

# Effective Online Order Acceptance Policies for Omni-Channel Fulfillment

Su Jia, Jeremy Karp, R. Ravi, Sridhar Tayur  
Tepper School of Business, Carnegie Mellon University

**Problem Definition:** Omni-channel retailing has led to the use of traditional stores as fulfillment centers for online orders. Omni-channel fulfillment problems have two components: (1) accepting a certain number of on-line orders prior to seeing store demands, and (2) satisfying (or filling) some of these accepted on-line 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 due to 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 multi-period settings.

**Academic/Practical Relevance:** Despite the growing importance of omni-channel 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 once the uncertain quantities of in-store purchases are revealed. We analyze two classes of threshold policies that accept online orders as long as the inventories are above a global threshold, or a local threshold per region.

**Results:** 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-store 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 give an effective way to manage fulfillment costs for online orders, demonstrating a significant reduction compared to policies that treat each store separately, reflecting the significant advantage of incorporating shipping in computing thresholds.

*Key words:* omni-channel retail, stochastic programming, infinitesimal perturbation analysis

---

## 1. Introduction

Omni-channel 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 the left-over inventory after filling in-store demands is zero) and may have to be cancelled by the retailer<sup>1</sup>. We study a new set of research questions related to acceptance and fulfillment of these online orders in omni-channel retail operations, taking into consideration shipping costs when they are filled and cancellation costs when they are not.

### 1.1. Omni-Channel Fulfillment Model

Fixed exogenous inventory  $I_i$  is available at each location  $i$  in a system of  $n$  stores. At each location  $i$ , there are two streams of demand, on-line ( $D_i^O$ ) and physical ( $D_i^P$ ), 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 store 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

<sup>1</sup> See e.g., <https://forums.bestbuy.com/t5/BestBuy-com-Knowledge-Base/Why-Was-My-Order-Cancelled/ta-p/956598>.

total number of online orders accepted but less fine-tuned control than with Local Thresholds over which orders are accepted.<sup>2</sup>

A Local Threshold policy  $[S_1, \dots, S_n]$  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.

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:

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

Consider first the objective function:

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

The expression  $\sum_{i=1}^n (I_i - \min(I_i, D_i^P) - \sum_{j=1}^n F_{ij})$  is the amount of remaining inventory after all orders have been fulfilled or cancelled.  $I_i$  is the starting inventory at location  $i$ ,  $D_i^P$  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  $\sum_{i=1}^n (D_i^O - A_i^O)$  is the number of online orders that were rejected.  $D_i^O$  is the amount of online demand at location  $i$ , and  $A_i^O$  is the number of online orders accepted from location  $i$  in the first stage. Consequently,  $\min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P) - \sum_{j=1}^n F_{ij}), \sum_{i=1}^n (D_i^O - A_i^O))$  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

<sup>2</sup> 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.

the sale price of the item. The selling price is an upper bound on the opportunity cost, since 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 they could possibly fill, i.e,  $c > \max_i f s_{ji} g$ .<sup>3</sup> 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  $\min(D_i^P, I_i) + R_i + \sum_{j=1}^n F_{ij} = I_i$  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  $C_i + \sum_{j=1}^n F_{ji} = A_i^O$  reflect that all accepted orders must be either cancelled or fulfilled. This is the omni-channel fulfillment problem we study in the rest of this paper.

## 1.2. Summary of Contributions

1. We formulate an analytical model for omni-channel fulfillment that incorporates uncertainty due to inventory pooling across sales channels as a multi-location, 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 store case and Section 3 for multiple stores).
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

<sup>3</sup>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  $\max_{i,j \in [n]} s_{ij} < p + c$

retail analytics firm, we use retail industry data to generate realistic problem instances. We conduct a series of experiments on two-store instances to demonstrate how certain instance attributes lead to strong performance of one class of threshold attributes relative to the other.

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 to other simple baseline policies currently employed by the retailer.

7. We formulate an extension to multiple periods. We develop exact methods for single store multi-period, and the two-store 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 Gallino and Moreno (2019) provides a comprehensive survey of literature related to omni-channel operations, with the chapter by Jasin, Sinha and Uichanco therein being the closest to this paper. Our work is related to the following streams of literature: (a) overbooking (b) fulfillment of on-line orders from stores and (c) multi-location transshipment. We briefly review here the closest papers to our work.

Harsha et al. (2016) 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 omni-channel retailing, such as the costs and benefits of “buy online and pick up in store” policies [Gao and Su (2016a)], information sharing [Gao and Su (2016b)], inventory optimization [Govindarajan et al. (2017)], and multi-channel price optimization [Harsha et al. (2016), Cattani et al. (2006)] have been studied by the Operations Management community, but ours is the first attempt to formulate and study stochastic models of cancellations caused by omni-channel 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 on-line orders that have been accepted in that period. Note that although the store inventory is shared across the on-line 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 on-line orders who risk not being fulfilled, even considering transshipments from other locations, is reminiscent of the multi-class overbooking

model with substitutions of Karaesmen and van Ryzin 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 on-line orders to fulfill. The models of Alishah et al. (2015) consider the omni-channel setting of fulfilling an order from a combination of a single omni-channel 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.

The multi-location transshipment problem has been previously formulated and studied in Herer et al. (2006), Herer and Rashit (1999). This work presents a stochastic multi-period 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, though the methodology of our work and the problem context are quite different from the setting of Herer et al. (2006), Herer and Rashit (1999).

Our models of omni-channel fulfillment thus incorporate some of the complexities considered in the multi-location transshipment problem, and multi-class overbooking problem, as they include order cancellations as a major component, that are the result of accepting on-line orders, which themselves should take into account the possibility of transshipments.

## 2. Single-Store 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. The second stage problem is straightforward:  $\min(I - \min(I, D^P), \min(D^O, S))$  accepted online orders are filled and the rest are cancelled. Using this observation, we can re-write the optimization problem (1) in a simpler form. The objective function simplifies to minimizing  $\min(I - \min(I, D^P) - F, D^O - \min(D^O, S)) + cC$ .

**PROPOSITION 1.** *(Proof in Appendix A) For a single-store instance of the Omni-Channel Fulfillment Model  $F = \min(D^O, S, I - \min(I, D^P))$  is satisfied in any optimal solution.*




Figure 1 Examples of outcomes in single-store model

By Proposition 1 and the constraint  $C + F = A^O$ , we see that for any optimal solution,  $C = \min(D^O, S) - \min(D^O, S, I - \min(I, D^P))$  and  $F = \min(D^O, S, I - \min(I, D^P))$ . Substituting these values, the value of the original objective function at an optimal solution becomes  $p \min(I - \min(I, D^P) - \min(D^O, S, I - \min(I, D^P)), D^O - \min(D^O, S)) + c(\min(D^O, S) - \min(D^O, S, I - \min(I, D^P)))$ . Consequently, we can combine the two stages of decision making:

$$S = \min_S E_{D^O, D^P} [p \min[I - \min(I, D^P) - \min(D^O, S, I - \min(I, D^P)), D^O - \min(D^O, S)] + c (\min[D^O, S] - \min[D^O, S, I - \min(I, D^P)])].$$

This has a closed-form optimal solution.

**THEOREM 1.** (Proof in Appendix A) For the above model,  $S = I - F_p^{-1}(\frac{c}{c+p})$ .

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

$$P[S < D^O](cP[D^P \leq S] - pP[D^P < S]).$$

The factor in this expression,  $cP[D^P \leq S] - pP[D^P < S]$ , 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 omni-channel fulfillment problem studied and the classical newsvendor model. Additionally, it is the presence of the other multiplicative factor,  $P[S < D^O]$ , that makes the function  $G(S)$  non-convex (though still unimodal).

### 3. Structural Properties of Multiple-Store 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-store instances in Section 2 to the more general multiple-location setting. We observed that the optimal threshold policy for single-store 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, though 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 non-zero 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 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.  *$G(S)$ , the expected value of the objective function of the Omni-Channel Fulfillment Model as a function of Global Threshold  $S$ , is unimodal.*

*Proof* Let  $S^*$  be an optimal solution to the Omni-Channel Fulfillment Model, a solution that minimizes  $G$ . We want to show that  $G(S+1) \leq G(S) + \delta S < S$  and  $G(S+1) \geq G(S) - \delta S > S^*$ . Our proof strategy will be to decompose  $G(S)$  into the sum of two functions  $P(S)$  and  $F(S)$ . Then, we will define functions  $G(S, D)$ ,  $P(S, D)$ , and  $F(S, D)$ , which are the functions  $G(S)$ ,  $P(S)$ , and  $F(S)$  under arbitrary demand distributions  $D = (D^O, D^P)$ , which may differ from the true demand distributions.  $P(S, D)$  represents the “missed sales” cost component of the model and  $F(S, D)$  represents the fulfillment cost component of the model.

First, we observe that, given the assumptions of our model, in the optimal solution accepted orders will be cancelled only if there is no available inventory to fulfill the orders. Consequently,  $\sum_{i,j} F_{ij} = \min(S, \sum_{i=1}^n D_i^O, \sum_{i=1}^n (I_i - \min(I_i, D_i^P)))$ . Then  $P(S, D) = E_D[p \min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \min(S, \sum_{i=1}^n D_i^O, \sum_{i=1}^n (I_i - \min(I_i, D_i^P))), \sum_{i=1}^n (D_i^O - A_i^O)]$ . Similarly, we can remove the missed sales term from the objective function of the second stage problem and the optimal solution will not change:

$$\begin{aligned} & \min \sum_{i=1}^n (c_i C_i + \sum_{j=1}^n s_{ji} F_{ji}) \\ & \text{such that } \min(D_i^P, I_i) + R_i + \sum_{j=1}^n F_{ij} = I_i, \quad \forall i \in [n] \\ & C_i + \sum_{j=1}^n F_{ji} = A_i^O, \quad \forall i \in [n] \\ & C_i, R_i, F_{ij} \geq 0, \quad \forall i, j. \end{aligned} \tag{2}$$

Now, let  $D(S)$  be the true demand distribution restricted to only outcomes where  $\sum_{i=1}^n D_i^O > S$ . We observe that  $G(S+1) > G(S)$  if and only if  $G(S+1, D(S)) > G(S, D(S))$  and likewise  $G(S+1) < G(S)$  if and only if  $G(S+1, D(S)) < G(S, D(S))$ . This is true because for specific realizations of demand where  $\sum_{i=1}^n D_i^O \leq S$ , the value of the model will be equal for Global Thresholds  $S$  and  $S+1$ . Then the realizations of demand where  $\sum_{i=1}^n D_i^O > S$  are the only ones needed to determine whether  $G(S+1)$  is greater or smaller than  $G(S)$ .

To complete the proof, we first show  $P(S + 1, D(S)) \geq P(S, D(S))$  and  $F(S + 1, D(S)) \geq F(S, D(S))$  are increasing in  $S$ . So is  $G(S + 1, D(S)) \geq G(S, D(S))$ . Finally we show  $G(S + 1) \geq G(S) \geq 0$  for  $S < S^*$  and  $G(S + 1) \leq G(S) \leq 0$  for  $S \geq S^*$  using the monotonicity.

Observe that  $P(S + 1, D(S)) \geq P(S, D(S)) = p \Pr(\sum_{i=1}^n \min(I_i, D_i^p) + S < \sum_{i=1}^n I_i)$ . This expression is clearly increasing in  $S$ , and so  $P(S + 1, D(S)) \geq P(S, D(S))$  is increasing in  $S$ . Next we show that  $F(S + 1, D(S)) \geq F(S, D(S))$  is also increasing in  $S$ . LP (2) may be viewed as a minimum cost, single-commodity flow problem in a bipartite network where nodes corresponding to each location are on one side of the bipartition and have supplies equal to the number of accepted orders at that location. Nodes corresponding to each location and a node corresponding to cancellations are on the other side. In this second set, all nodes have arcs directed to a sink node, and these arcs have capacities equal to the number of unsold units of inventory at the corresponding location and an unlimited capacity on the arc between the cancellation node and the sink. The sink node has demand equal to the sum of the supplies on the nodes representing accepted orders.

$F(S + 1, D(S)) \geq F(S, D(S))$  is the difference in expected value between fulfillment costs from accepting  $S + 1$  and  $S$  orders restricted to demand instances where there are at least  $S + 1$  online orders. An immediate consequence of Lemma 1 is that the marginal cost of fulfilling an order  $o$  in addition to a set of orders  $O$ , the difference between the minimum possible fulfillment cost of some set of orders  $O$  and the minimum possible fulfillment cost of orders  $O + fog$ , is at least as large as the marginal cost of fulfilling order  $o$  in addition to any subset of  $O$ . Suppose we have sets  $M$  and  $N$  of orders whose locations are selected at random from a common probability distribution and  $|N| > |M|$ . Then the expected marginal cost of filling an order  $o$  in addition to orders  $N$  will be greater than or equal to the expected marginal cost of filling order  $o$  in addition to set  $M$ . This is a consequence of the supermodularity property seen in Lemma 1 and happens because the probability distribution of the first  $|M|$  order locations in set  $N$  is the same as the distribution of order locations in set  $M$ . That  $F(S + 1, D(S)) \geq F(S, D(S))$  is increasing with  $S$  follows from the previous observation if we consider the case when  $|N| = |M| + 1$  and apply the observation over the probability distribution of possible orders  $o$ .

We have seen that both  $P(S + 1, D(S)) \geq P(S, D(S))$  and  $F(S + 1, D(S)) \geq F(S, D(S))$  are increasing with  $S$ , and so  $G(S + 1, D(S)) \geq G(S, D(S))$  is also increasing with  $S$ . Assuming that  $S^*$  is the optimal Global Threshold, it follows that  $G(S^* + 1) \geq G(S^*) \geq 0$ . Then  $G(S^* + 1, D(S^*)) \geq G(S^*, D(S^*)) \geq 0$ ; consequently,  $G(S^* + 1, D(S)) \geq G(S^*, D(S)) \geq 0$  and likewise  $G(S^* + 1) \geq G(S^*) \geq 0$  for  $S \geq S^*$ . Similarly,  $G(S^* - 1) \leq G(S^*) \leq 0$ . Then  $G(S^* - 1, D(S^* - 1)) \leq G(S^* - 1, D(S^* - 1)) \leq 0$ ; consequently,  $G(S^* - 1, D(S^* - 1)) \leq G(S^* - 1, D(S^* - 1)) \leq 0$  and likewise  $G(S^* - 1) \leq G(S^*) \leq 0$  for  $S \leq S^*$ .

This concludes the proof, as we have proved that  $G(S)$  is decreasing at all values of  $S$  below  $S^*$  and that  $G(S)$  is increasing at all values of  $S$  above  $S^*$ .

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 Appendix B)  $G(S_k)$ , the expected value of the objective function of the Omni-Channel 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. Recall that  $F(S)$  and  $F(S_k)$  be the linear interpolation of the integer values of functions  $G(S)$  and  $G(S_k)$ , defined in Theorems 2 and 3.

**LEMMA 2.**  *$F(S)$  and  $F(S_k)$  are quasi-convex and Lipschitz continuous.*

*Proof.* Observe that  $F(S)$  and  $F(S_k)$  are unimodal functions because  $G(S)$  and  $G(S_k)$  are unimodal functions and have global minima at integer values. The quasi-convexity of functions  $F(S)$  and  $F(S_k)$  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 a theorem 5.1 of Hazan et al. (2015) to prove that  $F(S)$  and  $F(S_k)$  can be efficiently minimized. The theorem proves that the Stochastic Normalized Gradient Descent algorithm will find an  $\epsilon$ -optimal minimum  $F(S)$  and  $F(S_k)$  with  $\text{poly}(\frac{1}{\epsilon})$  unbiased gradient estimates and optimization steps.

**THEOREM 4.** *An  $\epsilon$ -optimal minimum  $F(S)$  and  $F(S_k)$  can be obtained with  $\text{poly}(\frac{1}{\epsilon})$  unbiased gradient estimates and optimization steps by the Stochastic Normalized Gradient Descent algorithm.*

#### 4. Optimization Methods for Single-Period Multiple-Store Model

We now present an Infinitesimal Perturbation Analysis (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 Omni-Channel Fulfillment Model. Our key insight is, we can use dual values of an 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,  $A_i^O$ , 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:

$$\begin{aligned}
 & \max \sum_{i=1}^n (p \min(D_i^P, l_i) - c_i C_i + \sum_{j=1}^n (p - s_{ji}) F_{ji}) \\
 & \text{such that } \min(D_i^P, l_i) + R_i + \sum_{j=1}^n F_{ij} = l_i, \quad \forall i \in [n] \\
 & C_i + \sum_{j=1}^n F_{ji} = A_i^O, \quad \forall i \in [n] \\
 & C_i, R_i, F_{ij} \geq 0, \quad \forall i, j.
 \end{aligned} \tag{3}$$

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,  $\sum_{i=1}^n p \min(D_i^P, l_i)$ , reflects a profit of  $p$  for each unit sold in-store. The remaining terms of the objective function,  $-c_i C_i + \sum_{j=1}^n (p - s_{ji}) F_{ji}$ , reflect a profit of  $p$  for each unit sold online, with costs deducted for cancellations and shipping costs. This interpretation of LP (3) may at first seem counter-intuitive 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 (3) 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 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 (3) 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 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, because there could be orders which 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 (3) provides sensitivity information from the constraints  $C_i + \sum_{j=1}^n F_{ji} = A_i^O$ ,  $\forall i \in [n]$  using LP dual values. To access this sensitivity information, we will express these constraints as inequality constraints:  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  and  $C_i + \sum_{j=1}^n F_{ji} \geq A_i^O$ ,  $\forall i \in [n]$ . Then, we will dualize the first set of inequalities to obtain the following LP:

$$\begin{aligned} & \max \sum_{i=1}^n (p \min(D_i^P, l_i) - (c_i + M)C_i + MA_i + \sum_{j=1}^n (p - s_{ji} - M)F_{ji}) \\ & \text{such that } \min(D_i^P, l_i) + R_i + \sum_{j=1}^n F_{ij} = l_i, \forall i \in [n] \\ & C_i + \sum_{j=1}^n F_{ji} \leq A_i^O, \forall i \in [n] \\ & C_i, R_i, F_{ij} \geq 0, \forall i, j. \end{aligned} \tag{4}$$

We show in Proposition 3 that this LP (4) 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  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  to obtain gradient estimates for the Omni-Channel Fulfillment Model.

**PROPOSITION 3.** *(Proof in Appendix C) The linear program (4) has the same optimal solution as the linear program 3 when  $M > p$ .*

Using the network flow formulation of the original cost-minimization problem, we obtain:

**PROPOSITION 4.** *(Proof in 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  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  from linear program (4), where  $A_i^O$  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  $\lambda_i \mathbf{1}[S_i < D_i^O]$ , where  $\lambda_i$  is the optimal dual variable corresponding to constraint  $i$ ,  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$ , from linear program (4) is the negative partial derivative of the expected value of the objective of the multiple-store model with respect to  $S_i$ .*

*Proof* Let  $\lambda_i$  be the optimal dual variable corresponding to constraint  $i$ . Let  $\Pi(D^P, D^O) = \sum_{i=1}^n (p \min(D_i^P, I_i) - (c_i + M)C_i + MA_i + \sum_{j=1}^n (p - s_{ji} - M)F_{ji})$ , the optimal objective value of linear program (4). It is clear by inspection that  $\frac{\partial \Pi(D^P, D^O)}{\partial S_i} = \lambda_i \mathbf{1}[S_i < D_i^O]$ : If  $S_i < D_i^O$ , then  $A_i^O = S_i$ , and so an infinitesimal change to  $S_i$  will cause  $A_i^O$  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,  $\frac{\partial \Pi(D^P, D^O)}{\partial S_i} = \frac{\partial \Pi(D^P, D^O)}{\partial A_i^O} \frac{\partial A_i^O}{\partial S_i} = \frac{\partial \Pi(D^P, D^O)}{\partial A_i^O} = \lambda_i$ . In the other case,  $S_i \geq D_i^O$ , and so an infinitesimal change in  $S_i$  does not change  $A_i^O$ .  $\frac{\partial \Pi(D^P, D^O)}{\partial S_i} = \frac{\partial \Pi(D^P, D^O)}{\partial A_i^O} \frac{\partial A_i^O}{\partial S_i} = \frac{\partial \Pi(D^P, D^O)}{\partial A_i^O} \cdot 0 = 0$ . Therefore, taking expectations, we conclude  $\frac{\partial E[\Pi(D^P, D^O)]}{\partial S_i} = E[\lambda_i \mathbf{1}[S_i < D_i^O]]$

**THEOREM 5.** *From Theorem 5.1 of Hazan et al. (2015), using  $\frac{\partial \Pi(D^P, D^O)}{\partial S_i}$  as the normalized gradient, our algorithm obtains an  $\epsilon$ -optimal Local Threshold policy with  $\text{poly}(\frac{1}{\epsilon})$  total samples of linear program (4).*

*Proof* The following follows directly from the convergence Theorem 4 and the properties of the optimal dual variables in linear program (4) 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  $C_i + \sum_{j=1}^n F_{ji} - A_i^O$ ,  $\forall i \in [n]$  from linear program 4, 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  $\lambda_i$  is the mean online demand at location  $i$ ,  $f_i^O$ , the distribution from which random variable  $D_i^O$  is drawn, is a Poisson distribution, and  $g(i)$  is the dual value of constraint  $C_i + \sum_{j=1}^n F_{ji} - A_i^O$  after solving the maximization LP (4) for a demand sample, and  $G(S)$  is the expected value of the Omni-Channel Fulfillment Model with Global Threshold  $S$ . Then,  $\frac{dG(S)}{dS} = E[\sum_{i=1}^n \frac{\lambda_i}{\sum_{j=1}^n \lambda_j} (g(i)) \mathbf{1}[S < \sum_{k=1}^n D_k^O]]$ .*

*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.

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  $\sum_{i=1}^n \frac{\lambda_i}{\sum_{j=1}^n \lambda_j} (g(i)) \mathbf{1}[S < \sum_{k=1}^n D_k^O]$  as the normalized gradient, our algorithm obtains an  $\epsilon$ -optimal Global Threshold policy with  $\text{poly}(\frac{1}{\epsilon})$  total samples of linear program (4).*

*Proof* This result follows directly from the convergence Theorem 4 and the properties of the optimal dual variables in linear program (4) 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  $\tilde{A}(i) = \min\{D^O(i), S_i\}$  online orders (no different from the local threshold policy). If  $\sum_i \tilde{A}(i) > S$ , 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.

	Siloed Fulfillment	Reactive Fulfillment	Global Threshold	Local Threshold
Average Cost	471	369	116	104
Saving %	-	21.5%	75.3%	77.9%

Table 1 Average fulfillment costs across 20 full-network instances

## 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 inventory located at 30 to 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 to Siloed Fulfillment and Reactive Fulfillment policies.

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. 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 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 percent, 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 3 total inventory levels per item, and 10 items are tested. Every individual instance is evaluated by taking the average cost of each policy over 10000 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.

	Siloed Fulfill	Reactive Fulfill	Global Thresh	Local Thresh
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

Table 2 Overall average costs of each inventory balance condition for Inventory Balance experiments

Demand	Siloed Fulfill	Reactive Fulfill	Global Thresh	Local Thresh
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 3 Overall average costs of each in-store demand condition for demand magnitude experiments

Cancel Penalty	Siloed Fulfill	Reactive Fulfill	Global Thresh	Local Thresh
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

Table 4 Overall average costs of each cancellation penalty condition for cancellation cost magnitude experiments

**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 trial 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

In-Store Var	Online $\rho$	Siloed Fulfill	Reactive Fulfill	Global Thresh	Local Thresh
1.5	-.7	41.2	26.6	16.5	22.8
6	-.7	61.2	36.8	30.7	34.2
10.5	-.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	.7	25.3	15.9	15.5	15.9
6	.7	51.1	31.2	29.1	29.6
10.5	.7	65.6	38.2	36.4	36.7

**Table 5** Average cost of each covariance condition

along the diagonal of the covariance matrix 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 5. The covariance between the two stores is set such that the correlation coefficient of online demand is  $-.7$ ,  $0$ , and  $.7$  across 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-store instance and Local Threshold will no longer outperform Global Threshold, but 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 while 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  $-.7$ , compared to cost decreases of 7% and 4% for correlation coefficients  $0$  and  $.7$ , respectively. Similarly, Global Threshold policies have 13% lower cost than Local Threshold policies when in-store variance is 1.5, compared to cost decreases of 10% and 1% for variances of 6 and 10.5, respectively.

### 5.3. Hybrid Thresholds

We compare the performance of hybrid threshold with local and global thresholds in a two-store instance. As seen in Table 6, the Hybrid policy is indeed the one with the lowest cost. The Hybrid

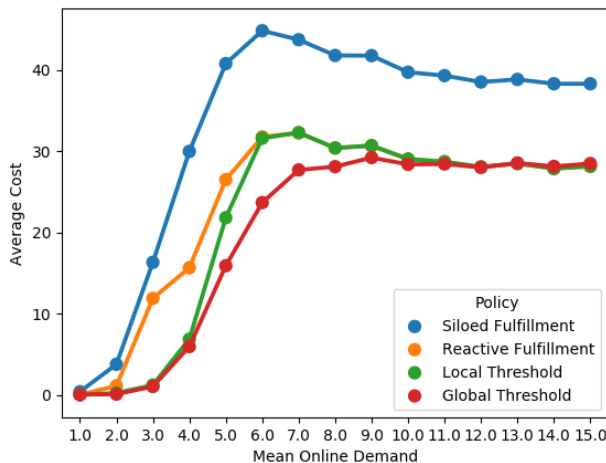


Figure 2 Average objective value as the mean online demand parameter is shifted

	Local	Global	Hybrid
Ship. Cost	3.61	5.05	4.48
Cancel. Cost	33.06	32.05	31.86
Opp. Cost	27.08	24.4	24.77
<b>Total Cost</b>	<b>63.74</b>	<b>61.5</b>	<b>61.11</b>

Table 6 Average cost of policies for  $c = 40; p = 20; s = 5$ . Initial inventory at each location is set to 20 with mean online and physical demand at Poisson rates of 10 each.

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

## 6. Multi-period 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 on-line orders; (b) physical demand in this time period is fulfilled first; (c) remaining inventory, if any, is used to satisfy on-line orders; (d) remaining on-line orders, if any, are cancelled. All future on-line orders and physical demands are lost as there is no inventory left at the location. Note that our model prioritizes current period on-line orders over next-period's physical demands.

The multi-location, multi-period model we consider here, likewise, treats on-line 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, on-line 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 on-line demand elsewhere, after its own physical demand and on-line demands are satisfied. That is, at the end of each period, on-line orders at a location are satisfied as best as possible with local inventory, may be satisfied completely by transshipped inventory, or cancelled. On-line orders are *not* carried over from period to another. Since transshipment is allowed in every period, there is no reason to “pre-transship” to balance inventories across locations.

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 Two-Period Model

For any given  $S_1$ , at the end of period 1, suppose the remaining inventory is  $I_1$ , then the decision problem for period-2 is a single-period problem with initial inventory  $I_1$ . By Theorem 1, she sets  $S_2 = I_1 - F_P^{-1}(\frac{c}{c+p})$  where  $F_P$  is the c.d.f. of the physical demand. We obtain a similar form for  $S_1$ .

DEFINITION 1. For any period-1 threshold  $S$ , define  $\lambda(S) \in [0, 1]$  s.t.  $P[I_1 > 0, I_2 > 0 | S = D_O^1] = \lambda(S) P[I_1 > 0 | S = D_O^1]$ .

THEOREM 7 (Optimality Condition). (Proof in Appendix D) Let  $I_0$  be the initial inventory. Then the optimal period-1 threshold  $S_1$  satisfies  $P[D_P^1 < I_0 - S_1] = \frac{c}{(S_1)(c+p)}$ .




Figure 3 Expected cost for different choices of  $S_1$ 's in the closed form threshold in Theorem 7 for  $N = 1; T = 2$ . We set  $c = 5; p = 10; l = 30; D_P = 5; D_O = 10$ .

The optimal  $S_1$  can thus be found by enumeration for the value of  $\lambda$ . We also consider the following myopic policy: at each period, pretend this is the last period and set the threshold using the closed-form formula in Theorem 1. This myopic policy is computationally efficient and can be easily generalized to  $T$  periods. As can be seen from Figure 3, the expected profit is unimodal as function of  $\lambda$ , and a value of  $\lambda$  between 0.64 and 0.74 matches the profit of the optimal solution. As seen in Figure 4, the myopic policy will set a higher period-1 threshold than the optimal policy. The profits are close to that of the optimal except at very low inventories.




Figure 4 Comparison between optimal and myopic policy for  $N = 1; T = 2$ . We set  $c = 20; p = 10; l = 30; D_P = 5; D_O = 10$ .

## 6.2. Single-Location Multi-Period Model

We present a dynamic programming (DP) approach for  $N = 1$  and arbitrary  $T$ . Let  $\Pi(I, t)$  be the maximum expected profit with inventory  $I$  and  $t$  fulfillment periods to go in the single store setting. Our DP estimates  $\Pi(I, t)$  for every combination  $I \in I_{max}$  and  $t \in T$  as follows.

Fix  $I, t$  and define  $V(S, D)$  to be the maximum expected net profit (revenues 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  $\Pi$  we have  $\Pi(I, t) = \max_S E_D[V(S, D)]$ , 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,  $V(S, D) = \Pi(I_1, t - 1) + \pi(I, D, S)$ , hence  $\Pi(I, t) = \max_S E_D[\Pi(I_1, t - 1) + \pi(I, D, S)]$ .*

The running time of DP increases linearly in  $T$ . Similar to the two-period situation, a myopic approach is possible here as well, where every period is treated as a final period. However, since

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 a heuristic – “*Look-ahead*” policy – that pretends that the remaining  $t$ -period situation can be viewed as a 2-period case with the second period being an aggregation of the remaining  $t - 1$  periods. In our numerical validation in Figure 5, we see that our look-ahead policy behaves similarly to the optimal policy in choosing thresholds at lower initial inventories, but yields close to optimal profit overall, while the myopic heuristic may be significantly worse.




Figure 5 Comparison between policies for  $N = 1; T = 3$  when  $c = 20; p = 10; l = 30; D_P = 5; D_O = 10$ . The myopic policy selects the higher thresholds, and incurs more cancellation risk, hence yields lower profit.

### 6.3. Two-store Two-Period Model

The two-store 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 a 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 demand vector (physical and online) is  $\mathbf{D}$ , with inventory  $\mathbf{I}$  and  $t$  periods to go. Fix  $t \geq 1, 2g$  and inventory  $\mathbf{I}$  and write  $V(\mathbf{S}, \mathbf{D}) = V(\mathbf{S}, \mathbf{D}; \mathbf{I}, t)$  for simplicity. Define

the *potential*  $\Pi(\mathbf{I}, t) = \max_{\mathbf{S}} E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D})]$ . Our objective is to maximize  $E_{\mathbf{D}}[V(\mathbf{S}, \mathbf{D})]$ . Again, our approach relies on a Bellman-type equation (which can be generalized directly for any  $N, T$ ).

**PROPOSITION 6 (General Bellman Equation).** *Let  $\mathbf{D} = (D(1), D(2))$  be the period-1 demands at the 2 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,*

$$V(\mathbf{S}, \mathbf{D}, \mathbf{I}, 2) = \max_x \left\{ \pi(\mathbf{S}, \mathbf{D}, \mathbf{I}, x) - s^T x + \Pi(\mathbf{I}_1(\mathbf{S}, \mathbf{D}, \mathbf{I}, x), 1) \right\}, \quad (5)$$

where write *max* is 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 (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. 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,  $x = \arg \max \left\{ (p + c - s)x + \Pi(R(1) - A(1) - x, 0) : 0 \leq x \leq \min\{R(1) - A(1), A(2) - R(2)\} \right\}$ .

To find the optimal  $\mathbf{S}$ , rather than enumerate over all  $\mathbf{S}$ , a faster approach is to perform gradient descent using 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  $\nabla_{\mathbf{S}} V(\mathbf{S}, \mathbf{D})$ . Figure 6 confirms that the gradient norms of the objective at each  $\mathbf{S} = (S(1), S(2))$  computed by our IPA-based heuristic almost perfectly match the discrete differences obtained via a brute-force approach.

**Models for general  $N, T$ .** As illustrated above, the key difficulties in extending our methods for the case of multi-period multi-store fulfillment is the requirement to encode the future profit potential of holding back excess inventory rather than ship them 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 multi-period problem, while 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.





**Figure 6** The gradient norms of the expected value function,  $E_D[V(S;D)]$ , for different period-1 thresholds  $S = (S(1); S(2))$ , when  $D_P = 5; D_O = 10; p = c = 10; s = 5$  and  $I = 15$ .

## 7. Conclusion

In this paper, we introduce a new stochastic two-stage model for omni-channel 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 store model. We then study the one-period multiple-store 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-store 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 while at lower inventories, a look-ahead heuristic may help. Extending our work to the full generality of multi-period models, as well as investigating alternate multi-period formulations are important directions for future work.

## Acknowledgments

## References

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

- Alishah, Elnaz Jalilipour, Kamran Moinzadeh, Yong-Pin Zhou. 2015. Inventory fulfillment strategies for an omni-channel retailer. *Available at SSRN 2659671* .
- Cattani, Kyle, Wendell Gilland, Hans Sebastian Heese, Jayashankar Swaminathan. 2006. Boiling frogs: Pricing strategies for a manufacturer adding a direct channel that competes with the traditional channel. *Production and Operations Management* **15**(1) 40.
- Gallino, Santiago, Antonio Moreno. 2019. *Operations in an Omnichannel World*. Springer.
- Gao, Fei, Xuanming Su. 2016a. Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Science* .
- Gao, Fei, Xuanming Su. 2016b. Online and offline information for omnichannel retailing. *Manufacturing & Service Operations Management* **19**(1) 84–98.
- Glasserman, Paul, Sridhar Tayur. 1995. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science* **41**(2) 263–281.
- Govindarajan, Aravind, Amitabh Sinha, Joline Uichanco. 2017. Inventory optimization for fulfillment integration in omnichannel retailing .
- Harsha, Pavithra, Shivaram Subramanian, Joline Uichanco. 2016. Omni-channel revenue management through integrated pricing and fulfillment planning. Tech. rep., Working Paper, Ross School of Business, University of Michigan.
- Hazan, Elad, Kfir Levy, Shai Shalev-Shwartz. 2015. Beyond convexity: Stochastic quasi-convex optimization. C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, R. Garnett, eds., *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc., 1594–1602.
- Herer, Yale T, Ayelet Rashit. 1999. Lateral stock transshipments in a two-location inventory system with fixed and joint replenishment costs. *Naval Research Logistics (NRL)* **46**(5) 525–547.
- Herer, Yale T, Michal Tzur, Enver Yücesan. 2006. The multilocation transshipment problem. *IIE transactions* **38**(3) 185–200.
- Jasin, Stefanus, Amitabh Sinha. 2015. An lp-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Operations Research* **63**(6) 1336–1351.
- Karaesmen, Itir, Garrett Van Ryzin. 2004. Overbooking with substitutable inventory classes. *Operations Research* **52**(1) 83–104.

## Appendix A: Proofs of Structural Properties of Single-Store Model

*Proof of Proposition 1* First, we show that for any feasible solution,  $(C, R, F)$ ,

$$F \leq \min(A^O, l - \min(l, D^P)).$$

By constraint  $C + F = A^O$ ,  $F \leq A^O$ . Similarly, by constraint  $\min(D^P, l) + R + F = l$ ,  $F \leq l - \min(D^P, l)$ . Then  $F \leq \min(A^O, l - \min(l, D^P))$ .

Now consider an arbitrary feasible solution where this inequality is not tight,  $F < \min(A^O, l - \min(l, D^P))$ . We will demonstrate that such a solution cannot be optimal. In such a solution,  $F < A^O$  and  $F < l - \min(D^P, l)$ . If  $F < A^O$  then  $C > 0$ , and if  $F < l - \min(D^P, l)$  then  $R > 0$ . Then, for any  $\epsilon < \min(R, C)$ ,  $(C - \epsilon, R - \epsilon, F + \epsilon)$  will be a feasible solution whose objective value is  $\epsilon(c + p)$  smaller than the objective value of  $(C, R, F)$ . Consequently,  $F = \min(A^O, l - \min(l, D^P))$  in any optimal solution. We recall that  $A^O = \min(D^O, S)$  by the definition of a Local Threshold policy. Therefore,

$$\begin{aligned} F &= \min(A^O, l - \min(l, D^P)) = \min(\min(D^O, S), l - \min(l, D^P)) \\ &= \min(D^O, S, l - \min(l, D^P)). \end{aligned}$$

*Proof of Theorem 1* Let us define the function  $G(S)$  as

$$\begin{aligned} G(S) &= E_{D^O, D^P} [p \min[l - \min(l, D^P) - \min(D^O, S, l - \min(l, D^P)), D^O - \min(D^O, S)] \\ &\quad + c (\min[D^O, S] - \min[D^O, S, l - \min(l, D^P)])]. \end{aligned}$$

We can also observe the derivative of this function with respect to threshold  $S$ :

$$\begin{aligned} &\frac{d}{dS} [E_{D^O, D^P} [p \min[l - \min(l, D^P) - \min(D^O, S, l - \min(l, D^P)), D^O - \min(D^O, S)] \\ &\quad + c (\min[D^O, S] - \min[D^O, S, l - \min(l, D^P)])] \\ &= \frac{d}{dS} \left[ \int_0^1 \int_0^1 f_P(x) f_O(y) (p \min[l - \min(l, x) - \min(y, S, l - \min(l, x)), y - \min(y, S)] \right. \\ &\quad \left. + c (\min[y, S] - \min[y, S, l - \min(l, x)]) dx dy \right] \\ &= \int_0^1 \int_0^1 f_P(x) f_O(y) (c \mathbf{1}[l - \min(l, x) \leq S - y] - p \mathbf{1}[S < \min(y, l - \min(l, x))]) dx dy \\ &= cP[l - \min(l, D^P) \leq S < D^O] - pP[S < \min(D^O, l - \min(l, D^P))] \\ &= cP[S < D^O]P[S \geq l - \min(l, D^P)] - pP[S < D^O]P[S < l - \min(l, D^P)] \\ &= P[S < D^O](cP[S \geq l - \min(l, D^P)] - pP[S < l - \min(l, D^P)]) \\ &= P[S < D^O](cP[D^P \geq l - S] - pP[D^P < l - S]). \end{aligned}$$

Then, assuming  $0 < F_O(x) < 1$  for  $x \in (0, I)$ ,  $G^0(S) = 0$  if and only if  $pF_P(I - S) = c(1 - F_P(I - S))$  or equivalently if  $I - S = F_P^{-1}(\frac{c}{c+p})$ . Let  $S^* = I - F_P^{-1}(\frac{c}{c+p})$ , or equivalently  $P[D^P < I - S^*] = \frac{c}{c+p}$ . If  $S < S^*$  then  $G^0(S) < 0$ :

$$\begin{aligned} G^0(S) &= P[S < D^O](cP[D^P < I - S] - pP[D^P < I - S]) \\ &= P[S < D^O](c(1 - P[D^P < I - S]) - pP[D^P < I - S]) \\ &< P[S < D^O](c - \frac{c^2}{c+p} - \frac{cp}{c+p}) \\ &= 0. \end{aligned}$$

Similarly, if  $S > S^*$  then  $G^0(S) > 0$ :

$$\begin{aligned} G^0(S) &= P[S < D^O](cP[D^P < I - S] - pP[D^P < I - S]) \\ &= P[S < D^O](c(1 - P[D^P < I - S]) - pP[D^P < I - S]) \\ &> P[S < D^O](c - \frac{c^2}{c+p} - \frac{cp}{c+p}) = 0. \end{aligned}$$

$G^0(\cdot)$  is decreasing when  $S < S^*$  and increasing when  $S > S^*$ , so  $S^*$  is an optimal threshold.

## Appendix B: Proofs of Structural Properties of Multiple-Store Model

*Proof of Lemma 1* Let  $s = (s_1, \dots, s_n)$  be the vector of supply at the source nodes, and let  $V(s) = V(s_1, \dots, s_n)$  be the objective value of the minimum cost flow in a fixed network as a function of  $s$ . We will complete the proof by arguing that

$$V(s_1 + 1, s_2 + 1, \dots, s_n) - V(s_1, s_2 + 1, \dots, s_n) \geq V(s_1 + 1, s_2, \dots, s_n) - V(s_1, s_2, \dots, s_n).$$

In other words, increasing the supply by one unit at a supply node cannot decrease the marginal cost of increasing the supply at another node. First observe that the marginal cost of increasing the supply at a specific supply node is the cost of the shortest path in the residual network obtained by computing the minimum cost flow in the network without this additional unit of demand. Suppose that the above inequality is not always true and there is a network where

$$V(s_1 + 1, s_2 + 1, \dots, s_n) - V(s_1, s_2 + 1, \dots, s_n) < V(s_1 + 1, s_2, \dots, s_n) - V(s_1, s_2, \dots, s_n).$$

Then adding a unit of supply to source node 2 decreases the marginal cost of adding supply to source node 1. For this to happen, the shortest path in the residual network between source node 1 and the sink under supplies  $(s_1, s_2, \dots, s_n)$  must be different from the shortest path in the residual network between source node 1 and the sink under supplies  $(s_1, s_2 + 1, \dots, s_n)$ . This requires the shortest path from source node 2 and the sink in the residual network under supplies  $(s_1, s_2, \dots, s_n)$  to create a new arc in the residual network by reversing flow in a saturated arc. Then, the shortest path from source node 1 and the sink in the residual network under supplies  $(s_1, s_2 + 1, \dots, s_n)$  must use this newly created arc. This cannot happen, however, because if this newly created path in the residual network (after adding supply to source node 2) is cheaper than the original path from source node 1 to the sink, this contradicts the fact that the augmenting path taken from source node 2 to the sink is a minimum cost path in that residual network.

*Proof of Theorem 3* The proof of this theorem follows from nearly the identical argument that was used to prove Theorem 2. Function  $G(S_k)$  is decomposed into the sum of  $P(S_k)$  and  $F(S_k)$ . Note that the values of  $G(S_k)$ ,  $P(S_k)$ , and  $F(S_k)$  depend on all Local Thresholds  $S_i$ ,  $\delta_i \geq [n]$ , but we express these functions as functions of  $S_k$  to indicate that  $S_k$  is a variable that can change while all other Local Thresholds  $S_j$  for  $j \neq k$  are fixed values. We observe that some of the equations change slightly, but the same arguments are true of these revised equations. Now:

$$\sum_{i,j} F_{ij} = \min\left(\sum_{i=1}^n \min(S_i, D_i^O), \sum_{i=1}^n (I_i - \min(I_i, D_i^P))\right)$$

and

$$P(S_k, D) = E_D\left[p \min\left(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \min\left(\sum_{i=1}^n \min(S_i, D_i^O), \sum_{i=1}^n (I_i - \min(I_i, D_i^P))\right), \sum_{i=1}^n (D_i^O - \min(S_i, D_i^O))\right]\right].$$

We can define  $D(S_k)$  as the true demand distribution restricted to outcomes where  $D_k^O > S_k$ , and we see that  $P(S_k + 1, D(S_k)) - P(S_k, D(S_k)) = -p \Pr(\sum_{i=1}^n \min(I_i, D_i^P) + \min(S_i, D_i^O) < \sum_{i=1}^n I_i)$ . The economic interpretation of the probability in this expression is  $p$  times the probability there is unsold inventory after physical and accepted online orders are filled. This probability is decreasing as we increase  $S_k$  and so  $P(S_k + 1, D(S_k)) - P(S_k, D(S_k))$  is increasing in  $S_k$ .  $F(S_k + 1, D(S_k)) - F(S_k, D(S_k))$  is also increasing in  $S_k$  as a direct consequence of Lemma 1 by the same argument used in the proof of Theorem 2. It follows that  $G(S_k)$  is decreasing at all values of  $S_k$  below  $S_k$  and that  $G(S_k)$  is increasing at all values of  $S_k$  above  $S_k$ , concluding the proof.

### Appendix C: Proofs Regarding Second Stage Assignment Problem

*Proof of Proposition 2* Consider again the objective function in the original minimization LP (1):

$p \min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \sum_{j=1}^n F_{ij}, \sum_{i=1}^n (D_i^O - A_i^O)) + \sum_{i=1}^n (c_i C_i + \sum_{j=1}^n s_{ji} F_{ji})$  We'll focus on the first term,

$$\min\left(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \sum_{j=1}^n F_{ij}, \sum_{i=1}^n (D_i^O - A_i^O)\right),$$

which reflects the amount of leftover inventory that could have been used to meet unfilled demand. By the constraints of the LP,  $C_i + \sum_{j=1}^n F_{ji} = A_i^O$ , so we can rewrite this as  $\min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \sum_{j=1}^n F_{ij}, \sum_{i=1}^n (D_i^O - C_i - \sum_{j=1}^n F_{ji}))$ .

We assume in the original formulation that if any cancellations occur, then there is no remaining inventory. This is equivalent to assuming that  $\max_{i,j \geq [n]} s_{ij} < p + c$ . In other words, the maximum ship cost in the network is small enough that it is always preferable to fill an online order rather than cancel the order. Note that once we have reformulated the problem we will be able to drop this assumption. Consequently, if  $C_i > 0$ , then the remaining inventory after satisfying all demands must necessarily be zero which makes the first term of the objective zero under this condition. In detail, if  $C_i > 0$ ,  $\sum_{i=1}^n (I_i - \min(I_i, D_i^P)) - \sum_{j=1}^n F_{ij} = 0$ . Suppose for a contradiction this were not true. Then we have a situation where we have cancelled an order at some

location  $j$  paying  $c$  while there is inventory left over after all fulfillment in some other location  $i$ ; However, by shipping this leftover inventory to fill this cancelled order, we avoid paying the cancellation cost  $c$ , while the increase in the cross-fulfillment reduces the first term by  $p$  and the shipping cost incurred is  $s_{ij}$ . Thus this changes the objective by  $-p - c + s_{ij}$  which is negative by our assumption, proving that our solution was non-optimal. But on the other hand, the second term of the minimizer is always non-negative:  $C_i + \sum_{j=1}^n F_{ji} = A_i^O - D_i^O$  so  $\sum_{i=1}^n (D_i^O - C_i - \sum_{j=1}^n F_{ji}) \geq 0$ . Therefore,  $\min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P) - \sum_{j=1}^n F_{ij}), \sum_{i=1}^n (D_i^O - C_i - \sum_{j=1}^n F_{ji}))$  is equivalent to  $\min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P) - \sum_{j=1}^n F_{ij}), \sum_{i=1}^n (D_i^O - \sum_{j=1}^n F_{ji}))$  for all optimal solutions to the LP. We can rewrite this as  $\min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)), \sum_{i=1}^n D_i^O - \sum_{i=1}^n \sum_{j=1}^n F_{ji})$ . This is the crucial observation that leads to the precise correspondence between LPs (1) and (3).

We rewrite the objective function using this reformulation and get the following:

$$p \min\left(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)), \sum_{i=1}^n D_i^O\right) + \sum_{i=1}^n (c_i C_i + \sum_{j=1}^n (s_{ji} - p) F_{ji}).$$

Observe that the first term,  $p \min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)), \sum_{i=1}^n D_i^O)$  no longer contains any decision variables, and all decision variables are part of the remaining terms  $\sum_{i=1}^n (c_i C_i + \sum_{j=1}^n (s_{ji} - p) F_{ji})$ . The first term is important for the economic interpretation of this LP, but only the later terms impact the quality of a feasible solution<sup>4</sup>. Then, maximizing the negative of this function and adding a constant will result in an equivalent problem. This results in the maximization LP (3).

*Proof of Proposition 3* By moving the constraints  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  to the objective function, we are relaxing the linear program (3) by allowing solutions that pay a penalty of  $M$  for each unit of violation of constraints  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$ . The cancellation variables  $C_i$  appear only in the the order acceptance constraints and have a negative coefficient in the objective, even with the violation penalty  $M$  removed. Consequently, any solution where any of constraints  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  have slack and any variable  $C_i > 0$  is not an optimal solution, because it can be improved by decreasing  $C_i$  by a sufficiently small value  $\epsilon$ .

If constraints  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  have slack,  $C_i = 0$   $\forall i \in [n]$ . In this case, there necessarily exists  $i, j$  such that  $F_{ij} > 0$ , and we will show that this solution also cannot be optimal. The solution obtained by reducing  $F_{ij}$  by  $\epsilon$  and increasing  $R_i$  by  $\epsilon$  (for a value of  $\epsilon$  smaller than the slack in the  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$  constraint) will be a feasible solution with an objective value at least  $M - p > 0$  greater than the prior solution. Consequently, no optimal solution will have slack in any of constraints  $C_i + \sum_{j=1}^n F_{ji} \leq A_i^O$ ,  $\forall i \in [n]$  therefore, this solution will also be feasible in linear program 3.

*Proof of Proposition 4* The integrality of the maximization LP (3) follows from the observation that this LP models a minimum cost feasible flow problem. The network shown in Figure 7 has exact flow requirements, indicated by the edge labels and unrestricted capacity on all unlabeled edges. Flow out of the top node on the left column of nodes represents cancellations, and all flow exiting this node incurs a cost of  $c$ . Flow into

<sup>4</sup> Note that we could add another constant term  $p \min(\sum_{i=1}^n (I_i - \min(I_i, D_i^P)), \sum_{i=1}^n D_i^O)$  to reflect the profit earned from the online orders accepted, but we left it out for simplicity.

the top node on the right column of nodes represents salvaged inventory. There is no cost to send flow to this node, which has unlimited capacity. The remaining nodes in the left column represent remaining inventory at each store after in-store demand is filled, and the remaining nodes in the right column represent accepted online orders. Flow from a left inventory node  $I_{i,a}$  to a right inventory node  $I_{j,b}$  represents inventory at location  $i$  used to fill orders at location  $j$  and has a cost  $s_{i,j} - p$ .

We will demonstrate that this minimum cost flow problem is equivalent to minimization LP (1) by writing out the flow problem as a linear program. Before we write the complete flow LP, let's enumerate its constraints moving from left to right in Figure 7. The first layer of edges sets an exact flow constraint from node  $s$  to each node  $I_{i,a}$  for all locations  $i$ . These constraints can be expressed as

$$x_{s,I_{i,a}} = l_i \quad \min(D_i^P, l_i), \quad \forall i \in [n]. \quad (6)$$

The next layer of nodes in the network has no capacity constraints, but each node has flow conservation constraints. We consider first the flow conservation constraint on node  $C$ :

$$x_{s,C} = \sum_{i=1}^n x_{C,I_{i,b}}. \quad (7)$$

The flow conservation constraints on the left inventory nodes are

$$x_{s,I_{i,a}} = x_{I_{i,a},R} + \sum_{j=1}^n x_{I_{i,a},I_{j,b}}, \quad \forall i \in [n]. \quad (8)$$

There are no edge capacities on the central edges, so we move on to the flow capacity constraints on the right side of the nodes. The flow conservation constraint on node  $R$  is

$$\sum_{i=1}^n x_{I_{i,a},R} = x_{R,t}. \quad (9)$$

The flow conservation constraints on the right inventory nodes are

$$x_{C,I_{i,b}} + \sum_{j=1}^n x_{I_{j,a},I_{i,b}} = x_{I_{i,b},t}, \quad \forall i \in [n]. \quad (10)$$

The edge capacity constraints from the right inventory nodes to  $t$  are

$$x_{I_{i,b},t} = A_i, \quad \forall i \in [n]. \quad (11)$$

Finally, flow conservation constraints on nodes  $s$  and  $t$  are

$$x_{t,s} = x_{s,C} + \sum_{j=1}^n x_{s,I_{j,a}} \quad (12)$$

$$x_{R,t} + \sum_{j=1}^n x_{I_{j,b},t} = x_{t,s}. \quad (13)$$

$x_{t,s}$  is otherwise unconstrained, so we replace this set of constraints with the following set of constraints:

$$x_{s,C} + \sum_{j=1}^n x_{s,I_{j,a}} = x_{R,t} + \sum_{j=1}^n x_{I_{j,b},t}. \quad (14)$$

Now we will reduce these constraints to the set of constraints in the minimization LP (1). Variable  $x_{s,I_{i,a}}$  is set to a fixed value by constraint 6, so we will replace  $x_{s,I_{i,a}}$  with  $l_i = \min(D_i^P, l_i)$  in all other constraints. We will rename variables  $x_{I_{i,a},I_{j,b}}$  to  $F_{i,j}$ , variable  $x_{C,I_{i,b}}$  to  $C_i$ , and  $x_{I_{i,a},R}$  to  $R_i$ ,  $\forall i \in [n]$ . We now re-state the constraints, using new variable names and replacing all variables that are constrained to a fixed value with that value:

$$x_{s,C} = \sum_{i=1}^n C_i \quad (15)$$

$$l_i = \min(D_i^P, l_i) = R_i + \sum_{j=1}^n F_{i,j}, \quad \forall i \in [n] \quad (16)$$

$$\sum_{i=1}^n R_i = x_{R,t} \quad (17)$$

$$C_i + \sum_{j=1}^n F_{j,i} = A_i, \quad \forall i \in [n] \quad (18)$$

$$x_{s,C} + \sum_{j=1}^n l_j = \min(D_j^P, l_j) = x_{R,t} + \sum_{j=1}^n A_j. \quad (19)$$

We can consolidate constraints 15, 17, and 19 into a single constraint:

$$\sum_{i=1}^n C_i + l_i = \min(D_i^P, l_i) = \sum_{i=1}^n R_i + A_i. \quad (20)$$

This results in the following constraint set for the minimum cost flow problem shown in Figure 7:

$$l_i = \min(D_i^P, l_i) = R_i + \sum_{j=1}^n F_{i,j}, \quad \forall i \in [n] \quad (21)$$

$$C_i + \sum_{j=1}^n F_{j,i} = A_i, \quad \forall i \in [n] \quad (22)$$

$$\sum_{i=1}^n l_i = \min(D_i^P, l_i) = \sum_{i=1}^n R_i + A_i = C_i. \quad (23)$$

Constraints 21 and 22 are exactly the constraints of minimization LP (1). Constraint 23 is redundant and is implied by constraints 21 and 22. Finally, we must verify that the objective functions of these two problems are the same or shifted by a constant. We will write out the objective function of the minimum cost flow problem:  $\sum_{i=1}^n c_i C_i + \sum_{j=1}^n (s_{ij} - p) F_{ij}$ .

We observed in the proof of Proposition 2 that this is equal to a constant plus the objective function of minimization LP (1). Therefore, the minimization LP (1) describes the minimum cost flow problem shown in Figure 7 and is consequently integral.

## Appendix D: Proof of Theorem 7

We fix  $D_O^1, D_P^1$  and analyze the effect of increasing  $S_1$  by  $\epsilon$ . If  $D_O^1 < S_1$ , then the objective will not change, so below we condition on the event  $D_1^O \geq S_1$ . Let  $I_1 = (I_0 - A_1^O - D_P^1)^+$  be the remaining inventory after period 1. There are two cases:



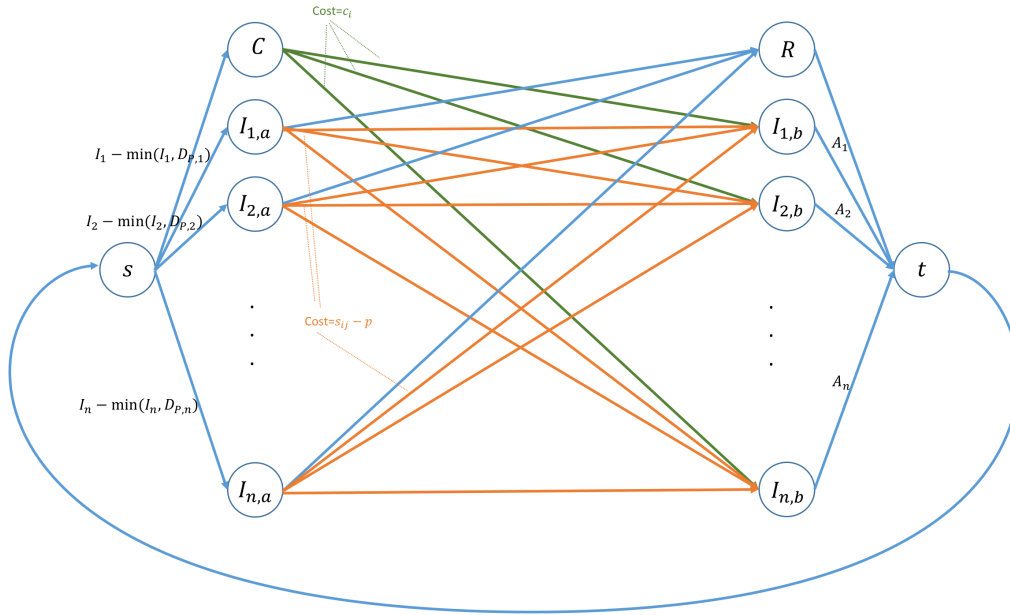


Figure 7 Minimum cost flow formulation of maximization LP 3

Case (1):  $I_1 = 0$ , i.e.  $A_O^1 + D_P^1 = I_0$ : in this case, we already have insufficient inventory to fulfill the accepted online orders, so by increasing  $S$  by  $\epsilon$ , we reduce total profit by an extra  $\epsilon c$  in cancellation fee.

Case (2):  $I_1 > 0$ , i.e.  $A_O^1 + D_P^1 < I_0$ . In this case, when we increase  $S_1$ , we sell more units in period 1, and  $I_1$  decreases. The analysis for opportunity cost of the reduced inventory requires two cases:

- (a) If  $I_2 > 0$  i.e.  $D_P^2 + A_O^2 < I_1$ , then we do save some lost sales and hence increase the profit by  $\epsilon p$ .
- (b) if  $I_2 = 0$  i.e.  $D_P^2 + A_O^2 = I_1$ , this only reduces the profit by  $\epsilon c$  due to more cancellation in period 2.

Now we can write down the exact optimality condition. Let  $P^0$  be the conditional distribution of  $P$  on the event  $\{S_1 = D_O^1\}$ . Then,

$$\begin{aligned} \frac{d}{dS_1} \mathbb{E}[Profit(S_1)] &= (c - p) P[S_1 = D_O^1 \text{ and } I_1 = 0] + p P[S_1 = D_O^1 \text{ and } I_1, I_2 > 0] - c P[S_1 = D_O^1 \text{ and } I_1 > 0 \text{ and } I_2 = 0] \\ &= (c - p) P^0[I_1 = 0] + p P^0[I_1 > 0, I_2 > 0] - c P^0[I_1 > 0, I_2 = 0] \\ &= (c - p) P^0[I_1 = 0] + p\lambda P^0[I_1 > 0] - c(1 - \lambda) P^0[I_1 > 0]. \end{aligned}$$

The last line holds for the value of  $\lambda = P^0[I_2 > 0]$ . Setting the derivative to zero and re-arranging, we have  $(c - p) P^0[I_1 = 0] = (\lambda p - (1 - \lambda)c) P^0[I_1 > 0]$ , i.e.

$$\frac{c - p}{\lambda(c + p)} = P^0[I_1 > 0] = P[A_O^1 + D_P^1 < I_0 | S_1 = D_O^1] = P[D_P^1 < I_0 - S_1 | S_1 = D_O^1].$$

Since  $D_O^1$  and  $D_P^1$  are independent, we have  $P[D_P^1 < I_0 - S_1 | S_1 = D_O^1] = P[D_P^1 < I_0 - D_O^1]$ .