1-y_1)^+ + (s_{12}-s)\,(D_2-y_2)^+ \\ & + (p_s-(p_o-s))\,(D_{1s}-y_1)^+ + (p_s-(p_o-s))\,(D_{2s}-y_2)^+ \Big] \end{aligned} \tag{18}$$

where $D = D_1 + D_2$ *is the total demand.*

Note that since the right-hand side of (18) does not have nested piecewise linear terms, the gradient of $C^{LB}(\mathbf{y})$ has a simple expression. Hence, the first-order conditions satisfied by the optimizer of C^{LB} can be written as:

$$(h+p_o-s_{12})F_D\left(\sum_{j=1,2} y_j\right) + (s_{12}-s)F_{D_i}(y_i) + (p_s-p_o+s)F_{D_{is}}(y_i) = p_s, \quad \forall i = 1,2. \tag{19}$$

Since C^{LB} is a convex function, these first-order conditions are necessary and sufficient. Equation 19 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 means that the inventory levels in each region must be

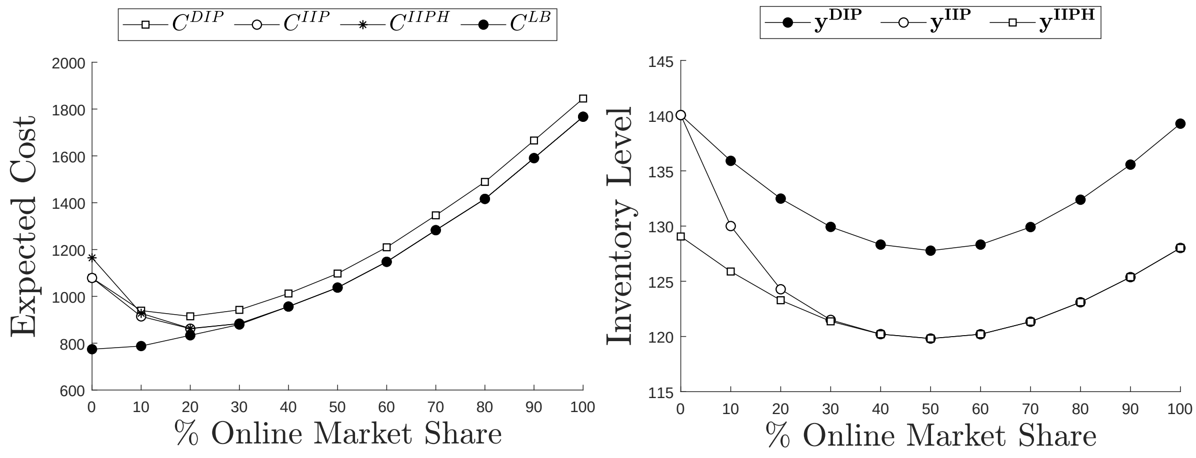

(a) (b)

Figure 2 **Shows the effect of online market share on expected costs (left) and the optimal store inventory levels (right).**

different, in contrast to Dong and Rudi (2004) where the optimality equation only yields a system-wide inventory level.

Note that the first-order condition (19) is a system of two equations with two variables. This system can be solved numerically to yield the inventory solution $\mathbf{y^{IIPH}}$. Since the first-order conditions are necessary and sufficient for optimality, $\mathbf{y^{IIPH}}$ minimizes the lower bound function $C^{LB}(\cdot)$. We use it as a heuristic solution to the IIP problem (15). We denote the expected cost of the heuristic under the IIP problem as $C^{IIPH} := C^{IIP}(\mathbf{y^{IIPH}})$.

The relaxation (17) to formulate the lower bound will be tight when the in-store demand is small compared to the online demand since the optimal inventory levels will be based on the total demands. We test this numerically by changing the mix of in-store and online demands in Figure 2. 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: $h = 10$, $p_s = 100$, $p_o = 100$, $s = 5$, $s_{12} = 7.5$.

In Figure 2a, we compare the expected costs of the heuristic, C^{IIPH}, the optimal IIP expected cost C^{IIP}, the expected cost of the decentralized inventory levels, $C^{DIP} := C^{IIP}(\mathbf{y^{DIP}})$, and the lower bound: $C^{LB} := C^{LB}(\mathbf{y^{IIPH}})$. We make a few observations. First, the heuristic provides savings over the decentralized inventory solution for most cases, except when the online market share is low ($< 10\%$). However, we note that when the online market share is low, the potential savings from centralized planning is limited, as seen from comparing C^{IIP} and C^{DIP}. Thus, in cases of very low online market share, the firm can simply plan for each region separately using the decentralized inventory strategy.

Second, centralized inventory planning is most valuable when there is a moderate mix of online and in-store demands (the expected cost C^{IIP} is minimum when the online market share is $\sim 20\%$). 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 regions work in synergy.

Third, as the in-store demand becomes smaller, the probability that there will be unfulfilled in-store demand decreases, so the relaxation (17) becomes tighter – we see that the lower bound is tight when the online market share is more than 30%. Correspondingly for this range of online market share, in Figure 2b, we see that the solution for the relaxed problem (IIPH) converges to the optimal IIP solution. We infer that when a significant portion of the demand occurs online ($> 30\%$), in-store demand can effectively be treated as online demand that can be fulfilled from any location (which makes the model tractable), as the probability of unfulfilled in-store demand becomes negligible.

The savings in cost in Figure 2a arises from a change in inventory levels in anticipation of pooling across customer regions. Proposition 5 addresses this observation from Figure 2b that the IIPH solution consistently stocks less than the DIP solution at each store.

Proposition 5. *For identical stores and normal demands,* $\mathbf{y}^{\mathbf{IIPH}} \leq (\geq)\, \mathbf{y}^{\mathbf{DIP}}$ *whenever* $\mathbf{y}^{\mathbf{DIP}} \geq (\leq)\, \mathbf{m}$*, where* $\mathbf{m}$ *is the vector of mean total demands at stores. Under perfect positive correlation across customer regions,* $\mathbf{y}^{\mathbf{IIPH}} = \mathbf{y}^{\mathbf{DIP}} = \mathbf{y}^{\mathbf{IIP}}$*.*

Similar to the intuition in newsvendor settings, $\mathbf{y}^{\mathbf{DIP}} \geq \mathbf{m}$ occurs 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 (12), which has a mixture distribution as compared to a simple normal distribution in newsvendor theory. Lastly, positive correlation across regions reduces the pooling benefits achieved by cross-shipping, and under perfect correlation, there is no benefit from pooling as all regions either have too much or too little inventory without any imbalance.

4.3. The Multi-Store Integrated Inventory Planning (IIP) Strategy

We now consider the integrated inventory planning problem IIP (8) for a general system with multiple customer regions (i.e., $N \geq 2$). The cross-shipping costs are taken to be $s_{ij} = s + f(d_{ij})$, where d_{ij} is the distance between region R_i and region R_j, and f is a non-negative, increasing function such that $f(d) \to 0$ as $d \to 0$. For the conditions (1) on the cost parameters to hold true, we assume $\sup_{d \in \mathcal{D}} f(d) < p_o - s$ where $\mathcal{D} := \{d_{ij} : i, j \in [N]\}$.

4.3.1. Hierarchical Shipping Cost Structures. The second-stage cost function in the IIP problem has a closed-form expression when $N = 2$ since the total shipping cost can be formulated in closed-form (see Equation 13). The closed-form equation was key in the development of the lower bound and the heuristic for the two-store case.

However, when $N > 2$, the second-stage cost function is the optimal value of a linear program, as seen in (7), and may not have a closed-form equation. We next show that if the shipping costs satisfy a hierarchical property (called *nested* cost structures in Govindarajan et al. 2020), then we can still express the expected cost in closed-form. Restricting our focus on these cost structures allow us to develop a heuristic for the multi-store case.

We say that the IIP problem has *nested shipping costs* if these costs can be represented by a tree structure with L levels ($2 \leq L \leq N$). The bottom-most level consists of N leaves, with each leaf representing a customer region. Each level of the tree corresponds to different shipping cost values; the cost is increasing when traversing the tree from bottom to top. The lowest level corresponds to the in-location shipping cost, s. The highest level corresponds to the most expensive shipping cost value. When determining the shipping cost between regions R_i and R_j, this is simply the cost value corresponding to the lowest level where the two regions are connected.

Nested shipping costs are useful since they imply a hierarchy to cross-shipping. Since the bottom-most level corresponds to the least expensive shipping cost, demand is fulfilled to the maximum extent using in-location fulfillment. If there are any unfulfilled demand or remaining inventory after this level of fulfillment, we can use the next level up for fulfillment, and so on. This is done until we reach the top-most level with the highest shipping cost. Hence, when the shipping costs are nested, the second-stage cost function (7) can be formulated as a sum of piecewise linear functions in $\mathbf{y}$ and $\tilde{\mathbf{D}}$.

To formally state this result, we introduce some notation. (This notation is consistent with Govindarajan et al. (2020) where nested structures are discussed in detail.) The set of regions $[N]$ is partitioned in to N_ℓ sets for each level $\ell = 0, 1, \ldots, L-1$, such that $N = N_0 > N_1 > \ldots > N_{L-1} = 1$. If two regions belong to the same set in level ℓ, then they are connected at that level of the tree. We denote the N_ℓ sets in partition ℓ as $\{\mathcal{I}_1^{(\ell)}, \mathcal{I}_2^{(\ell)}, \ldots, \mathcal{I}_{N_\ell}^{(\ell)}\}$. Define $\Xi = \{\mathbf{E}_0, \mathbf{E}_1, \ldots, \mathbf{E}_{L-1}\}$ as the set of assignment matrices, where the level ℓ assignment matrix $\mathbf{E}_\ell$ is a binary matrix of size $N_\ell \times N$ where the (k,i) entry is equal to 1 if and only if region R_i is 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. The tree structure follows from the assumption that any set in level ℓ is the union of sets in the preceding level $\ell - 1$.

If two regions are in set $\mathcal{I}_k^{(\ell)}$, then the shipping cost between the two regions is $s_{\ell,k}$. To induce the nested hierarchy for fulfillment, we assume that it is less costly to fulfill demand in lower levels. Mathematically, if $k^{(\ell)}(i)$ is the level ℓ set index of region R_i, since $s_{0,k^{(0)}(i)} = s$ for all $i \in [N]$, we assume that $s \le s_{1,k^{(1)}(i)} \le \cdots \le s_{L-1}$. We denote by $\mathbf{S} = \{s_{\ell,k}\}$ the set of all shipping costs. Note that the nested hierarchy, Ξ, and the shipping costs, $\mathbf{S}$, fully characterize the nested cost structure.

Proposition 6 **(Govindarajan et al. (2020))**. *Under the L-level nested structure,*

$$C^{IIP}(\mathbf{y}) = \mathbb{E}\Bigg[s \cdot \mathbf{e}^\top \mathbf{D}_o + h \cdot (\mathbf{e}^\top \mathbf{y} - \mathbf{e}^\top \mathbf{D})^+ + p_s \cdot \mathbf{e}^\top (\mathbf{D}_s - \mathbf{y})^+ + \sum_{\ell=0}^{L-1} \eta_\ell^\top \left(\mathbf{E}_\ell \mathbf{D}_o - \mathbf{E}_\ell (\mathbf{y} - \mathbf{D}_s)^+ \right)^+ \Bigg] \tag{20}$$

where $\mathbf{D}_o = (D_{io})_{i=1}^N$, $\mathbf{D}_s = (D_{is})_{i=1}^N$, $\mathbf{D} = (D_{io}, D_{is})_{i=1}^N$, $\eta_{L-1} = p_o - s_{L-1}$ *and, for* $\ell \le L-2$, $\eta_\ell = (\eta_{\ell,k})_{k \in [n_\ell]}$ *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* $\mathcal{I}_k^{(\ell)}$.

The reason that such a reformulation is possible under a nested structure is due to the fact that the total shipping cost can be expressed in closed-form by summing the shipping costs in each level. In level 0, the shipping cost is

$$\sum_{i \in [N]} s \cdot \min\left(D_{io}, (y_i - D_{is})^+\right) = s \cdot \mathbf{e}^\top \mathbf{D_o} - \sum_{i \in [N]} s \cdot \left(D_{io} - (y_i - D_{is})^+\right).$$

Because of the hierarchy in fulfillment induced by the costs, we know that for any level $\ell \ge 1$, the number of fulfilled units of demand from regions in set $\mathcal{I}_k^{(\ell)}$ at level ℓ is

$$\underbrace{\sum_{m \in \mathcal{K}_k^{(\ell)}} \left(\sum_{i \in \mathcal{I}_m^{(\ell-1)}} D_{io} - \sum_{i \in \mathcal{I}_m^{(\ell-1)}} (y_i - D_{is})^+ \right)^+}_{\text{unmet demand in } \mathcal{I}_k^{\ell} \text{ after level } \ell-1} - \underbrace{\left(\sum_{i \in \mathcal{I}_k^{(\ell)}} D_{io} - \sum_{i \in \mathcal{I}_k^{(\ell)}} (y_i - D_{is})^+ \right)^+}_{\text{unmet demand in } \mathcal{I}_k^{(\ell)} \text{ after level } \ell}, \tag{21}$$

where $\mathcal{K}_k^{(\ell)}$ is the set of level $\ell-1$ children of set $\mathcal{I}_k^{(\ell)}$. Note that the per-unit cost of this fulfillment is $s_{\ell,k}$. The same idea was used to quantify the number of units fulfilled through cross-shipping in the two-store case in Equation 14. This can also be reconciled with the fact that the two-store case naturally has a 2-level nested structure.

Thus, the total shipping cost is calculated as:

$$s\cdot\mathbf{e}^\top\mathbf{D}_o+\sum_{\ell=0}^{L-2}\sum_{k\in[N_\ell]}(s_{\ell+1,m^{(\ell+1)}(k)}-s_{\ell,k})\cdot\left(\sum_{i\in\mathcal{I}_k^{(\ell)}}D_{io}-\sum_{i\in\mathcal{I}_k^{(\ell)}}(y_i-D_{is})^+\right)^+$$

While nested cost structures are a good approximation for geographic distances in countries like the US, in general, shipping costs need not exhibit a nested structure. When shipping costs are not nested, they can be approximated with nested shipping costs using a hierarchical agglomerative clustering algorithm. Govindarajan et al. (2020) showed that such an algorithm resulted in only a small gap in approximation of the expected shipping costs. Hence, even though our focus in this section is on nested shipping costs, the heuristics that we develop are applicable to general shipping costs.

4.3.2. Lower Bound and Heuristic for the Multi-Location Problem. The key difference of our setting from the pure e-commerce setting in Govindarajan et al. (2020) is that the available inventory at region i for online demand fulfillment is $(y_i-D_{is})^+$ instead of y_i. This gives rise to nested piecewise linear terms in (20) that complicate the calculation of the first-order conditions for optimality. Similar to the two-store case, we make the following relaxation for each level ℓ to obtain a lower bound:

$$\underbrace{\mathbf{E}_\ell(\mathbf{D}_s-\mathbf{y})^++\left(\mathbf{E}_\ell\mathbf{D}_o-\mathbf{E}_\ell(\mathbf{y}-\mathbf{D}_s)^+\right)^+}_{\text{unfulfilled demand at level }\ell\text{ with in-store demand not pooled}}\quad\overset{(\geq)}{\longrightarrow}\quad\underbrace{(\mathbf{E}_\ell\mathbf{D}-\mathbf{E}_\ell\mathbf{y})^+}_{\text{unfulfilled demand when all demands are pooled}}$$

This relaxation implies that in-store demand is also pooled among the regions in set $I_k^{(\ell)}$ in a nested fashion. If a region R_i stocks out, any unmet in-store demand is routed to other locations with available inventory in the set $I_{k^{(\ell)}(i)}^{(\ell)}$, for increasing values of ℓ starting with $\ell=0$, and the demand is lost if no inventory exists after level $L-1$. Using this approximation, a lower bound to $C^{IIP}(\mathbf{y})$ is given by:

$$C_L^{LB}(\mathbf{y})=\mathbb{E}\Bigg[s\cdot\mathbf{e}^\top\mathbf{D}_o+h\cdot(\mathbf{e}^\top\mathbf{y}-\mathbf{e}^\top\mathbf{D})+(p_s-p_o+s)\cdot\mathbf{e}^\top(\mathbf{D}_s-\mathbf{y})^+ \\ +\sum_{\ell=0}^{L-1}\hat{\eta}_\ell^\top(\mathbf{E}_\ell\mathbf{D}-\mathbf{E}_\ell\mathbf{y})^+\Bigg] \qquad (22)$$

where $\hat{\eta}_\ell = \eta_\ell$ for all $\ell \leq L-2$, and $\hat{\eta}_{L-1} = h + p_o - s_{L-1}$. The subscript L on C_L^{LB} refers to the L levels to the nested shipping costs.

Having eliminated the nested piecewise linear terms, we can obtain the first order conditions as follows:

$$h + (p_s - p_o + s) \cdot (F_{D_{is}}(y_i) - 1) + \sum_{\ell=0}^{L-1} (\hat{\eta}_\ell^\top)_{k^{(\ell)}(i)} \cdot \left(F_{(\mathbf{E}_\ell \mathbf{D})_{k^{(\ell)}(i)}} \left((\mathbf{E}_\ell \mathbf{y})_{k^{(\ell)}(i)} \right) - 1 \right) = 0, \quad \forall i \in \mathcal{S}_{so}$$

$$h + \sum_{\ell=0}^{L-1} (\hat{\eta}_\ell^\top)_{k^{(\ell)}(i)} \cdot \left(F_{(\mathbf{E}_\ell \mathbf{D})_{k^{(\ell)}(i)}} \left((\mathbf{E}_\ell \mathbf{y})_{k^{(\ell)}(i)} \right) - 1 \right) = 0, \qquad \forall i \in \mathcal{S}_o \tag{23}$$

where $(\mathbf{x})_j$ denotes the j^{th} element of vector $\mathbf{x}$, and $k^{(\ell)}(i)$ is the level ℓ set index of region i. 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. Solving these system of N equations can be challenging, especially for large values of N seen in practice.

Suppose that the shipping costs $\mathbf{S} = (s_{ij})$ are not nested. By using hierarchical agglomerative clustering (described in Govindarajan et al. 2020), we can approximate these shipping cost values with the nested shipping costs $\mathbf{S}' = (s'_{ij})$. This leads to an approximation of the IIP cost $C^{IIP}(\mathbf{y})$ with function $C^{IIP'}(\mathbf{y})$ which we define to be the right-hand side of (20) with the approximate cost values $\mathbf{S}'$. The corresponding lower bound (22) is a lower bound on $C^{IIP'}(\mathbf{y})$, but not necessarily to IIP cost $C^{IIP}(\mathbf{y})$ of the original problem. This is because hierarchical clustering is not guaranteed to provide a lower bound to $C^{IIP}(\mathbf{y})$.

Alternatively, we can obtain a lower bound to $C^{IIP}(\mathbf{y})$ using a different nested structure approximation: we set $L = 2$, with $s_0 = s_1 = s$, i.e. the shipping costs are constant and equal to the in-location fulfillment cost, which gives us the following expected cost:

$$\begin{aligned} C_2^{LB}(\mathbf{y}) = \mathbb{E}\Big[& s \cdot \mathbf{e}^\top \mathbf{D}_o + h \cdot (\mathbf{e}^\top \mathbf{y} - \mathbf{e}^\top \mathbf{D})^+ + (p_o - s)\left(\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y}\right)^+ \\ & + (p_s - p_o + s) \sum_{i \in \mathcal{S}_{so}} (D_{is} - y_i)^+ \Big] \end{aligned} \tag{24}$$

This function is a lower bound on the IIP cost, as formalized in the following Proposition.

PROPOSITION 7. *When $\mathcal{S}_o = \emptyset$, the following are true for the cost function C_2^{LB}:*

1. $C_2^{LB}(\mathbf{y}) \leq C^{IIP}(\mathbf{y})$ *for any* $\mathbf{y} \geq 0$.
2. *Let* $\mathbf{y^{IIPH}} := \text{argmin}_{\mathbf{y}\geq 0} C_2^{LB}(\mathbf{y})$. $\mathbf{y^{IIPH}}$ *is unique, and is given by the solution to the system of* N *equations:*

$$(h + p_o - s) \cdot F_D\left(\mathbf{e}^\top \mathbf{y}\right) + (p_s - p_o + s) \cdot F_{D_{is}}(y_i) = p_s, \qquad \forall i \in \mathcal{S}, \tag{25}$$

where F_D *is the cumulative distribution of the aggregate demand* $D = \sum_{j\in\mathcal{S}}(D_{jo} + D_{js})$.
3. *When demands follow a multivariate normal distribution, the optimal solution* $\mathbf{y^{IIPH}}$ *has the following property: for some* $\nu \in [0,1]$,

$$y_i^{IIPH} = F_{D_{is}}^{-1}(\nu), \qquad \forall i \in \mathcal{S}. \tag{26}$$

4. *If demands are bounded and i.i.d. across regions, and if* $h > 0$ *is sufficiently small, then as the number of regions increases,* $\mathbf{y^{IIPH}}$ *is near-optimal in an asymptotic sense with a constant approximation factor. That is, if* $\mathbf{y^{IIP}} := \text{argmin}_{\mathbf{y}\geq 0} C^{IIP}(\mathbf{y})$, *then:*

$$1 \leq \frac{C^{IIP}(\mathbf{y^{IIPH}})}{C^{IIP}(\mathbf{y^{IIP}})} \leq 1 + \frac{h + p_o - s}{p_s - p_o + s}, \qquad \text{as } N \to \infty.$$

Proposition 7 illustrates the utility of the constant fulfillment cost approximation. First, it provides a valid lower bound to the multi-location two-stage cost function $C^{IIP}(\mathbf{y})$. Second, the finding the optimal solution to the lower bound is equivalent to solving a simple system of equations. Third, in the case where demands follow a normal distribution (a common assumption in practice), finding the optimal solution is further simplified by the reduction to a single variable optimization problem, namely over the common critical fractile (ν) of the in-store demands. Finally, the solution obtained by the approximation has an asymptotically bounded performance when the network size is large, thereby ensuring that the heuristic is not arbitrarily bad compared to the optimal solution.

When $\mathcal{S}_o \neq \emptyset$, Proposition 7(1) still holds; however, the first order conditions (25) fail to yield a feasible solution. We describe an algorithm for networks with $\mathcal{S}_o \neq \emptyset$ in Appendix C – the algorithm preserves the scalability property (Proposition 7(3)), and requires calculation of inventories at OFCs separately, while using them as input to calculate the store inventory levels. The scalability of the constant fulfillment cost heuristic ($\mathbf{y^{IIPH}}$) makes it a favorable candidate for inventory planning for networks of large sizes – we show in the numerical analysis (Section 6), the benefit of such centralized planning increases with network size, and hence a scalable solution is of utmost importance for real life networks with thousands of inventory nodes.

5. Determining the Reactive Fulfillment Policy

Recall the dynamics of the joint inventory and fulfillment setting described in Section 3.1. Specifically, after the initial inventory levels are chosen, the seller then has to decide how to dynamically fulfill the online demands that arrive over T fulfillment periods. The previous section focused on developing an inventory heuristic $\mathbf{y}^{\mathbf{IIPH}}$ that can be found efficiently by solving a system of equations. Hence, in this section, we next focus on developing a dynamic fulfillment policy for a given initial inventory level.

The simplest policy one can think of is myopic fulfillment – online demands are fulfilled as they arrive, from the closest location (in terms of shipping cost) with available inventory. Indeed, many firms follow this policy due to its simplicity. Note that this policy ignores future demands while making decisions on the fly. However, it may be *ex-post* optimal to withhold inventory at a location to fulfill future demands, than to fulfill an online order in the current period from a far-away location.

The natural question is to then ask, how much inventory should be withheld at a location in any period? We consider two fulfillment policies, which guide online order fulfillment:

1. the *myopic fulfillment (MF)* policy, where online demands in the current fulfillment period 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 demands, by halting online fulfillment from a location when the inventory level falls below a (time-dependent) threshold.

Since 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, but only if the thresholds are chosen correctly. Incorrectly setting aside too much inventory (by setting a high threshold) affects demand fulfillment, leading to increased lost online sales. 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. (2017) prove the existence of an optimal threshold rationing policy between in-store and online demands at a single store.

In the multi-location problem, it is not straightforward to estimate the underage cost for the low-priority (online) demand, as it is endogenized by the fulfillment policy and depends on where an order is fulfilled from. The optimal thresholds depend on in-store and online

Algorithm 1 Implementation of the Threshold Fulfillment (TF) Policy

1: In each fulfillment period t, each store first fulfills its own in-store demand to the maximum possible extent, and the leftover inventory at store is $\hat{x_i}^t, \forall i \in \mathcal{S}_{so}$.

2: The inventory available for online fulfillment at each store is $K_i^t = (\hat{x_i}^t - w_i^t)^+, \forall i \in \mathcal{S}_{so}$, where the thresholds w_i^t are calculated from (27).

3: Given the online demands D_{jo}, $j \in \mathcal{S}$, the online fulfillment decisions Z_{ij}^t, $i, j \in \mathcal{S}$ are obtained from solving the transportation LP:

$$\min \left\{ \sum_{i,j \in \mathcal{S}} (s_{ij} - p_o) Z_{ij}^t \ : \ \sum_{k \in \mathcal{S}} Z_{kj}^t \leq D_{jo}, \ \sum_{k \in \mathcal{S}} Z_{ik}^t \leq K_i^t, \ Z_{ij}^t \geq 0, \ \forall i, j \in \mathcal{S} \right\}$$

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. Alternatively, we leverage the fact that the IIP inventory model developed in the previous section captures these network-based trade-offs. We utilize the inventory heuristic ($\mathbf{y^{IIPH}}$) to inform fulfillment decisions in the following way: at store i, after in-store demands are fulfilled at the end of period t, use the excess inventory (if any) above thresholds w_i^t to fulfill online demands, where w_i^t is calculated as:

$$w_i^t = \begin{cases} \max\left(F^{-1}_{D_{is}^{[t+1,T]}} \left(\frac{p_s}{h+p_s} \right), \ \left(\mathbf{y^{IIPH,t+1}} \right)_i \right), & \text{if } i \in \mathcal{S}_{so}, \\ 0, & \text{if } i \in \mathcal{S}_o, \end{cases} \tag{27}$$

where $D_{is}^{[t+1,T]} := \sum_{t'=t+1}^{T} D_{is}^{t'}$, and $\mathbf{y^{IIPH,t+1}}$ is the inventory heuristic applied to the time horizon $[t+1, T]$. Thus, this policy sets aside inventory at each store for future demands as specified by the IIP inventory heuristic. In Figure 2, we noted that when in-store demand is dominant, the heuristic yields inventory levels that are lower than optimal. To correct for this phenomenon, we take the maximum of the heuristic inventory level, and a newsvendor quantity that caters to future in-store demands alone.

We formalize the TF policy in Algorithm 1. The fulfillment thresholds can be calculated at the start of the selling season, and only need be re-evaluated if the demand forecasts for the remaining periods in the horizon are updated. The calculation of fulfillment thresholds is computationally light due to the scalability of the inventory heuristic (Proposition 7(3)). The MF policy places no threshold restrictions on online fulfillment, and can simply be recovered from Algorithm 1 by setting the thresholds w_i^t to be zero for all $i \in \mathcal{S}$ in step 1.

We can evaluate the performance of these fulfillment policies for any given inventory decision by comparing them with the clairvoyant hindsight-optimal policy that can be calculated from (6). In Section 6, we show that the TF policy achieves a much smaller gap with respect to the hindsight-optimal lower bound, compared to the MF policy.

6. 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. Even though we developed our joint inventory and fulfillment heuristic policy based on relaxations of the problem, we will evaluate all policies based on the total expected cost of the T-fulfillment period problem, $C(\mathbf{y})$, through a Monte Carlo simulation with sample size of 10^4. By varying different problem parameters, we primarily compare our joint heuristic $\langle$IIPH,TF$\rangle$ with the traditional solution $\langle$DIP,MF$\rangle$ as a benchmark, to demonstrate the value of centralized planning. We also report the gap achieved by the fulfillment policies (MF,TF) from the lower bound following the hindsight-optimal policy (HF), for starting inventory $\mathbf{y}^{\text{IIPH}}$.

We considered alternative benchmarks and bounds which are not reported here. A deterministic solution that stocks the mean total demand at each location (which then informs fulfillment thresholds) was found to be an inferior benchmark compared to $\langle$DIP,MF$\rangle$ in most cases. A lower bound for the expected cost of the joint heuristic can be obtained by jointly optimizing (6) for inventory and fulfillment decisions; in most cases, this yielded an optimality gap of less than 20%. Due to looseness of this clairvoyant bound, in few cases (such as low online market share), the upper bound on the optimality gap was even as high as 70%, rendering this comparison non-informative without access to optimal costs.

6.1. Network Setup

Stores are taken to be located at the most populous cities in mainland US (Wikipedia 2016) and the OFCs are located according to the list of most efficient locations for warehouses in terms of possible transit lead-times (Chicago Consulting 2016). 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. Other cost parameters used are: $h = 10$, $p_s = p_o = 100$, $s = 9.182$. The demands are taken to be independent and normally distributed with parameters proportional to the population of the cities, with α being the

No. of OFCs	No. of Stores	⟨IIPH,TF⟩ vs ⟨DIP,MF⟩			% Gap vs ⟨IIPH,HF⟩	
		% Savings	% Imbalance Reduction	% Turnover Increase	⟨IIPH,MF⟩	⟨IIPH,TF⟩
1	50	12.0	5.9	20.8	13.7	1.2
1	100	16.5	13.1	25.3	19.2	1.7
1	150	19.0	20.2	27.8	22.2	1.9
2	50	14.4	16.5	23.1	13.9	1.2
2	100	17.5	23.2	26.5	18.9	1.7
2	150	19.7	29.4	28.8	21.7	1.8
5	50	16.6	32.2	26.4	12.9	1.1
5	100	18.8	35.2	28.3	16.9	1.5
5	150	21.1	37.7	30.3	20.6	1.8
10	50	17.4	35.3	27.3	12.5	1.0
10	100	19.5	36.7	28.8	17.1	1.5
10	150	21.4	38.0	30.6	20.1	1.7

Table 1 Effect of network size (number of stores and OFCs)

proportion of the total demand that occurs online. The coefficient of variation of the total selling-season demands at each location are fixed at 0.2. The total demand is split evenly across the T fulfillment periods into identical and independent normal random variables. We denote the number of physical stores by n_s, and the number of OFCs by n_o. In the base case, we take $\alpha = 0.5, T = 5, n_s = 50, n_0 = 2$. Further details on the numerical setup and a brief overview of the simulation process can be found in Appendix D.

6.2. Effect of Network Size

Table 1 shows that increasing the network size has a positive and marginally decreasing effect on the cost savings of ⟨IIPH,TF⟩ relative to the traditional solution of decentralized inventory planning and myopic fulfillment ⟨DIP,MF⟩. As the network size increases, centralized inventory planning and strategic fulfillment is increasingly valuable, as there is more pooling and flexibility in terms of options available in fulfillment.

We also compare the strategies based on two important metrics: inventory imbalance and inventory efficiency. Higher imbalance can lead to costly spillovers and local stockouts (Acimovic and 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 fulfillment period, and taking the average value over the selling horizon. 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

s_{max}/s	% Savings	% Gap vs ⟨IIPH,HF⟩	
	⟨IIPH,TF⟩ vs ⟨DIP,MF⟩	⟨IIPH,MF⟩	⟨IIPH,TF⟩
1.16	13.6	13.4	1.2
2	13.1	13.0	1.2
3	12.7	12.6	1.2
4	12.4	12.2	1.3
5	11.6	12.5	1.4

Table 2 **Effect of the sensitivity of cross-shipping costs to distance ($s_{\max}/s$)**

level of the system in the selling horizon (calculated as the mean of the starting inventory level and expected ending inventory at the end of the horizon). Higher efficiency achieved by the heuristic ⟨IIPH,TF⟩ stems from a reduction in the starting inventory levels without a considerable decrease in service levels, due to planning in advance for cross-shipping. This offers a potential solution to the decreasing trend in turnovers in the retail industry in recent years (Samuel 2017). The last two columns of Table 1 show that the threshold-based fulfillment (TF) policy significantly outperforms myopic fulfillment (MF) when compared relative to the hindsight-optimal fulfillment (HF). While increasing number of stores increases the gap of both policies to the hindsight-optimal lower bound, the threshold-based policy achieves a significantly smaller gap compared to the myopic policy. Note that the achieved gap with respect to the hindsight-policy is an upper bond on the optimality gap (with respect to the optimal non-anticipating fulfillment policy).

6.3. Effect of Cross-shipping Costs

We next vary the slope of shipping costs with respect to distance, thereby increasing the ratio $s_{\max}/s$ (value of 1.16 corresponds to the base case setting), where $s_{\max} = \max_{i,j} s_{ij}$. As expected, the relative performance of the heuristic decreases as shipping costs become more sensitive to distance (Table 2). In practice, the range of shipping costs is not too large: for a 5lb package, the ratio $s_{\max}/s$ 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. Hence the heuristic provides significant savings for most existing shipping cost structures, while still being near-optimal ($< 1.5\%$ optimality gap).

6.4. Effect of Online Market Share

We observe similar trends (as previously seen in Figure 2a) when the online market share is varied (Table 3). As expected, we see that when the online market share is low, the benefit achieved by the heuristic is limited. However, the savings increases sharply when

Online Market Share (α)	% Savings ⟨IIPH,TF⟩ vs ⟨DIP,MF⟩	% Gap vs ⟨IIPH,HF⟩	
		⟨IIPH,MF⟩	⟨IIPH,TF⟩
10%	1.7	51.6	6.1
20%	17.9	65.5	11.0
30%	18.3	44.1	7.2
40%	16.1	24.2	2.9
50%	13.9	13.0	1.2
60%	12.2	7.8	0.5
70%	11.3	4.9	0.3
80%	11.0	2.9	0.2
90%	10.9	1.5	0.1

Table 3 Effect of online market share (α)

No. of Periods	% Savings ⟨IIPH,TF⟩ vs ⟨DIP,MF⟩	% Gap vs ⟨IIPH,HF⟩	
		⟨IIPH,MF⟩	⟨IIPH,TF⟩
1	14.2	0.0	0.0
5	14.1	13.8	1.3
10	14.6	23.3	2.6
15	14.4	28.5	4.3
20	14.3	32.2	5.8
25	14.1	34.3	7.3
30	14.2	36.5	9.1
35	13.9	38.9	10.7
40	13.6	42.0	12.2

Table 4 Effect of number of fulfillment periods (T)

the online market share increases, thus demonstrating that firms can obtain considerable savings through centralized inventory strategies in the current state of the industry.

Surprisingly, the lower bound gap of the fulfillment policies are also non-monotone. In fact, the myopic policy is quite close to the lower bound when the online market share is high, indicating that a simple myopic policy can be effective for pure-play e-commerce firms. However, the TF policy is far superior, especially for omnichannel firms with a moderate mix of online and in-store demands.

6.5. Effect of Number of Fulfillment Periods (T)

By increasing the number of times online fulfillment decisions are made, we can closely model the continuous-time case where fulfillment decisions are made as and when online orders arrive. We keep the parameters of the total demand over the selling season constant, and keep demands across fulfillment periods independent and identically distributed.

Table 4 tabulates the results. We see that the savings achieved by the heuristic is fairly stable and robust to the chosen value of T, whereas the lower bound gap achieved by the fulfillment policies are increasing in T. The MF policy is punished for failing to reserve inventory for future in-store demands, and the lower bound gap increases at a higher rate compared to the TF policy's gap, since the TF policy takes future demands into account.

It is worth noting here that the hindsight-optimal solution is only a proxy for the optimal non-anticipating fulfillment policy – as T increases, the variability of demands increase (since the variability of the total demand is kept constant), and hence we can expect that the hindsight-optimal bound may get looser with respect to the non-anticipating optimal fulfillment policy. Thus, while the upper bound to the optimality gap increases with T, it is not certain that the actual optimality gap does too.

7. 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 — network inventory management, by demonstrating the value in utilizing the pooling benefits offered by omnichannel retailing, through a combined inventory and fulfillment policy. Our heuristic solutions are highly scalable and easy to understand, and provide significant savings over strategies that are traditionally used in practice. Our solutions are generalizable to demands originating from arbitrary customer regions, by treating them as facilities that carry zero inventory.

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. Multi-period models for regular products may be considered, however, the complexity arising from fulfillment decisions may render these models intractable. A heuristic control for managing multiple products is also an interesting and important extension. Our models can also inform important network design decisions in determining optimal locations for new facilities. In conclusion, we believe that the scalability, interpretability and generalizability of our solutions make them capable of serving as helpful decision tools for practitioners.

Acknowledgments

This work was started while the authors Dr. Aravind Govindarajan and Dr. Amitabh Sinha were at the Ross School of Business at University of Michigan, and was completed independently of their employment at Target and at Amazon, respectively. This material is based upon work supported by the National Science Foundation under Grant Number 1561791. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Acimovic, Jason, Stephen C Graves. 2014. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Oper. Management* **17**(1) 34–51.

Acimovic, Jason, Stephen C Graves. 2017. Mitigating spillover in online retailing via replenishment. *Manufacturing & Service Oper. Management* **19**(3) 419–436.

Ansari, Asim, Carl F Mela, Scott A Neslin. 2008. Customer channel migration. *J. Marketing Res.* **45**(1) 60–76.

Bayram, Armagan, Bahriye Cesaret. 2020. Order fulfillment policies for ship-from-store implementation in omni-channel retailing. *European Journal of Operational Research* (forthcoming).

Bell, David R, Santiago Gallino, Antonio Moreno. 2017. Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Sci.* **64**(4) 1629–1651.

Brynjolfsson, Erik, Yu Jeffrey Hu, Mohammad S Rahman. 2013. Competing in the age of omnichannel retailing. *MIT Sloan Management Review* **54**(4) 23.

Chaffey, Dave. 2019. Forecast growth in percentage of online retail/ecommerce sales 2017 to 2023. *smartinsights.com.* Retrieved on 25-Dec-2019 from: `https://www.smartinsights.com/digital-marketing-strategy/online-retail-sales-growth/`.

Chen, Jian, Youhua Chen, Mahmut Parlar, Yongbo Xiao. 2011. Optimal inventory and admission policies for drop-shipping retailers serving in-store and online customers. *IIE Transactions* **43**(5) 332–347.

Chicago Consulting. 2016. Ten best warehouse networks. URL `http://www.chicago-consulting.com/ten-best-warehouse-networks/`. [Online; accessed 12-Dec-2016].

Chopra, Sunil, Peter Meindl. 2007. Supply chain management. strategy, planning and operation. *Das Summa Summarum des Management*. Springer, 265–275.

DeValve, Levi, Yehua Wei, Di Wu, Rong Yuan. 2018. Understanding the value of fulfillment flexibility in an online retailing environment. Working Paper.

Dong, Lingxiu, Nils Rudi. 2004. Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Sci.* **50**(5) 645–657.

Gallino, Santiago, Antonio Moreno. 2014. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Sci.* **60**(6) 1434–1451.

Gallino, Santiago, Antonio Moreno, Ioannis Stamatopoulos. 2017. Channel integration, sales dispersion, and inventory management. *Management Sci.* **63**(9) 2813–2831.

Gao, Fei, Xuanming Su. 2016. Online and offline information for omnichannel retailing. *Manufacturing & Service Oper. Management* **19**(1) 84–98.

Gao, Fei, Xuanming Su. 2017. Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Sci.* **63**(8) 2478–2492.

Govindarajan, Aravind, Amitabh Sinha, Joline Uichanco. 2020. Distribution-free inventory risk pooling in a multi-location newsvendor. *Management Sci.* (forthcoming).

Harrison, J Michael, Jan A van Mieghem. 1999. Multi-resource investment strategies: Operational hedging under demand uncertainty. *Eur. J. Oper. Res.* **113**(1) 17–29.

Hoeffding, Wassily. 1963. Probability inequalities for sums of bounded random variables. *J. American Statistical Association* **58**(301) 13–30.

Jalilipour Alishah, Elnaz, Kamran Moinzadeh, Yong-Pin Zhou. 2017. Store fulfillment strategy for an omnichannel retailer. Working Paper.

Jasin, Stefanus, Amitabh Sinha. 2015. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Oper. Res.* **63**(6) 1336–1351.

Kleijn, Marcel J, Rommert Dekker. 1999. An overview of inventory systems with several demand classes. *Lecture Notes in Economics and Mathematical Systems* 253–266.

Lei, Yanzhe, Stefanus Jasin, Amitabh Sinha. 2018. Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Oper. Management* **20**(2) 269–284.

Lim, Yun Fong, Song Jiu, Marcus Ang. 2020. Integrating anticipative replenishment-allocation with reactive fulfillment for online retailing using robust optimization. *Manufacturing & Service Oper. Management* (forthcoming).

Nageswaran, Leela, Soo-Haeng Cho, Alan Andrew Scheller-Wolf. 2020. Consumer return policies in omnichannel operations. *Management Sci.* (forthcoming).

Nash, Kim. 2015. Walmart build supply chain to meet e-commerce demands. *The Wall Street Journal.* Retrieved from: `http://www.wsj.com/articles/wal-mart-builds-supply-chain-to-meet-e-commerce-demands-1431016708`.

Paterson, Colin, Gudrun Kiesmüller, Ruud Teunter, Kevin Glazebrook. 2011. Inventory models with lateral transshipments: A review. *Eur. J. Oper. Res.* **210**(2) 125–136.

Rigby, Darrell. 2011. The future of shopping. *Harvard Business Review* **89**(12) 65–76.

Samuel, John. 2017. Catch the bus. Then steer it. *risnews.com* Retrieved on 2-Dec-2020 from: `https://risnews.com/catch-bus-then-steer-it`.

Seifert, Ralf W, Ulrich W Thonemann, Marcel A Sieke. 2006. Relaxing channel separation: Integrating a virtual store into the supply chain via transshipments. *IIE Trans.* **38**(11) 917–931.

Tagaras, George. 1989. Effects of pooling on the optimization and service levels of two-location inventory systems. *IIE Trans.* **21**(3) 250–257.

Tagaras, George, Morris A Cohen. 1992. Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Sci.* **38**(8) 1067–1083.

UPS. 2017. 2017 UPS rate and service guide. URL `https://www.ups.com/media/en/daily_rates.pdf`. [Online; accessed 10-Oct-2017].

U.S. Census Bureau. 2019. Quarterly retail e-commerce sales. Retrieved on 25-Dec-2019 from: `https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf`.

van Mieghem, Jan A, Nils Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management* **4**(4) 313–335.

Wikipedia. 2016. List of United States cities by population. URL `https://en.wikipedia.org/wiki/List_of_United_States_cities_by_population`. [Online; accessed 12-Dec-2016].

Yang, Jian, Zhaoqiong Qin. 2007. Capacitated production control with virtual lateral transshipments. *Oper. Res.* **55**(6) 1104–1119.

Yao, David D, Sean X Zhou, Weifen Zhuang. 2016. Joint initial stocking and transshipment–Asymptotics and bounds. *Production and Oper. Management* **25**(2) 273–289.

Young, Jessica. 2019. Global ecommerce sales to reach nearly $3.46 trillion in 2019. *digitalcommerce360.com.* Retrieved from: `https://www.digitalcommerce360.com/article/global-ecommerce-sales/`.

Appendices

Appendix A: Proofs of Propositions

A.1. Proof of Proposition 1

Consider the case where items are ordered at the start of the selling horizon, and online demands are fulfilled over T fulfillment periods. Assume that $C_{T+1}(\mathbf{x^{T+1}}, \tilde{\mathbf{D}}^{T+1}) = 0$ without loss of generality. Thus, from (2),(3),(4), $C_T(\mathbf{x^T}, \tilde{\mathbf{D}}^T)$ is the optimal value of a linear program which is jointly convex in $(\mathbf{x^T}, \tilde{\mathbf{D}}^T)$. This leads to the base case result that $C_T(\mathbf{x^T}, \tilde{\mathbf{D}}^T)$ is convex in x^T given any $\tilde{\mathbf{D}}^T$. By backward induction, we need to show that $C_t(\mathbf{x^t}, \tilde{\mathbf{D}}^t)$ is convex in $\mathbf{x^t}$ for any given $\tilde{\mathbf{D}}^t$, with the assumption that $C_{t+1}(\mathbf{x^{t+1}}, \tilde{\mathbf{D}}^{t+1})$ is convex in $\mathbf{x^{t+1}}$ given any $\tilde{\mathbf{D}}^{t+1}$. The cost-to-go function can be represented by $C_t(\mathbf{x^t}, \tilde{\mathbf{D}}^t) = \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \mathcal{G}(\mathbf{x^t}, \tilde{\mathbf{D}}^t, \mathbf{z^t}, \mathbf{Z^t})$, where

$$\mathcal{G}(\mathbf{x^t}, \tilde{\mathbf{D}}^t, \mathbf{z^t}, \mathbf{Z^t}) = \left[P(\mathbf{x^t}, \tilde{\mathbf{D}}^t, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E}C_{t+1}(x_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t, \tilde{\mathbf{D}}^{t+1})\right] \tag{28}$$

Consider any $\mu \geq 0$, and $\mathbf{x_1^t}, \mathbf{x_2^t} \geq 0$. Let $(\mathbf{z_i^t}, \mathbf{Z_i^t}) = \underset{\mathbf{z^t}, \mathbf{Z^t} \in \Delta}{\arg\min}\, \mathcal{G}(\mathbf{x_i^t}, \tilde{\mathbf{D}}^t, \mathbf{z^t}, \mathbf{Z^t})$. Note that P is a linear function in its variables (Equation 3), and $\mathbb{E}C_{t+1}(\mathbf{x^{t+1}}, \tilde{\mathbf{D}}^{t+1})$ 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:

$$\begin{aligned}
C_t(\bar{\mathbf{x}}^\mathbf{t}, \tilde{\mathbf{D}}^t) &= \min_{\mathbf{z^t}, \mathbf{Z^t} \in \Delta} \left[P(\bar{\mathbf{x}}^\mathbf{t}, \tilde{\mathbf{D}}^t, \mathbf{z^t}, \mathbf{Z^t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - z_i^t - \sum_{j=1}^{N} Z_{ij}^t, \tilde{\mathbf{D}}^{t+1})\right] \\
&\leq P(\bar{\mathbf{x}}^\mathbf{t}, \tilde{\mathbf{D}}^t, \bar{\mathbf{z}}^\mathbf{t}, \bar{\mathbf{Z}}^\mathbf{t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{\mathbf{D}}^{t+1}) \\
&\leq \mu P(\mathbf{x_1^t}, \tilde{\mathbf{D}}^t, \mathbf{z_1^t}, \mathbf{Z_1^t}) + (1-\mu)P(\mathbf{x_2^t}, \tilde{\mathbf{D}}^t, \mathbf{z_2^t}, \mathbf{Z_2^t}) + \mathbb{E}C_{t+1}(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{\mathbf{D}}^{t+1})
\end{aligned} \tag{29}$$

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

$$\begin{aligned}
\mathbb{E}C_{t+1}\big(\bar{x}_i^t - \bar{z}_i^t - \sum_{j=1}^{N} \bar{Z}_{ij}^t, \tilde{\mathbf{D}}^{t+1}\big) &= \mathbb{E}C_{t+1}\left[\mu\left(x_{1,i}^t - z_{1,i}^t - \sum_{j=1}^{N} Z_{1,ij}^t\right) + (1-\mu)\left(x_{2,i}^t - z_{2,i}^t - \sum_{j=1}^{N} Z_{2,i}^t\right), \tilde{\mathbf{D}}^{t+1}\right] \\
&\leq \mu\mathbb{E}C_{t+1}\left[x_{1,i}^t - z_{1,i}^t - \sum_{j=1}^{N} Z_{1,ij}^t, \tilde{\mathbf{D}}^{t+1}\right] + (1-\mu)\mathbb{E}C_{t+1}\left[x_{2,i}^t - z_{2,i}^t - \sum_{j=1}^{N} Z_{2,i}^t, \tilde{\mathbf{D}}^{t+1}\right]
\end{aligned} \tag{30}$$

Thus, from Equation 28, we have:

$$\begin{aligned}
C_t(\bar{\mathbf{x}}^\mathbf{t}, \tilde{\mathbf{D}}^t) &\leq \mu\mathcal{G}(\mathbf{x_1^t}, \tilde{\mathbf{D}}^t, \mathbf{z_1^t}, \mathbf{Z_1^t}) + (1-\mu)\mathcal{G}(\mathbf{x_2^t}, \tilde{\mathbf{D}}^t, \mathbf{z_2^t}, \mathbf{Z_2^t}) \\
&= \mu C_t(\mathbf{x_1^t}, \tilde{\mathbf{D}}^t) + (1-\mu)C_t(\mathbf{x_2^t}, \tilde{\mathbf{D}}^t)
\end{aligned} \tag{31}$$

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.2. Proof of Lemma 1

By recursion on x_i^t, we have: $x_i^T - z_i^T - \sum_j Z_{ij}^T = y_i - \sum_{t=1}^{T} z_i^t - \sum_{t=1}^{T}\sum_j Z_{ij}^t$. Thus, we have the following coefficients for the decision variables in the objective:

$$\begin{aligned}
&z_i^t : -p_s - h \quad , && \forall i, \forall t \leq T \\
&Z_{ii}^t : s - p_o - h \quad , && \forall i, \forall t \leq T \\
&Z_{ij}^t : s_{ij} - p_o - h && \forall i, j \neq i, \forall t \leq T
\end{aligned}$$

Note that based on the assumptions in Equation 1, we have: $-p_s - h > s - p_o - h \geq s_{ij} - p_o - h$. Then, by greedy allocation for each i, we will have $\sum_{t=1}^{T} z_i^t = \min(y_1, \sum_{t=1}^{T} D_{is}^t$, followed by $\sum_{t=1}^{T} Z_{ii}^t = \min\left(\left(y_i - \sum_{t=1}^{T} D_{is}^t\right)^+, \sum_{t=1}^{T} D_{io}^t\right)$. Finally, $\sum_{t=1}^{T}\sum_{i,j} Z_{ij}^t = \min\left(\sum_{i=1}^{N}\left(y_i - \sum_{t=1}^{T} D_{is}^t\right)^+, \sum_{i=1}^{N}\sum_{t=1}^{T} D_{io}^t\right)$. $\square$

A.3. Proof of Proposition 2

First we eliminate x_i^t variables using $x_i^t = y_i - \sum_{t'=1}^{t-1} z_i^{t'} - \sum_{t'=1}^{t-1} Z_{ij}^{t'}$. Thus, (6) is equivalent to:

$$\begin{aligned}
\underline{C}(\mathbf{y}, \tilde{\mathbf{D}}) = \min_{\mathbf{z}^t, \mathbf{Z}^t} \quad & \sum_{t=1}^{T} \left[\sum_{i=1}^{N} p_s (D_{is}^t - z_i^t) + \sum_{j=1}^{N} p_o \left(D_{jo}^t - \sum_{i=1}^{N} Z_{ij}^t \right) \right. \\
& \left. + \sum_{i=1}^{N} s Z_{ii}^t + \sum_{i=1}^{N} \sum_{j=1, j\neq i}^{N} s_{ij} Z_{ij}^t \right] + \sum_{i=1}^{N} h \left(y_i - \sum_{t=1}^{T} z_i^t - \sum_{t=1}^{T}\sum_{j=1}^{N} Z_{ij}^t \right) \\
\text{s.t.} \quad & \sum_{t'=1}^{t} z_i^{t'} + \sum_{t'=1}^{t}\sum_{j=1}^{N} Z_{ij}^{t'} \leq y_i, \quad \forall i \in [N], \forall t \in [T], \\
& z_i^t \leq D_{is}^t, \quad \forall i \in [N], \forall t \in [T], \\
& \sum_{i=1}^{N} Z_{ij}^t \leq D_{jo}^t, \quad \forall j \in [N], \forall t \in [T], \\
& \mathbf{z^t}, \mathbf{Z^t} \geq 0, \quad \forall t \in [T]
\end{aligned} \tag{32}$$

First, note that the first constraint can be replaced by $\sum_{t'=1}^{T} z_i^{t'} + \sum_{t'=1}^{T}\sum_{j=1}^{N} Z_{ij}^{t'} \leq y_i, \quad \forall i \in [N]$, since $\mathbf{z^t}, \mathbf{Z^t} \geq 0$. Since the objective in (6) contains the decision variables z_i^t, Z_{ij}^t only occurring in the sum over T (i.e. as $\sum_{t=1}^{T} z_i^t$ and $\sum_{t=1}^{T} Z_{ij}^t$), we can replace the second and third constraints by $\sum_{t=1}^{T} z_i^t \leq \sum_{t=1}^{T} D_{is}^t, \quad \forall i \in [N]$ and $\sum_{t=1}^{T}\sum_{i=1}^{N} Z_{ij}^t \leq \sum_{t=1}^{T} D_{jo}^t, \quad \forall j \in [N]$ respectively. Note that this replacement relaxes the problem, but we show that the objective solution does not change in value. Consider the second constraint involving z_i^t variables. Any feasible solution to the relaxed problem can be modified to be feasible in the original problem without altering the objective, as the objective only contains terms of the form $\sum_{t=1}^{T} z_i^t$. The proof is by contradiction, as if the solution cannot be modified to be feasible in the original problem, then it cannot be feasible in the relaxed problem. Similar arguments can be made for the third constraint involving Z_{ij}^t variables. Thus, an equivalent formulation of (6) is:

$$\begin{aligned}
\underline{C}(\mathbf{y}, \tilde{\mathbf{D}}) = \min_{\mathbf{z}^t, \mathbf{Z}^t} \quad & \sum_{t=1}^{T} \left[\sum_{i=1}^{N} p_s (D_{is}^t - z_i^t) + \sum_{j=1}^{N} p_o \left(D_{jo}^t - \sum_{i=1}^{N} Z_{ij}^t \right) \right. \\
& \left. + \sum_{i=1}^{N} s Z_{ii}^t + \sum_{i=1}^{N} \sum_{j=1, j\neq i}^{N} s_{ij} Z_{ij}^t \right] + \sum_{i=1}^{N} h \left(y_i - \sum_{t=1}^{T} z_i^t - \sum_{t=1}^{T}\sum_{j=1}^{N} Z_{ij}^t \right) \\
\text{s.t.} \quad & \sum_{t=1}^{T} z_i^t + \sum_{t=1}^{T}\sum_{j=1}^{N} Z_{ij}^t \leq y_i, \quad \forall i \in [N], \\
& \sum_{t=1}^{T} z_i^t \leq \sum_{t=1}^{T} D_{is}^t, \quad \forall i \in [N], \\
& \sum_{t=1}^{T}\sum_{i=1}^{N} Z_{ij}^t \leq \sum_{t=1}^{T} D_{jo}^t, \quad \forall j \in [N], \\
& \mathbf{z^t}, \mathbf{Z^t} \geq 0, \quad \forall t \in [T]
\end{aligned} \tag{33}$$

Applying the transformations completes the proof:

$$D_{is} \leftarrow \sum_{t=1}^{T} D_{is}^{t}, \quad D_{io} \leftarrow \sum_{t=1}^{T} D_{io}^{t}$$
$$z_i \leftarrow \sum_{t=1}^{T} z_i^{t}, \quad Z_{ij} \leftarrow \sum_{t=1}^{T} Z_{ij}^{t}$$

□

A.4. Proof of Proposition 3

Proof: Consider the linear program representation $\tilde{C}(\mathbf{y}, \tilde{\mathbf{D}})$, 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.

$$\begin{aligned}
\tilde{C}(\mathbf{y}, \tilde{\mathbf{D}}) = \min_{z_i, Z_{ij}} \sum_i h \quad & (y_i - z_i - \sum_j Z_{ij}) + \sum_i p_s(D_{is} - z_i) \\
& + \sum_i p_o(D_{io} - \sum_j Z_{ji}) + \sum_i sZ_{ii} + \sum_i \sum_{j \neq i} s_{ij} Z_{ij} \\
\text{subject to } z_i + \sum_j Z_{ij} & \leq y_i, \qquad \forall i \\
z_i & \leq D_{is}, \qquad \forall i \\
\sum_j Z_{ji} & \leq D_{io}, \qquad \forall i \\
z_i, Z_{ij} & \geq 0, \qquad \forall i, j
\end{aligned} \tag{34}$$

Note that $C^{IIP}(\mathbf{y}) = \mathbb{E}(\tilde{C}(\mathbf{y}, \tilde{\mathbf{D}}))$. 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 $\tilde{C}(\mathbf{y}, \tilde{\mathbf{D}})$ with respect to $\mathbf{y} = (y_1, y_2)$ can be written as:

$$\nabla_{\mathbf{y}} \tilde{C}(\mathbf{y}, \tilde{\mathbf{D}}) = (h, h)^T - \lambda(\mathbf{y}, \tilde{\mathbf{D}}) \tag{35}$$

where $\lambda(\mathbf{y}, \tilde{\mathbf{D}})$ is the dual-price vector corresponding to the constraints with y_1 and y_2 in (34). For a given $\mathbf{y}$, the 4-dimensional demand space $(D_{1s}, D_{1o}, D_{2s}, D_{2o})$ can be divided into domains $(\Omega_k(\mathbf{y}))_{k=1}^{20}$ such that in each domain, the optimal values of the decision variables z_i, z_{ii} and z_{ij} are linear in y_i, and hence the dual-price vector $\lambda(\mathbf{y}, \tilde{\mathbf{D}})$ is constant (refer to Appendix B for a discussion). The first-order conditions are:

$$0 = \nabla_{\mathbf{y}} C^{IIP}(\mathbf{y}) = \nabla_{\mathbf{y}} \mathbb{E}\left(\tilde{C}\left(\mathbf{y}, \tilde{\mathbf{D}}\right)\right) \tag{36}$$

We can interchange the gradient and expectation on the right hand side of Equation 36 (see Harrison and van Mieghem (1999) for a proof), and thus Equation 36 becomes

$$\begin{aligned}
0 = \nabla_{\mathbf{y}} C^{IIP}(\mathbf{y}) = \mathbb{E}_{\tilde{D}} \nabla_{\mathbf{y}} \tilde{C}\left(\mathbf{y}, \tilde{\mathbf{D}}\right) &= (h, h)^T - \mathbb{E}_{\tilde{D}} \lambda\left(\mathbf{y}, \tilde{\mathbf{D}}\right) \\
&= (h, h)^T - \sum_k \lambda^k \mathbb{P}\left(\Omega_k\left(\mathbf{y}\right)\right)
\end{aligned} \tag{37}$$

where λ^k is the constant $\lambda(\mathbf{y}, \tilde{\mathbf{D}})$ for $\tilde{\mathbf{D}} \in \Omega_k(\mathbf{y})$. □

A.5. Proof of Proposition 4

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 $s_{12} < h + p_o$.

A.6. Proof of Proposition 5

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 19, we have:

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

where Φ is the CDF of the standard normal distribution. The equality follows from the fact that $\mathbf{y^{DIP}}$ satisfies Equation 12, 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.7. Proof of Proposition 6

The proof follows from Govindarajan et al. (2020), by noting that the nested structure provides a closed-form expression for the total shipping cost, as opposed to a linear program, by summing the shipping costs in each level. The key difference from Govindarajan et al. (2020) is that the available inventory levels at location i is $(y_i - D_{is})^+)$, rather than just y_i, which gives rise to nested piecewise linear terms in the cost function.

In level 0, the shipping cost is $\sum_{i\in[N]} s \cdot \min\Big(D_{io}, (y_i - D_{is})^+\Big) = s\cdot \mathbf{e}^\top \mathbf{D_o} - \sum_{i\in[N]} s\cdot\Big(D_{io} - (y_i - D_{is})^+\Big)$. For any level $\ell \geq 1$, the number of fulfilled units of demand from regions in set $\mathcal{I}_k^{(\ell)}$ at level ℓ is

$$\underbrace{\sum_{m\in\mathcal{K}_k^{(\ell)}}\left(\sum_{i\in\mathcal{I}_m^{(\ell-1)}} D_{io} - \sum_{i\in\mathcal{I}_m^{(\ell-1)}}(y_i - D_{is})^+\right)^+}_{\text{unmet demand in } \mathcal{I}_k^\ell \text{ after level } \ell-1} - \underbrace{\left(\sum_{i\in\mathcal{I}_k^{(\ell)}} D_{io} - \sum_{i\in\mathcal{I}_k^{(\ell)}}(y_i - D_{is})^+\right)^+}_{\text{unmet demand in } \mathcal{I}_k^{(\ell)} \text{ after level } \ell}, \tag{38}$$

where $\mathcal{K}_k^{(\ell)}$ is the set of level $\ell-1$ children of set $\mathcal{I}_k^{(\ell)}$. Note that the per-unit cost of this fulfillment is $s_{\ell,k}$. The total cost is thus given by:

$$\begin{aligned} C^{IIP}(\mathbf{y}) = \mathbb{E}\Bigg[& h\cdot(\mathbf{e}^\top\mathbf{y} - \mathbf{e}^\top\mathbf{D})^+ + p_s\cdot\mathbf{e}^\top(\mathbf{D}_s - \mathbf{y})^+ + p_o\cdot\left(\mathbf{e}^\top\mathbf{D}_o - \mathbf{e}^\top(\mathbf{y}-\mathbf{D}_s)^+\right)^+ \\ & + s\cdot\mathbf{e}^\top\mathbf{D}_o + \sum_{\ell=0}^{L-2}\sum_{k\in[n_\ell]}(s_{\ell+1,m^{(\ell+1)}(k)} - s_{\ell,k})\cdot\left(\sum_{i\in\mathcal{I}_k^{(\ell)}} D_{io} - \sum_{i\in\mathcal{I}_k^{(\ell)}}(y_i - D_{is})^+\right)^+\Bigg] \end{aligned}$$

where $m^{(\ell+1)}(k) \in [n_{\ell+1}]$ is the level $\ell+1$ parent of $k \in [n_\ell]$. The proof is completed using the definition of η_ℓ as given in the Proposition statement. ■

A.8. Proof of Proposition 7

Proof of (1): The proof is similar to that of Proposition 4 and is hence omitted.

Proof of (2): C_2^{LB} is convex in the inventory levels, and its first order conditions can be solved to yield a heuristic solution y^{IIPH} characterized by the first order conditions:

$$(h+p_o-s)F_D\left(\sum_{j\in\mathcal{S}} y_j^{IIPH}\right)+(p_s-p_o+s)F_{D_{is}}(y_i^{IIPH})=p_s, \quad \forall i\in\mathcal{S} \tag{39}$$

Rewriting the above equation, we have:

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

Let $m=p_s-(h+p_o-s)\cdot F_{D_\mathcal{S}}\left(\sum_{j\in\mathcal{S}} y_j^{IIPH}\right)$. Thus, we have:

$$y_i^{IIPH}=F_{D_{is}}^{-1}\left(\frac{m}{p_s-p_o+s}\right) \tag{40}$$

Substituting the above equation into the definition of m, we have:

$$\sum_{j\in\mathcal{S}} F_{D_{is}}^{-1}\left(\frac{m}{p_s-p_o+s}\right)=F_D^{-1}\left(\frac{p_s-m}{h+p_o-s}\right) \tag{41}$$

The left hand side is increasing in m, whereas the right hand side is decreasing in m. Note that $p_s-(h+p_o-s)\le m\le p_s-(p_o-s)$. Due to the monotonicity of the left and right hand sides and their extreme values in this range, there must be a unique value of m that satisfies this equation, thus yielding a unique solution from (40). □

Proof of (3): Since we can solve for a unique solution for m in (41) which yields a unique solution $\mathbf{y^{IIPH}}$ from (40), it directly follows that stores stocks at the same critical fractile of their in-store demand. □

Proof of (4): 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.

Since the solution $\mathbf{y^{IIPH}}$ yields identical quantities at each location when the demands and costs are identical across locations, we simplify notation for the sake of this proof by replacing $C(\mathbf{y})$ by $C(y)$, where y is the inventory level at each location as specified by the solution $\mathbf{y}$. 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, $C^{IIP}(y)\le C^{IIP_c}(y)$ and $C^{LB'}(y)\le C^{LB'_c}(y)$ for any $y\ge 0$. Let $g(y,N)$ denote the cost incurred by N stores starting with inventory y each, without the option of cross-shipping:

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

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 $C_2^{LB}(y^{IIPH}) \leq C^{LB'}(y^{IIPH}) \leq C^{IIP}(y^{IIPH}) \leq C^{IIP_c}(y^{IIPH})$. We first show that $\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} \to 1$ as $N \to \infty$. We have:

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

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

$$\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} - 1 \leq \frac{f\left(\frac{\sqrt{2}}{N^{\frac{1}{4}}}\right)}{s} + \left(\frac{h+p_o-s}{s\mu_o\sqrt{N}}\right)\mathbb{E}\left(\sum_{i=1}^{\sqrt{N}} D_{io} - \left(y^{IIPH} - D_{is}\right)^+\right)^+ \tag{42}$$

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

$$y^{IIPH} \to F_s^{-1}\left(\frac{p_s - p_o + s - h}{p_s - p_o + s}\right) \in (0,\infty), \text{ as } N \to \infty \tag{43}$$

Proof: Lemma 1 is proved from the optimality equations of C^{LBN} (Equation 25) for identical stores:

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

From the above equation, when $h < p_o - s$, we have $p_s < 2\,(p_o-s)\,\mathbb{P}\left(\sum_{i=1}^{N} D_i \le Ny^{IIPH}\right) + (p_s - p_o + s)$. This simplifies to yield $y^{IIPH} > \mu$. Now, by applying the central limit theorem as $N \to \infty$ and $y^{IIPH} > \mu$, $\mathbb{P}\Big(\sum_{i=1}^{N} D_i/N \le y^{IIPH}\Big) \to 1$, 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 3. *When $h < \min(p_o - s,\ p_s - p_o + s)$, and the demands are bounded above as $D_{is} \le M_s$ and $D_{io} \le M_o$ for all i,*

$$\mathbb{P}\left(\sum_{i=1}^{\sqrt{N}} D_{io} > \sum_{i=1}^{\sqrt{N}} \left(y^{IIPH} - D_{is}\right)^+\right) \le \exp\left\{\frac{-2\sqrt{N}(y^{IIPH}-\mu)^2}{M_o+M_s}\right\} \tag{44}$$

Proof:

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

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 42 can be bounded as follows:

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

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

Thus, we have:

$$\begin{aligned}\frac{C^{IIP_c}(y^{IIPH})}{C^{LB'}(y^{IIPH})} &\le 1 + \frac{f\left(\frac{\sqrt{2}}{N^{\frac{1}{4}}}\right)}{s} + \left(\frac{h+p_o-s}{s\mu_o}\right)\left(M_o\sqrt{N}\exp\left\{\frac{-2\sqrt{N}(y^{IIPH}-\mu)^2}{M_o+M_s}\right\}\right)\\ &\to 1,\ as\ N\to\infty\end{aligned} \tag{45}$$

The next step is to show the C_2^{LB} is off by a constant factor from the $C^{LB'}$. From the proof of Proposition 4, the difference simplifies to:

$$C^{LB'}(y^{IIPH}) - C_2^{LB}(y^{IIPH}) \\ = (h+p_o-s)\,\mathbb{E}\left[\left(\sum_{i=1}^{N} D_{io} - \left(y^{IIPH} - D_{is}\right)^+\right)^+ + \sum_{i=1}^{N}\left(D_{is} - y^{IIPH}\right)^+ - \left(D - \sum_{i=1}^{N} y^{IIPH}\right)^+\right]$$

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

Similar to what was done to bound the second term in Equation 42, we can show that whenever the conditions in Lemma 2 are satisfied, $\mathbb{E}\left(\sum_{i=1}^{N} D_{io} - (y^{IIPH} - D_{is})^+\right)^+ \le M_o N \exp\left\{\frac{-2N(y^{IIPH}-\mu)^2}{M_o+M_s}\right\}$. Thus, we have:

$$C^{LB'}(y^{IIPH}) - C_2^{LB}(y^{IIPH}) \le (h+p_o-s)\left[M_o N \exp\left\{\frac{-2N(y^{IIPH}-\mu)^2}{M_o+M_s}\right\} + \sum_{i=1}^{N}\left(D_{is} - y^{IIPH}\right)^+\right]$$

Using $C_2^{LB}(y^{IIPH}) \ge s\mu_o N$ and $C_2^{LB}(y^{IIPH}) \ge (p_s - p_o + s)\sum_{i=1}^{N}(D_{is} - y^{IIPH})^+$, we have:

$$\frac{C^{LB'}(y^{IIPH})}{C^{IIPH}(y^{IIPH})} - 1 \le \left(\frac{h+p_o-s}{s\mu_o}\right)\left(M_o \exp\left\{\frac{-2N(y^{IIPH}-\mu)^2}{M_o+M_s}\right\}\right) + \left(\frac{h+p_o-s}{p_s-p_o+s}\right) \tag{46}$$

Thus, from Equations 45 and 46, as $N \to \infty$, we have

$$\frac{C^{IIP_c}(y^{IIPH})}{C_2^{LB}(y^{IIPH})} \le 1 + \frac{h+p_o-s}{p_s-p_o+s}$$
$$\Rightarrow \frac{C^{IIP}(\mathbf{y^{IIPH}})}{C^{IIP}(\mathbf{y^{IIP}})} \le \frac{h+p_s}{p_s-p_o+s}$$

The final step follows from $C^{IIP_c}(y^{IIPH}) \ge C^{IIP}(y^{IIPH})$, and $C_2^{LB}(\mathbf{y^{IIPH}}) \le C^{IIP}(\mathbf{y^{IIP}})$. □

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.

Appendix B: 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 3.1.3) 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 (34), Table 5 shows the different cases that are possible given y_1 and y_2. From these cases, the independent demand regions are listed in Table 6 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. The dual prices λ_1, λ_2 are the shadow prices of the constraints which contain y_1 and y_2 respectively, namely the first set of constraints $z_i + \sum_{j=1}^{2} z_{ij} \le y_i, \forall i$ in the linear program in (34), and can be obtain in a standard fashion from linear programming theory. For example, for the demand regions with the case D1, that is, $y_1 \ge D_1 + D_{2o}$, 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 $z_1 + \sum_{j=1}^{2} z_{1j} \le y_1$ will not bind, and hence $\lambda_1 = 0$.

Table 5 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	$y_1 + y_2 < D_1 + D_2$	$y_1 + y_2 \geq D_1 + D_2$		

Table 6 Table showing the various demand regions and the corresponding constant dual-prices. (underlined notation indicates redundant cases)

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

Appendix C: Heuristic based on Constant Shipping Costs for a Network of Omnichannel Stores and OFCs

We obtain the heuristic solution y^{IIPH} for multiple locations with $\mathcal{S}_o \neq \emptyset$ by calculating order quantities for the OFCs separately, and using them in Equation 39 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 $D_{\mathcal{S}_o} = \sum_{i \in \mathcal{S}_o} D_{io}$.

$$\sum_{j \in \mathcal{S}_o} y_j^{IIPH} = F_{D_{\mathcal{S}_o}}^{-1}\left(\frac{p_o - s}{h + p_o - s}\right) \tag{47}$$

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 shipping cost (s_o) compared to the stores (s), as the inventory calculation for the OFCs is done separately.

To calculate the individual order quantities at the OFCs, 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). Each unit from $\sum_{j\in\mathcal{S}_o} y_j^{IIPH}$ 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 48 to determine order-up-to levels for other omnichannel stores.

$$(h+p_o-s)F_{D_{\mathcal{S}}}\left(\sum_{j\in\mathcal{S}} y_j^{IIPH}\right) + (p_s - p_o + s)\,F_{D_{is}}\left(y_i^{IIPH}\right) = p_s, \quad \forall i \in \mathcal{S}_{so} \tag{48}$$

Calculating this heuristic solution $\mathbf{y^{IIPH}}$ is also computationally fast, as Proposition 7(3) still applies to Equation 48. The cost of the heuristic solution is given by $C^{IIPH} = C^{IIP}(\mathbf{y^{IIPH}})$. We capture the effect of virtual pooling among the facilities in this heuristic, and the systematic approach is shown in Algorithm 2.

Algorithm 2 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 $MC_i\left(y_i^{IIPH}\right) = -\left(p_o - s\right)\left(1 - F_{D_{io}}(y_i^{IIPH})\right) + hF_{D_{io}}\left(y_i^{IIPH}\right)$
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: $(h+p_o-s)\,F_{D_{\mathcal{S}}}\left(\sum_{j\in\mathcal{S}} y_j^{IIPH}\right) + (p_s - p_o + s)\,F_{D_{is}}\left(y_i^{IIPH}\right) = p_s$, $\forall i \in \mathcal{S}_{so}$.

Appendix D: 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. We take 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 mean in-store and online total demands over the horizon are calculated as $1-\alpha$ and α proportions respectively of this mean market size in each region. 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 (2016).

D.1. Simulation Procedure

A brief overview of the simulation is listed below:

1. The parameters for demands in each fulfillment period are calculated based on demands over the horizon 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 12 and Algorithm 1 respectively.
2. We generate a sample of size 10^4, where each sample is a realization of demands over the entire selling horizon, although fulfillment decisions in each fulfillment period 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 27. 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 period, penalty and shipping costs are calculated. The ending inventory at a location becomes the starting inventory for the next fulfillment period.
7. The total cost is the sum of the costs in each fulfillment period over the selling horizon.