This article was downloaded by: [128.2.149.108] On: 10 March 2022, At: 14:02
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

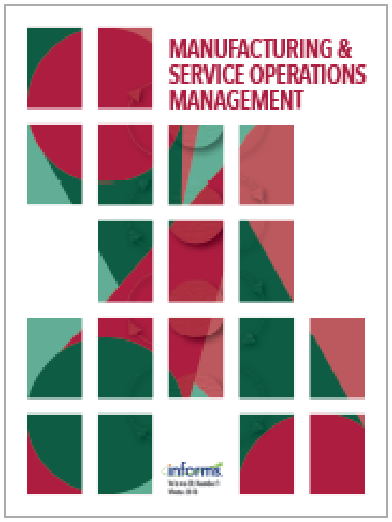

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Effective Online Order Acceptance Policies for Omnichannel Fulfillment

Su Jia, Jeremy Karp, R. Ravi, Sridhar Tayur

To cite this article:
Su Jia, Jeremy Karp, R. Ravi, Sridhar Tayur (2021) Effective Online Order Acceptance Policies for Omnichannel Fulfillment. Manufacturing & Service Operations Management

Published online in Articles in Advance 01 Dec 2021

. https://doi.org/10.1287/msom.2021.1024

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2021, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–14
http://pubsonline.informs.org/journal/msom
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Effective Online Order Acceptance Policies for Omnichannel Fulfillment

Su Jia,[a] Jeremy Karp,[b] R. Ravi,[a] Sridhar Tayur[a]

[a] Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; [b] Marketplace Optimization, Lyft Inc, San Francisco, California

Contact: sjia1@andrew.cmu.edu, https://orcid.org/0000-0002-8948-0834 (SJ); jeremy.a.karp@gmail.com (JK); ravi@andrew.cmu.edu, https://orcid.org/0000-0001-7603-1207 (RR); stayur@andrew.cmu.edu, https://orcid.org/0000-0002-8008-400X (ST)

Received: April 30, 2019
Revised: February 1, 2021; May 28, 2021
Accepted: June 3, 2021
Published Online in Articles in Advance: December 1, 2021

https://doi.org/10.1287/msom.2021.1024

Copyright: © 2021 INFORMS

Abstract. ***Problem definition:*** Omnichannel retailing has led to the use of traditional stores as fulfillment centers for online orders. Omnichannel fulfillment problems have two components: (1) accepting a certain number of online orders prior to seeing store demands and (2) satisfying (or filling) some of these accepted online demands as efficiently as possible with any leftover inventory after store demands have been met. Hence, there is a fundamental trade-off between store cancellations of accepted online orders and potentially increased profits because of more acceptances of online orders. We study this joint problem of online order acceptance and fulfillment (including cancellations) to minimize total costs, including shipping charges and cancellation penalties in single-period and limited multiperiod settings. ***Academic/practical relevance:*** Despite the growing importance of omnichannel fulfillment via online orders, our work provides the first study incorporating cancellation penalties along with fulfillment costs. ***Methodology:*** We build a two-stage stochastic model. In the first stage, the retailer sets a policy specifying which online orders it will accept. The second stage represents the process of fulfilling online orders after the uncertain quantities of in-store purchases are revealed. We analyze threshold policies that accept online orders as long as the inventories are above a global threshold, a local threshold per region, or a hybrid. ***Results:*** For a single period, total costs are unimodal as a function of the global threshold and unimodal as a function of a single local threshold holding all other local thresholds at constant values, motivating a gradient search algorithm. Reformulating as an appropriate linear program with network flow structure, we estimate the derivative (using infinitesimal perturbation analysis) of the total cost as a function of the thresholds. We validate the performance of the threshold policies empirically using data from a high-end North American retailer. Our two-location experiments demonstrate that local thresholds perform better than global thresholds in a wide variety of settings. Conversely, in a narrow region with negatively correlated online demand between locations and very low shipping costs, global threshold outperforms local thresholds. A hybrid policy only marginally improves on the better of the two. In multiple periods, we study one- and two-location models and provide insights into effective solution methods for the general case. ***Managerial implications:*** Our methods provide effective algorithms to manage fulfillment costs for online orders, demonstrating a significant reduction over policies that treat each location separately and reflecting the significant advantage of incorporating shipping in computing thresholds. Numerical studies provide insights as to why local thresholds perform well in a wide variety of situations.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2021.1024.

Keywords: omnichannel retail • stochastic programming • infinitesimal perturbation analysis • dynamic programming

1. Introduction

Omnichannel retailing, the combination of online and traditional store channels, uses traditional stores as shipping centers for originating online orders and customer pickup points for online orders, thus using the inventory at the store in a pooled manner across channels. Online orders arrive over time, but satisfying them (with the store inventory) is usually done in a periodic manner after also accounting for in-store orders, which are given strict priority. Thus, a key decision is the number of online orders to accept, with the understanding that some may not be satisfied (as when the leftover inventory after filling in-store demands is zero) and may have to be cancelled by the retailer.[1] Using proprietary data from high-end North American retailers, we observed cancellation rates above 20%, whereas inventory positions of the most popular Stock-Keeping Units (SKUs) across physical store locations varied from a few units to over a hundred units, motivating the potential for better-performing joint acceptance and fulfillment policies. We study a new set of research questions related to acceptance and fulfillment of these online orders in omnichannel retail operations, taking into consideration shipping costs when they are filled and cancellation costs when they are not.

1.1. Omnichannel Fulfillment Model

Fixed exogenous inventory $\mathcal{I}_i$ is available at each location i in a system of N locations. At each location i, there are two streams of demand, online ($\mathcal{D}_i^O$) and physical ($\mathcal{D}_i^P$), which both draw from the same pool of inventory. Physical demand is fulfilled with higher priority than online demand, and online orders are cancelled if there is not sufficient inventory to fill them. There is a cost, c, associated with canceling an order, and there is also a penalty cost, p, associated with not accepting an order that could have been filled. The shipping cost per unit is s_{ij} from location i to location j.

The goal for the retailer is to set a policy that minimizes total costs in expectation. We capture this process through a two-stage stochastic model. The first stage of the problem occurs as online orders arrive at the retailer. The retailer must decide whether to accept or reject each order as it arrives, in an online manner. Many retailers already use threshold policies to manage their online sales channels, so it is natural to focus on this policy class. More complex policies might utilize the arrival times of orders, but this is not our focus in this analysis. The first stage concludes after all online orders have arrived and are accepted or rejected by the retailer.

We consider two types of threshold policies, local thresholds and global thresholds. Local threshold policies have a parameter for each location, allowing the retailer fine-tuned control over which areas are accepting online orders. Global threshold policies have a single parameter for the full network of stores. Global threshold policies allow the retailer more control over the total number of online orders accepted but less fine-tuned control than with local thresholds over which orders are accepted.[2]

Definition 1. A local threshold policy $[S_1, \ldots, S_n]$ accepts the first S_i online orders from location $i \in [n]$ and rejects all remaining orders. A global threshold policy S accepts the first S online orders (from all locations) and rejects all remaining orders.

Setting the acceptance thresholds too high increases the risk of eventual cancellation of these accepted online orders, whereas setting the thresholds to conservatively low levels results in incurring unnecessary rejection penalty costs on leftover inventory that could have been used to fill rejected online orders, thus setting up a trade-off for the optimal thresholds.

After the first stage concludes, the retailer learns the amount of in-store demand it received as the online orders arrived. The retailer then must decide whether to cancel or fulfill each accepted online order and from which store inventory will be used to fill these orders. This second stage can be naturally formulated as a network flow optimization problem:

$$\begin{aligned} \text{minimize } & p \min\left(\sum_{i=1}^{n}\left(\mathcal{I}_i - \min(\mathcal{I}_i, \mathcal{D}_i^P) - \sum_{j=1}^{n} F_{ij}\right),\right. \\ & \left.\sum_{i=1}^{n}(\mathcal{D}_i^O - A_i)\right) + \sum_{i=1}^{n}\left(c_i C_i + \sum_{j=1}^{n} s_{ji} F_{ji}\right) \\ \text{such that } & \min(\mathcal{D}_i^P, \mathcal{I}_i) + R_i + \sum_{j=1}^{n} F_{ij} = \mathcal{I}_i, \quad \forall i \in [n] \\ & C_i + \sum_{j=1}^{n} F_{ji} = A_i, \quad \forall i \in [n] \\ & C_i, R_i, F_{ij} \geq 0, \quad \forall i, j. \end{aligned} \tag{1}$$

In this formulation, A_i is the number of online orders accepted from location i in the first stage, and it is set differently in local and global policies. For local threshold policies, we have each region i accepting up to S_i orders giving $A_i = \min\{\mathcal{D}_i^O, S_i\}$. In the global threshold policy, we set $A_i = \mathcal{D}_i^O$ as long as $\sum_i \mathcal{D}_i^O \leq S$, the global threshold. Otherwise, for a given set of online demand realizations $\{\mathcal{D}_i^O\}_i$, there are many ways of truncating them to accept only S. For concreteness, we assume that the demand arrivals occur uniformly over the first stage and hence, scale the demands proportionally so that the total is S. In other words, when $\sum_i \mathcal{D}_i^O > S$, we set $A_i = \frac{\mathcal{D}_i^O}{\sum_i \mathcal{D}_i^O} \cdot S$. We denote the objective function for the second-stage program as $G(S_1, \ldots, S_n)$ in the local threshold case and $G(S)$ in the global threshold case. The first-stage problem is to set the corresponding thresholds to minimize the expected value of G in both cases.

Consider first the objective function:

$$p \min\left(\sum_{i=1}^{n}\left(\mathcal{I}_i - \min\left(\mathcal{I}_i, \mathcal{D}_i^P\right) - \sum_{j=1}^{n} F_{ij}\right), \sum_{i=1}^{n}\left(\mathcal{D}_i^O - A_i\right)\right) + \sum_{i=1}^{n}\left(c_i C_i + \sum_{j=1}^{n} s_{ji} F_{ji}\right).$$

The expression $\sum_{i=1}^{n}(\mathcal{I}_i - \min(\mathcal{I}_i, \mathcal{D}_i^P) - \sum_{j=1}^{n} F_{ij})$ is the amount of remaining inventory after all orders have been fulfilled or cancelled. $\mathcal{I}_i$ is the starting inventory at location i, $\mathcal{D}_i^P$ is the in-store demand at location i, and F_{ij} is the number of filled online orders received at location j and filled from inventory at location i. The expression $\sum_{i=1}^{n}(\mathcal{D}_i^O - A_i)$ is the number of online orders that were rejected. $\mathcal{D}_i^O$ is the amount of online demand at location i, and A_i is the number of online orders accepted from location i in the first stage. Consequently, $\min(\sum_{i=1}^{n}(\mathcal{I}_i - \min(\mathcal{I}_i, \mathcal{D}_i^P) - \sum_{j=1}^{n} F_{ij}),$ $\sum_{i=1}^{n}(\mathcal{D}_i^O - A_i))$ is the number of rejected online orders that could have been fulfilled had they been accepted. The objective function assigns a cost of p to each of these orders, reflecting the sale price of the item. The selling price is an upper bound on the opportunity cost because an accepted order may incur shipping costs in being fulfilled from leftover inventory in another location. The expression $\sum_{i=1}^{n} c_i C_i$ reflects the sum of all cancellation penalties. c_i is the cost parameter of a cancelled order from location i, and C_i is the decision variable for the number of online orders cancelled from location i. We assume that the cancellation cost is high enough that the retailer would never want to cancel any order it could possibly fill (i.e., $c > \max\{s_{ji}\}$).[3] Lastly, the expression $\sum_{i=1}^{n}\sum_{j=1}^{n} s_{ji}F_{ji}$ represents the shipping costs for all online orders that were accepted and fulfilled. s_{ji} is the unit shipping cost from location j to i, and F_{ji} is the decision variable for the number of online orders filled from inventory at j and shipped to customers at i.

The constraints of the form $\min(\mathcal{D}_i^P, \mathcal{I}_i) + R_i + \sum_{j=1}^{n} F_{ij} = \mathcal{I}_i$ express that all inventory at location i must be used to fulfill in-store demand, be saved, or be used to fulfill online demand. R_i is a decision variable reflecting the amount of inventory that is left over. The constraints of the form $C_i + \sum_{j=1}^{n} F_{ji} = A_i$ reflect that all accepted orders must be either cancelled or fulfilled. This is the omnichannel fulfillment problem we study in the rest of this paper.

1.2. Summary of Contributions

1. We formulate an analytical model for omnichannel fulfillment that incorporates uncertainty because of inventory pooling across sales channels as a multilocation, two-stage stochastic optimization problem.

2. We introduce local threshold and global threshold policy classes for the first-stage problem and present a sampling-based optimization method to set these policies.

3. We show that the expected retailer costs are unimodal as a function of a single global threshold or of a single local threshold holding all other local thresholds constant, allowing us to do a gradient-based search (Section 2 for a single-location case and Section 3 for multiple locations).

4. Our optimization method uses infinitesimal perturbation analysis (IPA) to estimate derivatives of the objective function with respect to these threshold policy parameters (Section 4). To obtain derivative estimates for the IPA, we rely on the dual values of a linear program related to the second-stage problem, which we reformulate appropriately to derive the required estimates.

5. We present empirical results from numerical experiments to provide insights and demonstrate the effectiveness of policies generated by our methods (Section 5). Through a partnership with a retail analytics firm, we use retail industry data to generate realistic problem instances. We conduct a series of experiments on two-location instances to demonstrate how certain instance attributes lead to strong performance of one class of threshold attributes relative to the other. We find that local threshold policies have better performance than global threshold policies in a majority of situations. However, this advantage is diminished and sometime reversed, especially at low shipping costs, when online demands are negatively correlated and variance of physical demand is high.

6. We use retail data from a high-end North American retailer to formulate realistic full-network problem instances through which we show that both local thresholds and global thresholds achieve a considerable reduction in costs compared with other simple baseline policies currently employed by the retailer.

7. We formulate an extension to multiple periods. We develop exact methods for single-location, multiperiod, and the two-location, two-period cases. We propose a look-ahead heuristic and compare it with myopic and optimal policies.

1.3. Related Work

The recent edited volume by Gallino and Moreno (2019) provides a comprehensive survey of literature related to omnichannel operations, with the chapter by Jasin et al. therein being the closest to this paper. Our work is related to the following streams of literature: (a) overbooking, (b) fulfillment of online orders from stores, and (c) multilocation transshipment. We briefly review here the closest papers to our work.

Harsha et al. (2019) models how customer demand responds to changes in price, allowing the retailer to optimally set clearance prices in all sales channels by solving an integer program. In our work, prices are fixed, and our focus is on fulfillment with uncertain inventory and stochastic demand. Other aspects of omnichannel retailing, such as the costs and benefits of "buy online and pick up in store" policies (Gao and Su 2017), information sharing (Gao and Su 2016), inventory optimization (Govindarajan et al. 2017), and multichannel price optimization (Cattani et al. 2006, Harsha et al. 2019), have been studied by the operations management community, but our study is the first attempt to formulate and study stochastic models of cancellations caused by omnichannel fulfillment, where inventory at a location (which is a retail store and not a warehouse) is shared across the walk-in customers who are given priority in any period and online orders that have been accepted in that period. Note that although the store inventory is shared across the online and physical walk-in demands, the strict prioritization of the latter makes it less of an inventory rationing problem (or the uncertain/phantom inventory situation) and more like an overbooking problem. Indeed, this feature of accepting online orders that risk not being fulfilled, even considering transshipments from other locations, is reminiscent of the multiclass overbooking model with substitutions of Karaesmen and Van Ryzin (2004); however, the specifics of the mechanisms and the related stochastic considerations are different. Furthermore, our analyses use different techniques that have been successful in other areas within operations management, including transshipment problems by Herer et al. (2006) and sensitivity analysis by Glasserman and Tayur (1995).

Models that study fulfillment (Acimovic and Graves 2014, Jasin and Sinha 2015) in an e-commerce setting in real time so as to minimize the cost of picking and shipping items do not consider the possibility of optimally choosing which online orders to fulfill, even though they typically study multiple products and the effect of consolidation of different products in an order during fulfillment. The models of Alishah et al. (2015) consider the omnichannel setting of fulfilling an order from a combination of a single omnichannel fulfillment location and a single offline store in a continuous time setting and show that for the rationing decision of whether to use offline inventory for online demand, threshold-type policies are optimal for fulfillment. DeValve et al. (2021) consider the value of adding flexibility to a network in terms of cost reduction for spillover fulfillment. Bayram and Cesaret (2021) considered the fulfillment problem with one fulfillment center and multiple stores and proposed an efficient Markov Decision Process–based heuristic.

The multilocation transshipment problem has been previously formulated and studied in Herer and Rashit (1999) and Herer et al. (2006). This work presents a stochastic multiperiod model, and the theoretical result is that optimal inventory replenishment policies in this model are "order-up-to S" policies. Reduced costs of a linear program are used in this paper to iteratively update safety stock values. The optimization algorithms presented in our paper use dual values of a linear program to update our threshold policies, although the methodology of our work and the problem context are quite different from the settings of Herer and Rashit (1999) and Herer et al. (2006).

Our models of omnichannel fulfillment thus incorporate some of the complexities considered in the multilocation transshipment problem and the multiclass overbooking problem, as they include order cancellations as a major component that are the result of accepting online orders, which themselves should take into account the possibility of transshipments.

2. Single-Location Model

Note that for a single-location instance, local threshold and global threshold policies are equivalent. In this section, we remove the subscripts denoting store location. For simplicity, we assume that the shipping cost of fulfilling an online order from the store in this location is zero (i.e., $s = 0$). The second-stage problem is straightforward: $\min(\mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P), \min(\mathcal{O}^O, S))$ accepted online orders are filled, and the rest are cancelled. Using this observation, we can rewrite the optimization problem (1) in a simpler form (Figure 1). The objective function simplifies to minimizing $p\min(\mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P) - F, \mathcal{D}^O - \min(\mathcal{D}^O, S)) + cC$.

Proposition 1 (Proof in Online Appendix A). *For a single-location instance of the omnichannel fulfillment model,* $F = \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P))$ *is satisfied in any optimal solution.*

By Proposition 1 and the constraint $C + F = A$, we see that for any optimal solution, $C = \min(\mathcal{D}^O, S) - \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P))$ and $F = \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P))$. Substituting these values, the value of the objective function of (1) at an optimal solution becomes

$$\begin{aligned} &p\min(\mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P) \\ &\quad - \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P)), \mathcal{D}^O - \min(\mathcal{D}^O, S)) \\ &\quad + c(\min(\mathcal{D}^O, S) - \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P))). \end{aligned}$$

Consequently, we can combine the two stages of decision making:

$$\begin{aligned} S^* = \min_S\ & E_{\mathcal{D}^O, \mathcal{D}^P}[p \cdot \min[\mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P) \\ & - \min(\mathcal{D}^O, S, \mathcal{I} - \min(\mathcal{I}, \mathcal{D}^P)), \mathcal{D}^O - \min(\mathcal{D}^O, S)] \\ & + c \cdot (\min[\mathcal{D}^O, S] - \min[\mathcal{D}^O, S, \mathcal{I} - \min[\mathcal{I}, \mathcal{D}^P]])]. \end{aligned}$$

This has a closed-form optimal solution.

Theorem 1 (Proof in Online Appendix A). *For the model,* $S^* = \left(\mathcal{I} - F_P^{-1}\left(\frac{c}{c+p}\right)\right)^+$.

Note that the optimal threshold for this problem takes a similar form to the solution to the newsvendor problem.

An important consequence of this theorem is that the optimal threshold depends only on the distribution of physical demand, not that of online demand.

The proof of the optimal threshold also provides insight on the structure in this optimization problem. The theorem is proved by showing that $G(S)$, the expected value of the optimization problem as a function of threshold S, is unimodal and finding the point where the derivative of $G(S)$ with respect to S changes signs. The derivative of $G(S)$ reduces to

$$P[S < \mathcal{D}^O](cP[\mathcal{D}^P \geq \mathcal{I} - S] - pP[\mathcal{D}^P < \mathcal{I} - S]).$$

Figure 1. (Color online) Examples of Outcomes in a Single-Location Model

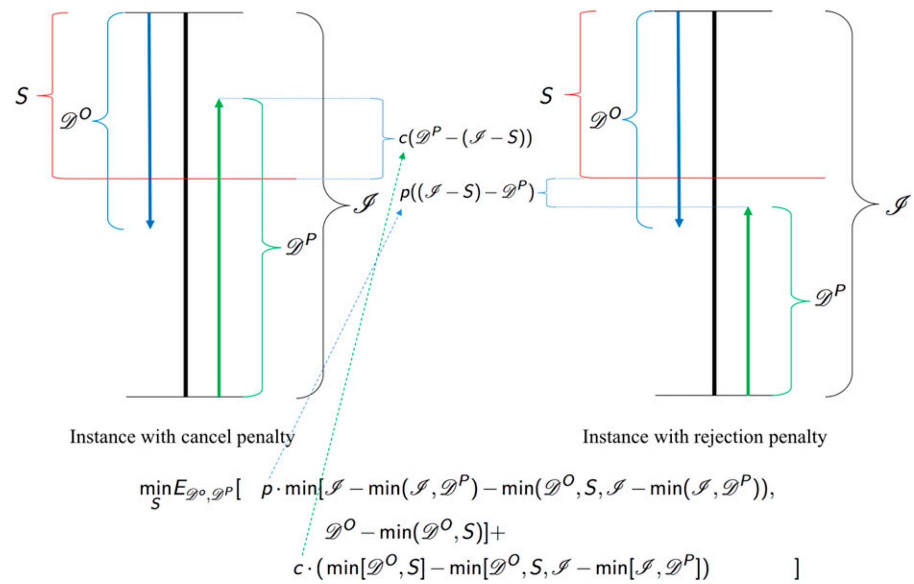

The factor in this expression, $cP[\mathcal{D}^P \geq \mathcal{I} - S] - pP[\mathcal{D}^P < \mathcal{I} - S]$, is nearly identical to the derivative of the classical newsvendor problem's expected objective with respect to the quantity purchased. This reveals a close connection between the omnichannel fulfillment problem studied and the classical newsvendor model. Additionally, it is the presence of the other multiplicative factor, $P[S < \mathcal{D}^O]$, that makes the function $G(S)$ nonconvex (although still unimodal).

3. Structural Properties of the Multiple-Location Model

Theorem 2 states that the expected objective function value as a function of a global threshold is unimodal, and Theorem 4 provides an efficient optimization algorithm for our fulfillment model, contingent on an oracle that produces unbiased gradient estimates. We use the same proof technique to show that when all but one S_i of a local threshold policy are fixed, the expected value of the model as a function of the free local threshold parameter is unimodal. In Section 4, we show how to compute these estimates, demonstrating an efficient optimization procedure to the thresholds.

The intuition behind these proofs comes from extending our observations about single-location instances in Section 2 to the more general multiple-location setting. We observed that the optimal threshold policy for single-location instances is the $\frac{c}{c+p}$ fractile of the in-store demand distribution. The online demand distribution influences the shape of the expected cost as a function of the threshold, although it does not influence the value of the optimal threshold. An informative way to think about this property is to consider the marginal effect on cost with respect to the threshold. For realizations where online demand is below the threshold, this marginal effect is zero, and this marginal effect will have the same nonzero value for all realizations where online demand is above the threshold. Then, the sign of this marginal effect on cost is determined entirely by the demand distribution restricted to realizations where online demand is greater than the threshold value.

We lift these observations to the multiple-store setting and use them to analyze the marginal effect of increasing the threshold. This marginal effect on penalties for missed sale opportunities is negative and decreasing in magnitude, and we use Lemma 1 to establish that the marginal effect on fulfillment costs (including cancellation costs) is also always increasing in the threshold. These properties are sufficient to argue that the total expected cost as a function of global threshold S (or local threshold S_k at location k with all other local threshold elements fixed) is unimodal.

Lemma 1 (Proof in Online Appendix B). *In a minimum-cost, single-commodity flow problem with multiple sources, one sink, and integer supplies, demands, and capacities, the objective value of a minimum cost feasible flow as a function of the supplies at the source nodes is supermodular.*

Theorem 2 (Proof in Online Appendix B). *$G(S)$, the expected value of the objective function of the omnichannel fulfillment model as a function of global threshold S, is unimodal.*

We apply a similar argument to show that when all but one S_i of a local threshold policy are fixed, the expected value of the model as a function of S_i is unimodal.

Theorem 3 (Proof in Online Appendix B). $G(S_k)$, *the expected value of the objective function of the omnichannel fulfillment model as a function of local threshold* S_k, *is unimodal when all other local threshold parameters,* S_j *for* $j \neq k$, *are fixed values.*

Theorems 2 and 3 establish that the expected value of the objective function as a function of a single global or local threshold variable is unimodal. We conclude this section by demonstrating that this property is sufficient to show that global threshold policies and single thresholds of local threshold policies can be set optimally and efficiently. Let $L(S)$ and $L(S_k)$ be the linear interpolation of the integer values of functions $L(S)$ and $L(S_k)$, defined in Theorems 2 and 3.

Lemma 2. *$L(S)$ and $L(S_k)$ are quasiconvex and Lipschitz continuous.*

Proof. Observe that $L(S)$ and $L(S_k)$ are unimodal functions because $G(S)$ and $G(S_k)$ are unimodal functions and have global minima at integer values. The quasiconvexity of functions $L(S)$ and $L(S_k)$ follows trivially from Theorems 2 and 3, as these are unimodal functions of a single variable. The Lipschitz continuity of these functions is also trivial, as the absolute value of the slope of these functions cannot exceed the maximum of cost parameters c, p, and s_{ij}. □

These technical conditions allow us to apply theorem 5.1 of Hazan et al. (2015) to prove that $L(S)$ and $L(S_k)$ can be efficiently minimized. The theorem proves that the stochastic normalized gradient descent algorithm will find an ϵ-optimal minimum $L(S)$ and $L(S_k)$ with poly$(\frac{1}{\epsilon})$ unbiased gradient estimates and optimization steps.

Theorem 4. *An ϵ-optimal minimum $L(S)$ and $L(S_k)$ can be obtained with* poly$(\frac{1}{\epsilon})$ *unbiased gradient estimates and optimization steps by the stochastic normalized gradient descent algorithm.*

4. Optimization Methods for the Single-Period, Multiple-Store Model

We now present an IPA algorithm that converges to optimal policies for certain policy classes. This IPA method can be used to obtain unbiased gradient estimates of global threshold and local threshold policy parameters for the omnichannel fulfillment model. Our key insight is that we can use dual values of an linear program (LP) related to the second-stage problem to produce unbiased estimates of the derivative of the objective function in the threshold parameters. We apply this method to set both global and local threshold policies.

4.1. Second-Stage Assignment Problem

Recall the original assignment problem modeled in LP (1) defined in Section 1.1. We would like to compute derivative estimates using the dual values of this constraint set. However, we cannot do this with this LP because the value that changes when relaxing the right-hand side of this constraint, A_i, also appears in the LP's objective function.

We formulate an alternative LP with equivalent optimal solutions, in which we can compute gradient estimates more straightforwardly:

$$\max \sum_{i=1}^{n}\left(p\min(\mathcal{D}_i^P,\mathcal{I}_i) - c_iC_i + \sum_{j=1}^{n}(p - s_{ji})F_{ji}\right)$$
$$\text{such that}\quad \min(\mathcal{D}_i^P,\mathcal{I}_i) + R_i + \sum_{j=1}^{n} F_{ij} = \mathcal{I}_i,\quad \forall i \in [n]$$
$$C_i + \sum_{j=1}^{n} F_{ji} = A_i,\quad \forall i \in [n]$$
$$C_i, R_i, F_{ij} \geq 0,\quad \forall i, j. \tag{2}$$

We can immediately observe that it has an economic interpretation consistent with that of the original LP (1). The first term in the objective function, $\sum_{i=1}^{n} p\min(\mathcal{D}_i^P,\mathcal{I}_i)$, reflects a profit of p for each unit sold in store. The remaining terms of the objective function, $-c_iC_i + \sum_{j=1}^{n}(p - s_{ji})F_{ji}$, reflect a profit of p for each unit sold online, with costs deducted for cancellations and shipping costs. This interpretation of LP (2) may at first seem counterintuitive because p was originally defined as the penalty cost for missed sales. However, we can observe that the damage incurred by the retailer from missing a sale is the profit it would get from making an additional sale, which explains why the economic interpretations of LPs (1) and (2) are consistent with each other. Both interpretations implicitly assume that the retailer gets no value from holding on to excess inventory. In particular, recall that we assume that if any cancellations occur, then there is no remaining inventory. This is equivalent to assuming that $\max_{i,j\in[n]} s_{ij} < p + c$. In other words, the maximum ship cost in the network is small enough that it is always preferable to fill an online order (paying at most $\max_{i,j\in[n]} s_{ij}$) rather than cancel the order (and incur a cancellation cost of c plus an additional missed sale penalty of p).

Proposition 2 (Proof in Online Appendix C). *As long as* $\max_{i,j\in[n]} s_{ij} < p + c$, *for any optimal solution to the original minimization LP (1), there is an optimal solution to the maximization LP (2) that yields an identical assignment of orders to stores.*

This correspondence means that we can find the optimal thresholds for the maximization problem and use these to compute the optimal (expected) costs for the minimization problem. In this way, we can use the dual values of the maximization LP to optimize our policy in the first-stage problem.

The original minimization problem placed assumptions that the cancellation cost must be high enough that the retailer would never want to cancel an order they could possibly fill. Without this assumption, it is unclear what it means to "unnecessarily reject" an order; there could be orders that could theoretically be filled but doing so would not be the profit-maximizing decision for the retailer. This new maximization LP formulation models the full profits received by the retailer from both online and in-store sales, rather than just costs, so it is no longer necessary to make this assumption if we wish to use only this formulation of the second-stage problem.

Maximization LP (2) provides sensitivity information from the constraints $C_i + \sum_{j=1}^{n} F_{ji} = A_i,\ \forall i \in [n]$ using LP dual values. To access this sensitivity information, we will express these constraints as inequality constraints: $C_i + \sum_{j=1}^{n} F_{ji} \le A_i,\ \forall i \in [n]$ and $C_i + \sum_{j=1}^{n} F_{ji} \ge A_i,\ \forall i \in [n]$. Then, we will dualize the first set of inequalities to obtain the following LP:

$$\max \sum_{i=1}^{n}\Big(p \min(\mathcal{D}_i^P, \mathcal{I}_i) - (c_i + M)C_i + MA_i + \sum_{j=1}^{n}(p - s_{ji} - M)F_{ji}\Big)$$

$$\text{such that } \min(\mathcal{D}_i^P, \mathcal{I}_i) + R_i + \sum_{j=1}^{n} F_{ij} = \mathcal{I}_i,\ \forall i \in [n]$$

$$C_i + \sum_{j=1}^{n} F_{ji} \ge A_i,\ \forall i \in [n]$$

$$C_i, R_i, F_{ij} \ge 0,\ \forall i, j. \tag{3}$$

We show in Proposition 3 that this LP (3) has the same optimal solution as the original second-stage and maximization LPs. This means we will be able to use the dual values of constraints $C_i + \sum_{j=1}^{n} F_{ji} \ge A_i,\ \forall i \in [n]$ to obtain gradient estimates for the omnichannel fulfillment model.

Proposition 3 (Proof in Online Appendix C). *The linear program* (3) *has the same optimal solution as the linear program* (2) *when* $M > p$.

Using the network flow formulation of the original cost-minimization problem, we obtain Proposition 4.

Proposition 4 (Proof in Online Appendix C). *The minimization LP* (1) *is integral.*

4.1.1. Infinitesimal Perturbation Analysis Method.

Our IPA algorithm can be viewed as a stochastic normalized gradient descent method. We begin by specifying a starting policy P and a value U, which will be the number of samples we use to compute a single gradient estimation iteration. U demand samples are drawn, and online orders are accepted and rejected according to policy P for each of the U samples. We assume that policy P is a local or global threshold policy, but the method may apply to additional policy classes (such as a hybrid policy we define at the end of this section). Then, we solve the maximization assignment LP to fulfill the accepted orders in each of the demand samples. We use the dual values of the maximization assignment LP to compute unbiased estimates of the gradients of the fulfillment profit with respect to the policy parameters (S_i in the case of local thresholds). We then update the threshold parameters with their normalized gradients, multiplied by a step size value.

4.1.2. Local Threshold Derivative Estimates.

We compute derivative estimates by looking at the dual values corresponding to the LP constraints $C_i + \sum_{j=1}^{n} F_{ji} \ge A_i,\ \forall i \in [n]$ from linear program (3), where A_i and C_i are the number of accepted and cancelled online orders at location i, respectively, and F_{ji} is the number of online orders at location i filled from inventory from store j. The dual value from one of these constraints indicates the rate of increase in the objective function from relaxing the constraint. For the case of local threshold policies, if demand at location i exceeds threshold S_i, then this dual value is precisely the gradient on the total profit of the LP with respect to threshold S_i. We average these gradient estimates over the U samples to get an unbiased estimate of the gradient each time we update the threshold values.

Lemma 3. *The expected value of* $p - g(i)\mathbf{1}[S_i < \mathcal{D}_i^O]$ *where* $g(i)$ *is the optimal dual variable corresponding to constraint* i, $C_i + \sum_{j=1}^{n} F_{ji} \ge A_i,\ \forall i \in [n]$, *from linear program* (3) *is the negative partial derivative of the expected value of the objective of the multiple-store model with respect to* S_i.

Proof. Let $g(i)$ be the optimal dual variable corresponding to constraint i. Let $\Pi^*(\mathcal{D}^P, \mathcal{D}^O) = \sum_{i=1}^{n} (p \min(\mathcal{D}_i^P, \mathcal{I}_i) - (c_i + M)C_i^* + MA_i + \sum_{j=1}^{n}(p - s_{ji} - M)F_{ji}^*)$, the optimal objective value of linear program (3). It is clear by inspection that $\frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial S_i} = -g(i)\mathbf{1}[S_i < \mathcal{D}_i^O]$. If $S_i < \mathcal{D}_i^O$, then $A_i = S_i$, and so, an infinitesimal change to S_i will cause A_i to change an equal amount in the same direction. An immediate consequence of linear programming duality theory is that the optimal dual value of a constraint is equal to marginal change to the optimal objective value from relaxing this constraint. More precisely, $\frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial S_i} = \frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial A_i}\frac{\partial A_i}{\partial S_i} = \frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial A_i} = p - g(i)$. In the other case, $S_i \ge \mathcal{D}_i^O$, and so, an infinitesimal change in S_i does not change A_i. $\frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial S_i} = \frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial A_i}\frac{\partial A_i}{\partial S_i} = \frac{\partial \Pi^*(\mathcal{D}^P, \mathcal{D}^O)}{\partial A_i} 0 = 0$. Therefore, taking expectations, we conclude $\frac{\partial E[\Pi^*(\mathcal{D}^P, \mathcal{D}^O)]}{\partial S_i} = p - E[g(i)\mathbf{1}[S_i < \mathcal{D}_i^O]]$. □

Theorem 5. *From theorem* 5.1 *of Hazan et al.* (2015), *using* $p - \frac{g}{|g|}$ *as the normalized gradient, our algorithm obtains an* ϵ*-optimal local threshold policy with* $\text{poly}(\frac{1}{\epsilon})$ *total samples of linear program* (3).

Proof. The following follows directly from the convergence Theorem 4 and the properties of the optimal dual variables in linear program (3) established in Lemma 3. □

4.1.3. Global Threshold Derivative Estimates. We use the same dual values used to estimate derivatives with respect to local threshold parameters, those corresponding to constraints $C_i + \sum_{j=1}^{n} F_{ji} \geq A_i, \ \forall i \in [n]$ from linear program (3), to estimate derivatives of the objective function with respect to a global threshold parameter. The dual value from one of these constraints indicates the rate of increase in the objective function from relaxing the constraint. For a global threshold policy, if total demand exceeds threshold S, then the derivative of the total profit of the LP with respect to threshold S is the sum of these dual values, weighted by the probabilities that the first rejected order is from each store.

Lemma 4. *Suppose* λ_i *is the mean online demand at location i, and* f_i^O*, the distribution from which random variable* $\mathcal{D}_i^O$ *is drawn is a Poisson distribution. Also suppose that g(i) is the dual value of constraint* $C_i + \sum_{j=1}^{n} F_{ji} \geq A_i$ *after solving the maximization LP* (3) *for a demand sample, and G(S) is the expected value of the omnichannel fulfillment model with global threshold S. Then,* $\frac{dG(S)}{dS} = E\left[p - \sum_{i=1}^{n} \frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j}(-g(i))\mathbf{1}\left[S < \sum_{k=1}^{n} \mathcal{D}_i^O\right]\right]$.

Proof. Among instances when the first rejected order is at location i, $-g(i)$ is the unbiased derivative estimate, so in general, the unbiased derivative estimate is the weighted sum of dual values across all locations, weighted by arrival probability. □

Note that the Poisson arrival assumption is important here. This ensures that the probability of where the first order arrives after the threshold is filled can be computed independently of S. In general, this works for arrival distributions that can be proportionally thinned over time. We average these gradient estimates over the U samples to get an unbiased estimate of the gradient each time we update the threshold values.

Theorem 6. *From theorem* 5.1 *of Hazan et al.* (2015), *using* $p - \sum_{i=1}^{n} \frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j}(-g(i))\mathbf{1}\left[S < \sum_{k=1}^{n} \mathcal{D}_i^O\right]$ *as the normalized gradient, our algorithm obtains an* ϵ*-optimal global threshold policy with* $\text{poly}(\frac{1}{\epsilon})$ *total samples of linear program* (3).

Proof. This result follows directly from the convergence Theorem 4 and the properties of the optimal dual variables in linear program (3) established in Lemma 4. □

4.2. Hybrid Policy

In addition to the thresholds S_i at each location i, we also use a global threshold S, which affects the final number of accepted orders as follows. Each location i tentatively accepts $\tilde{A}_i = \min\{\mathcal{D}_i^O, S_i\}$ online orders (no different from the local threshold policy). If $\sum_i \tilde{A}_i \leq S$, then the final accepted amount is $A_i = \tilde{A}_i$. Otherwise, we shrink each $\tilde{A}_i$ proportionally until their sum decreases to S as follows: $A_i = \frac{S}{\sum_i \tilde{A}_i}\tilde{A}_i$. In our numerical study of the local and global threshold policies, we found that in nearly all cases the hybrid policy closely tracked either local thresholds or global thresholds, whichever method performs better on the instance.

5. Computational Results on One-Period Models

We assess the empirical performance of these policies on full-size and two-store problem instances. This will give us insight into the strengths of each policy class while also verifying that our IPA method is of practical use.

5.1. Complete Network Results

We use demand distributions that are estimated from sales and inventory data of an upscale North American retailer. We use the demand data across the full retail network from the top 20 bestselling items at this retailer to generate a realistic instances. A typical instance will have 30–40 store locations. The cancellation cost parameter is set to two times the price parameter, and ship costs are proportional to distance. Inventories are set to two units at each retail location to generate instances where careful supervision of online fulfillment is necessary. For each of these test instances, we compare local threshold and global threshold policies with siloed fulfillment and reactive fulfillment policies.

Definition 2. The siloed fulfillment policy treats each store location as a separate retail network and computes the optimal global threshold policy for each store as its own instance.

Siloed fulfillment policies might be used in practice if a retailer is not aware or sophisticated enough to implement a coordinated full-network ship from store program.

Definition 3. The reactive fulfillment policy is the local threshold policy that uses the thresholds from the siloed fulfillment policy as its threshold parameters.

Reactive fulfillment policies use the same set of thresholds for the first-stage problem as those computed

Table 1. Average Fulfillment Costs Across 20 Full-Network Instances

	Siloed fulfillment	Reactive fulfillment	Global threshold	Local threshold
Average cost	471	369	116	104
Saving, %	—	21.5	75.3	77.9

Table 2. Overall Average Costs of Each Inventory Balance Condition for Inventory Balance Experiments

	Siloed fulfill	Reactive fulfill	Global	Local
25%:75%	157.9	145.7	128.0	117.2
50%:50%	64.7	53.2	51.1	46.4
75%:25%	53.1	46.2	40.6	35.4

by the siloed fulfillment policy, but the retailer is still able to execute long-distance shipments when solving the second-stage problem.

The four fulfillment algorithms are tested in 100 trials for each of the 20 items to generate the results described in this section and Table 1. We find that across our test instances, local thresholds and global thresholds provide significant (78% and 75%, respectively) improvement over the siloed fulfillment and reactive fulfillment policies.

5.2. Insights from Two-Store Instances

We are specifically interested in four questions.

1. What is the effect of balanced and imbalanced inventory?
2. How does the magnitude of in-store demand affect performance?
3. Does the relative performance of policies vary with cancellation costs?
4. What conditions result in good performance of global thresholds?

To answer each of these questions, we run our policies across several instances that vary in a deliberate way across a small number of specific parameters. To assess the effect of inventory balance, we vary how evenly inventory is distributed between stores, whether this inventory is aligned with demand, and whether the total amount of inventory available modulates with this effect. We investigate the effect of in-store demand magnitude by testing our algorithms on four different in-store demand levels, each tested on instances with four different inventory levels. Lastly, we test our algorithms on an instance where cancellation cost c is varied from 50% to 400% of the unnecessary rejection penalty p to understand the impact of the ratio of cancellation cost to rejection penalty on the relative performance of our methods. We conduct each of these two-store experiments on demand distributions fit to 10 real-life items.

5.2.1. Inventory Balance. For each item, we fit Poisson demand distributions for its two top-selling store locations (with respect to online demand). We allow inventory to vary at three levels ranging from 50% of mean total demand to 150% of mean total demand. We also let inventory balance vary from 25% to 75% of total inventory in the first location, across three conditions. This results in nine trials for each item evaluated. Our primary finding is that local thresholds provide the greatest improvement over global thresholds when inventory is not aligned with demand. We also observe that this effect is magnified by low inventory levels. As the total amount of starting inventory increases, the performances of the two methods become very similar.

We report overall results in Table 2, averaged by inventory balance condition. Each inventory balance condition is evaluated at three total inventory levels per item, and 10 items are tested. Every individual instance is evaluated by taking the average cost of each policy over 10,000 samples of demand. We call the inventory balance conditions "25%:75%," "50%:50%," and "75%:25%." In these conditions, the first percentage refers to the percent of total inventory located at the location with the higher online demand rate, and the second percentage indicates the percent of total inventory located at the location with the lower online demand rate.

Across all instances, local thresholds slightly outperform global thresholds, and both threshold policies substantially outperform the two benchmark policies, siloed fulfillment and reactive fulfillment. The gap between our IPA threshold policies (local thresholds and global thresholds) and these benchmarks grows in absolute terms yet shrinks in percentage terms as inventory is most out of balance with online demand.

5.2.2. Magnitude of In-Store Demand. To answer this question, we set inventory equal at each location, but we tested four levels of inventory at each location: 5, 10, 15, and 20. For each inventory level, we test three Poisson rate parameters of in-store demand: 25%, 50%, and 75% of inventory. This results in 12 trials for each item. We observe that local thresholds outperform global thresholds across all scenarios, but the performance of local threshold policies is more sensitive to increases in in-store demand. The summary results are in Table 3.

5.2.3. Impact of Cancellation Costs. In this experiment, we test three inventory levels: 5, 10, and 15 units at each store location. For each of these inventory levels, we compare our policies at the following cancellation costs: 20, 40, 60, and 80. This results in 12 total

Table 3. Overall Average Costs of Each In-Store Demand Condition for Demand Magnitude Experiments

Demand, %	Siloed fulfill	Reactive fulfill	Global	Local
25	54.4	47.0	50.8	41.7
50	58.8	47.1	51.2	43.9
75	63.1	47.9	52.5	46.9

Table 5. Average Cost of Each Covariance Condition

In-store variance	Online ρ	Siloed fulfill	Reactive fulfill	Global	Local
1.5	−0.7	41.2	26.6	16.5	22.8
6	−0.7	61.2	36.8	30.7	34.2
10.5	−0.7	74.1	42.4	37.2	40.4
1.5	0	32.1	20.0	15.5	19.9
6	0	56.0	31.8	28.9	31.8
10.5	0	67.9	38.5	37.1	38.5
1.5	0.7	25.3	15.9	15.5	15.9
6	0.7	51.1	31.2	29.1	29.6
10.5	0.7	65.6	38.2	36.4	36.7

trials for each item. The price of the item is set to 20 for all trials. The overall results from these experiments are presented in Table 4.

5.2.4. Global Thresholds Performance. In this section, we explore potentially artificial scenarios that result in global threshold policies outperforming local threshold policies to see which extreme situations could lead to this. We consider a set of two-store instances where inventory is fixed at 20 at both locations. $c = p = 20$, and the shipping cost between the two stores is 0.5. Demand distributions are multivariate normal, and we will vary the covariance matrix across several conditions. The in-store demand distribution has mean demand 15 at each location, and values along the diagonal of the covariance matrix are 1.5, 6, and 10.5 across three conditions tested. The covariance of in-store demand between the two locations is zero. The online demand distribution has mean demand 5 at each location, and the values along the diagonal of the covariance matrix are five. The covariance between the two stores is set such that the correlation coefficients of online demand are –0.7, 0, and 0.7 across the three conditions tested. Results are in Table 5.

Several factors influence the results of these trials in favor of the global threshold policies. For one, we set shipping costs to be relatively low. If shipping costs are zero, then the model becomes equivalent to a single-location instance, and local threshold will no longer outperform global threshold; however, this alone is not always enough to give global threshold a distinct advantage. In particular, for global threshold to have an advantage, online demand rates should be in a narrow range where there is enough demand to differentiate the policies but not enough demand for a local threshold policy to always accept up to its threshold value at all locations. The average costs at each demand rate are plotted in Figure 2, confirming that the advantage global threshold policies have are in a narrow range.

Negatively correlated online demand hurts local threshold policies by cancellations, whereas global threshold policies are largely unaffected especially at low shipping costs. Our results are consistent with this observation, where global threshold policies have 19% lower cost than local threshold policies when online demand has a correlation coefficient of –0.7, compared with cost decreases of 7% and 4% for correlation coefficients 0 and 0.7, respectively. Similarly, global threshold policies have 13% lower cost than local threshold policies when in-store variance is 1.5, compared with cost decreases of 10% and 1% for variances of 6 and 10.5, respectively.

5.3. Hybrid Thresholds

We compare the performance of the hybrid threshold with local and global thresholds in a two-location instance in Table 6 across a range of shipping costs s. The global threshold policy outperforms the local counterpart for small s but becomes worse because of increased shipping costs incurred in reconciliation.

Table 4. Overall Average Costs of Each Cancellation Penalty Condition for Cancellation Cost Magnitude Experiments

Cancel penalty	Siloed fulfill	Reactive fulfill	Global	Local
20	60.6	50.9	52.5	45.7
40	71.3	60.1	58.8	53.1
60	77.4	66.4	62.7	57.1
80	81.8	71.1	65.8	60.0

Figure 2. (Color online) Average Objective Value as the Mean Online Demand Parameter Is Shifted

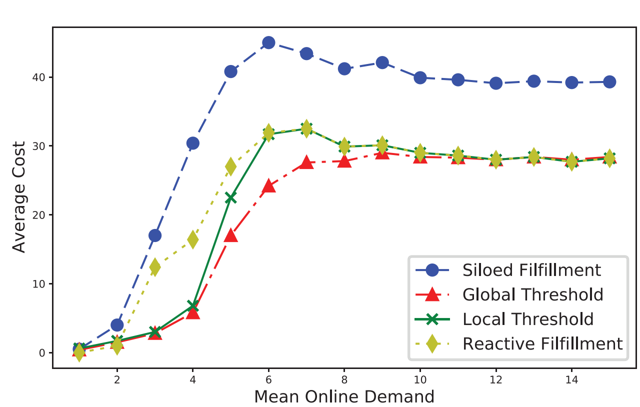

Table 6. Impact of Increasing Shipping Costs s When $c = 40, p = 20, I = 20, D^O = D^P = 10$

s	Reactive	Local	Global	Hybrid	s	Reactive	Local	Global	Hybrid
5					20				
Revenue (physical)	400.33	399.8	399.71	400.27	Revenue (physical)	399.90	399.83	399.81	399.92
Revenue (online)	329.36	322.62	325.50	324.97	Revenue (online)	329.56	315.13	325.55	320.89
Revenue (total)	729.69	722.42	725.21	725.24	Revenue (total)	729.46	714.96	725.36	720.81
Shipping cost	3.76	3.61	5.05	4.48	Shipping cost	15.08	13.15	20.33	14.9
Cancellation cost	40.89	33.06	32.05	31.86	Cancellation cost	40.55	26.41	31.96	29.22
Opportunity cost	20.14	27.08	24.40	24.77	Opportunity cost	20.1	34.72	24.31	28.92
Profit	685.04	685.75	688.11	688.90	Profit	673.83	675.4	673.07	676.69
Total cost	64.79	63.74	61.50	61.11	Total cost	75.73	74.28	76.6	73.04
10					40				
Revenue (physical)	400.02	399.89	400.06	399.88	Revenue (physical)	399.85	399.89	399.88	399.84
Revenue (online)	329.63	322.59	325.48	323.72	Revenue (online)	329.61	315.07	315.44	315.08
Revenue (total)	729.65	722.48	725.54	723.60	Revenue (total)	729.46	714.96	715.32	714.92
Shipping cost	7.55	7.21	10.15	8.26	Shipping cost	30.21	26.28	37.54	26.31
Cancellation cost	40.66	32.97	32.16	30.85	Cancellation cost	40.6	26.42	24.1	26.32
Opportunity cost	20.09	27.12	24.20	26.02	Opportunity cost	20.13	34.69	34.41	34.69
Profit	681.44	682.3	683.23	684.49	Profit	658.65	662.26	653.68	662.29
Total cost	68.29	67.30	66.51	65.14	Total cost	90.93	87.39	96.05	87.33

Additional experiments confirm that the hybrid improves upon the better of local and global only marginally in a wide variety of settings.

6. Multiperiod Models

In this section, we extend our approach to the N-location, T-period model. The details of the one-location, T-period model are as follows. (a) In any given period, we accept a maximum number of online orders. (b) Physical demand in this time period is fulfilled first. (c) Remaining inventory, if any, is used to satisfy online orders. (d) Remaining online orders, if any, are cancelled. All future online orders and physical demands are lost as there is no inventory left at the location. Note that our model prioritizes current period online orders over the next period's physical demands.

The multilocation, multiperiod model we consider here, likewise, treats online orders in a given period with priority over future period physical demand at the same location. That is, we allow transshipments at the end of each period, but they are not obligated to satisfy demands at other locations. Thus, because transshipments are costly, online orders at a location may be cancelled, although inventory exists in another location, as it is profitable to keep that inventory in its location to satisfy local demand later. On the other hand, if another location has sufficiently high inventory, some transshipment may indeed take place at the end of each period to (partially or completely) satisfy the online demand elsewhere, after its own physical demand and online demands are satisfied. That is, at the end of each period, online orders at a location are satisfied as best as possible with local inventory, may be satisfied completely by transshipped inventory, or cancelled. Online orders are *not* carried over from period to another. Because transshipment is allowed in every period, there is no reason to "pretransship" to balance inventories across locations.

A T-period problem begins with the first period and ends after period T. All references to variables in different periods use the period number as the superscript and retain the subscript to denote the location (e.g., the inventories at location j at the end of period t are denoted I_j^t with I_j^0 being the initial inventory, and demands at location j over time are denoted by $\mathcal{D}_j^{O,t}$ (online) and $\mathcal{D}_j^{P,t}$ (physical)). The online order acceptance thresholds in location j over time are $S_j^1, \ldots, S_j^T$. We study only (time-varying) local threshold policies in this section.

We employ a profit-maximization formulation and focus on local threshold policies that were effective in the single-period case for realistic data. The global threshold and hybrid variants for this case can be addressed in a similar vein.

6.1. Single-Location, Multiperiod Model

We first characterize the optimal period 1 threshold for the T-period model similar to Theorem 1. Then, we introduce some heuristics and compare their performance. We suppress the subscript corresponding to the single location.

6.1.1. Characterization of the Optimal Threshold. The additional complexity in analyzing the optimal period 1 threshold for the multiperiod model is because of the optimal threshold S^2 for the second period depending on the demand realization in period 1.

Define $\Pi(I, t)$ to be the maximum expected profit for the single-location, t-period model with initial inventory I. Let $I^1 = (I^0 - A^1 - \mathcal{D}^{P,1})^+$ be the (random) remaining inventory at the end of period 1. As in Theorem 1, the following theorem involves finding the condition under which the derivative of the expected profit is zero. A key quantity here is $\frac{\partial}{\partial I}\Pi(I, T-1)$, which measures how the optimal expected profit of the $(T-1)$-period problem would change, given a unit change in its initial inventory I.

Theorem 7 (Optimality Condition; Proof in Online Appendix D). *Let I^0 be the initial inventory. Then, the optimal period 1 threshold S^1 for the single-location, T-period model satisfies*

$$S^1 = \max\left\{0, I^0 - F_{P,1}^{-1}\left(\frac{c}{c+p-\tilde{p}(S^1)}\right)\right\},$$

where $\tilde{p} = \tilde{p}(S^1) = E\left[\frac{\partial}{\partial I}\Pi(I^1, T-1) \mid \mathcal{D}^{O,1} > S^1\right]$ and $F_{P,1}$ is the cumulative density function of $\mathcal{D}^{P,1}$.

Note that for $T = 2$, this differs from the optimal single-period threshold (Theorem 1) only in having an extra $\tilde{p}$ term in the denominator. Further, observe that $\tilde{p}(S) \in [0, p]$ for any S because $\frac{\partial}{\partial I}\Pi \in [0, p]$. Hence, $1 \geq \frac{c}{c+p-\tilde{p}} \geq \frac{c}{c+p}$, and we obtain the following monotonicity result.

Corollary 1. *Consider a single-location, single-period instance X and a single-location, T-period instance Y with arbitrary $T \geq 1$. Suppose X and Y have the same c, p, I^0 and $F_{P,1}$. Then, the optimal period 1 thresholds $S^*(X), S^*(Y)$ satisfy $S^*(X) \geq S^*(Y)$.*

We note that when the online demand is infinite, there exists an optimal policy with a first-stage optimal threshold equal to zero as long as $T > 1$ via a simple transfer argument that defers the thresholds S in all but the last period to the last. Our limited simulation experiments confirm that even when online demand is not infinite, the first-stage optimal thresholds are nonincreasing with T, the number of periods of independently and identically distributed demand. However, these effects may arise because there is no (holding) cost for deferring fulfillment and also because depleting inventory using online orders before fulfilling physical demands in future periods incurs no penalty in this model. The study of the more realistic multiperiod models with different profit margins for online and offline orders as well as holding costs between periods is left for future work.

6.1.2. Algorithms. We first present a dynamic programming (DP) approach for $n = 1$ and arbitrary T. Our DP estimates $\Pi(I, t)$ for every combination $I \leq I_{max}$ and $t \leq T$ as follows.

Define $V(S, D, I, t)$ to be the maximum expected net profit (i.e., revenue minus cancellation and shipping costs) if the period 1 threshold is S and demand vector (physical and online) is D, with inventory I and t periods. By definition of Π, we have $\Pi(I, t) = \max_S E_D[V(S, D, I, t)]$, and our goal is to find $\Pi(I^0, T)$.

Proposition 5 (Bellman Equation). *Denote the period 1 demand vector $D = (D^O, D^P)$. Let $I^1 = I^1(D, S)$ be the remaining inventory after period 1 and $\pi(I, D, S)$ be the one-period profit. Then, $V(S, D, I, t) = \Pi(I^1, t-1) + \pi(I, D, S)$, and hence, $\Pi(I, t) = \max_S E_D[\Pi(I^1, t-1) + \pi(I, D, S)]$.*

6.1.2.1. Heuristics. Because the run time of the DP scales linearly in T, we propose two efficient heuristics. First, consider a myopic approach; at each period, view this period as the last period, and set the threshold using the closed-form formula in Theorem 1. However, because such a heuristic fails to consider future demands, it is inclined to set high thresholds and take unnecessary cancellation risks. To circumvent this issue, we propose another heuristic (the *look-ahead* policy) that pretends that the remaining t-period situation can be viewed as a two-period model, with the second period being an aggregation of the remaining $(t-1)$ periods. In our numerical validation in Figure 3, our look-ahead policy behaves similarly to the optimal policy in choosing thresholds at lower initial inventories but yields close to optimal profit overall. In contrast, the myopic heuristic may be significantly worse (in the worst case, by 15%).

6.2. Two-Location, Two-Period Model

The two-location, two-period model becomes significantly more difficult because of the optimal shipping decisions. In the one-period, N-location model, if there is any backlog, an optimal seller would ship the available inventory from surplus locations to avoid as many cancellations as possible, as long as the margin after shipping is higher than cancellation fee (which is usually true in practice and assumed throughout). Hence, the shipping problem reduces to a network flow problem.

In the two-period case, however, by shipping from location 1 to 2, at the end of period 1, we lose some "potential" of earning at location 1 in period 2. In particular, it is better not to ship when the loss of potential at location 1 outweighs the immediate benefit of avoiding cancellations at 2. This intertwining nature makes the general N, T version significantly harder.

Define the *value function* $V(\mathbf{S}, \mathbf{D}, \mathbf{I}, t)$ to be the maximum expected profit if in period 1, the threshold vector is $\mathbf{S}$ and the demand vector (physical and online) is $\mathbf{D}$, with inventory $\mathbf{I}$ and t periods to go. Fix $t \in \{1, 2\}$ and inventory $\mathbf{I}$, and write $V(\mathbf{S}, \mathbf{D}) = V(\mathbf{S}, \mathbf{D}, \mathbf{I}, t)$ for simplicity. Define the *potential* $\Pi(\mathbf{I}, t) = \max_{\mathbf{S}} E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D}, \mathbf{I}, t)]$. Our objective is to maximize $E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D}, \mathbf{I}, t)]$. Again, our approach relies on a Bellman-type equation (which can be generalized directly for any N, T).

Figure 3. (Color online) Comparison Between Policies for $N = 1, T = 3$ when $c = 20, p = 10, I = 30, D_P = 5, D_O = 10$

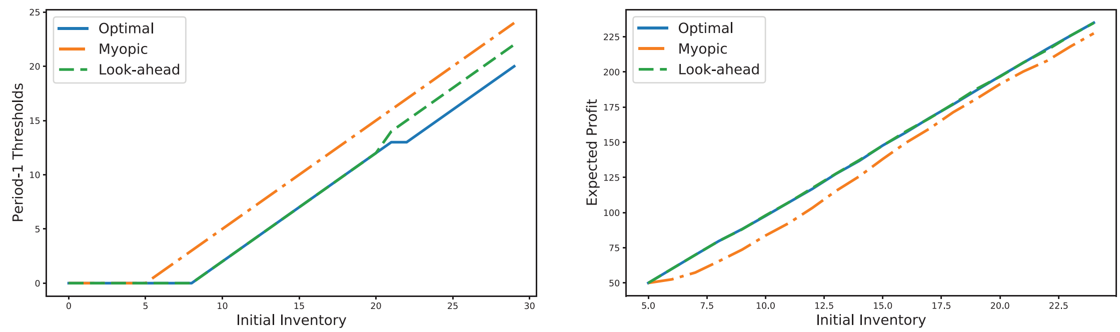

Note. The myopic policy selects the higher thresholds and incurs more cancellation risk; hence, it yields lower profit.

Proposition 6 (General Bellman Equation). *Let* $\mathbf{D}$ *be the period 1 demands at the two locations. Let* $\mathbf{I}^1(\mathbf{S}, \mathbf{D}, \mathbf{I}, x)$ *be the inventory at the end of period 1, after implementing a shipping decision* x*, with initial inventory* $\mathbf{I}$ *and period 1 demand* $\mathbf{D}$ *and thresholds* $\mathbf{S}$*. Let* $\pi(\mathbf{S}, \mathbf{D}, \mathbf{I}, x)$ *be the one-period profit with shipping decision* x *excluding the shipping costs. Then,*

$$V(\mathbf{S}, \mathbf{D}, \mathbf{I}, 2) = \max_x \Big\{\pi(\mathbf{S}, \mathbf{D}, \mathbf{I}, x) - \mathbf{s}^T \boldsymbol{x} + \Pi(\mathbf{I}^1(\mathbf{S}, \mathbf{D}, \mathbf{I}, x), 1)\Big\}, \quad (4)$$

where the maximum is taken over feasible shipping decisions x*. Further,* $\Pi(\mathbf{I}, 2) = \max_{\mathbf{S}} E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D}, \mathbf{I}, 2)]$.

We describe how to find an optimal shipping decision x^* (for $N = T = 2$). Let $\mathbf{A} = (A_1, A_2)$ and $\mathbf{R} = (R_1, R_2)$ be the number of accepted orders and the remaining inventory after fulfilling *physical* orders in period 1 at the two locations, respectively. The nontrivial case is when $R_i > A_i$ at one location i and $R_j < A_j$ at the other location j, where in addition to determining whether to ship items, we must also determine how many to ship. W.l.o.g suppose $R_2 < A_2$ and $R_1 > A_1$. The optimal shipping decision strikes the best trade-off between the drop of potential at location 1 and avoiding cancellations at 2. Formally, let x^* be the units shipped from location 1 to 2 in an optimal policy. Then, $x^* = \arg\max\{(p + c - s)x + \Pi(I_1^1 = R_1 - A_1 - x, I_2^1 = 0, t = 1) : 0 \le x \le \min\{R_1 - A_1, A_2 - R_2\}\}$.

Rather than enumerate over all $\mathbf{S}$, a faster approach to find the optimal $\mathbf{S}$ is by IPA. Our IPA-based procedure views the expected profit as a function of the period 1 threshold vector $\mathbf{S} = (S_1, S_2)$ and returns $\mathbf{S}$ where the gradient (almost) vanishes. It is not hard to compute an explicit formula for $\nabla_{\mathbf{S}} V(\mathbf{S}, \mathbf{D}, \mathbf{I}, t)$. Figure 4 confirms that the thresholds with almost vanishing gradients computed by our IPA-based heuristic almost match that of a brute-force approach.

Figure 4. (Color online) The Gradient Norms of the Expected Value Function, $E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D})]$, for Different Period 1 Thresholds $\mathbf{S} = (S_1, S_2)$ when $p = c = 10, s = 5, I = 15$, and $\mathcal{D}_t^{P,i} = 5, \mathcal{D}_t^{O,i} = 10$ for $i = 1, 2, t = 1, 2$

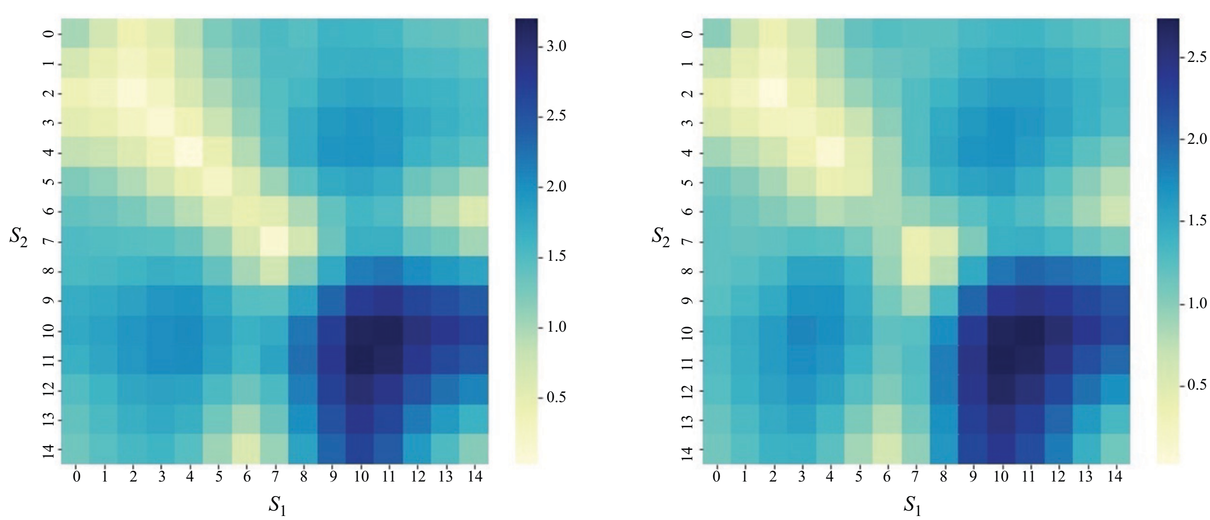

Note. The left and right panels are returned by the IPA-based heuristic and brute-force approaches, respectively.

6.2.1. Models for General N, T. As illustrated, the key difficulties in extending our methods for the case of multiperiod, multilocation fulfillment is the requirement to encode the future profit potential of holding back excess inventory rather than ship it to fulfill backlogged orders elsewhere in the network. For the two-location, T-period case, we can adapt the look-ahead heuristic of one location by synthetically merging the remaining $(t-1)$ periods into one and using the IPA-based approach that we proposed for $N = T = 2$. At sufficiently high inventories, we expect myopic policies to be effective in practice for the general multiperiod problem. At lower inventories, the look-ahead heuristic may help reduce the loss of profits of the myopic policy. We leave the detailed study of the general N, T model for future work.

7. Conclusion

In this paper, we introduce a new stochastic two-stage model for omnichannel fulfillment. We incorporate new risks that occur when fulfillment operations are combined for in-store and online demand. We obtain a closed-form optimal solution for a one-period, single-location model. We then study the one-period, multiple-location setting, with local threshold and global threshold policies. We also present a sampling-based IPA algorithm to optimize threshold policies within each of these policy classes. Next, we evaluate our methods in a variety of test instances based on North American retailers and find that our IPA-optimized local threshold policies consistently outperform global threshold policies and other benchmark policies under realistic conditions. In a synthetic two-location setting, we explored several factors and discuss how changes along these dimensions affect the performance of our policies. We extend our study to multiple periods and find that at high inventories, myopic policies are effective, whereas at lower inventories, a look-ahead heuristic may help. Extending our work to the full generality of multiperiod models and investigating alternate multiperiod formulations are important directions for future work.

Endnotes

[1] See, for example, https://forums.bestbuy.com/t5/BestBuy-com-Knowledge-Base/Why-Was-My-Order-Cancelled/ta-p/956598.

[2] In Section 4.2, we define a hybrid policy that uses both local thresholds and an additional global threshold to moderate the local acceptances proportionally to capture the best of both controls.

[3] We note that for the retailer to not cancel any order it can possibly fill from any other location, it suffices for the maximum shipping cost to be at most the sum of the cancel cost and the opportunity cost of a lost sale, which translates to the weaker condition $\max_{i,j\in[n]} s_{ij} < p + c$.

References

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

Alishah EJ, Moinzadeh K, Zhou Y-P (2015) Inventory fulfillment strategies for an omni-channel retailer. Preprint, submitted September 17, http://dx.doi.org/10.2139/ssrn.2659671.

Bayram A, Cesaret B (2021) Order fulfillment policies for ship-from-store implementation in omni-channel retailing. *Eur. J. Oper. Res.* 294(3):987–1002.

Cattani K, Gilland W, Heese HS, Swaminathan J (2006) Boiling frogs: Pricing strategies for a manufacturer adding a direct channel that competes with the traditional channel. *Production Oper. Management* 15(1):40–56.

DeValve L, Wei Y, Di Wu RY (2021) Understanding the value of fulfillment flexibility in an online retailing environment. *Manufacturing Service Oper. Management*, ePub ahead of print July 12, https://doi.org/10.1287/msom.2021.0981.

Gallino S, Moreno A, eds. (2019) *Operations in an Omnichannel World* (Springer, Cham, Switzerland).

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

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

Glasserman P, Tayur S (1995) Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Sci.* 41(2):263–281.

Govindarajan A, Sinha A, Uichanco J (2017) Inventory optimization for fulfillment integration in omnichannel retailing. Preprint, submitted March 1, http://dx.doi.org/10.2139/ssrn.2924850.

Harsha P, Subramanian S, Uichanco J (2019) Dynamic pricing of omnichannel inventories. *Manufacturing Service Oper. Management* 21(1):47–65.

Hazan E, Levy K, Shalev-Shwartz S (2015) Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, eds. *Beyond Convexity: Stochastic Quasi-Convex Optimization*, Advances in Neural Information Processing Systems, vol. 28 (Curran Associates, Inc., Red Hook, NY), 1594–1602.

Herer YT, Rashit A (1999) Lateral stock transshipments in a two-location inventory system with fixed and joint replenishment costs. *Naval Res. Logist.* 46(5):525–547.

Herer YT, Tzur M, Yücesan E (2006) The multilocation transshipment problem. *IIE Trans.* 38(3):185–200.

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

Karaesmen I, Van Ryzin G (2004) Overbooking with substitutable inventory classes. *Oper. Res.* 52(1):83–104.