

Dynamic pricing of omnichannel inventories

Pavithra Harsha

IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, pharsha@us.ibm.com

Shivaram Subramanian

IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, subshiva@us.ibm.com

Joline Uichanco

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

Omnichannel retail refers to a seamless integration of the e-commerce channel and the network of brick-and-mortar stores. New cross-channel interactions emerge from the integration: (i) cross-channel fulfillment allows inventory to be shared between channels, and (ii) accessible price information induces price-based channel substitution. However, existing price optimization systems have not been able to keep pace with these new interactions. In this paper, we propose a novel multistage stochastic program for the dynamic pricing of a product offered in an omnichannel network incorporating the new interactions. We propose a deterministic and a robust heuristic where, in each period, omnichannel prices and cross-channel fulfillment inventories are jointly optimized via computationally tractable mixed integer programs. In extensive simulations, the heuristics incur an average revenue loss of less than 4% and 8%, respectively. The key benefits of these approaches arise from inventory rebalancing using omnichannel prices (more cross-channel fulfillment inventories from stores with low sell-through rates) and through better management of channel demands.

In experiments with historical consumer electronics data of a major U.S. retailer, the proposed analytics shows a 6–12% increase in markdown revenue over the retailer’s actual revenue across different categories. A proprietary implementation of the analytics is now commercially available as part of the IBM Commerce Markdown Price Solution and has been piloted for clearance pricing at the retailer. A causal model analysis on the live pilot data shows a 12% increase in clearance period revenue.

Key words: Omnichannel, pricing, attraction demand, markdown pricing, e-commerce fulfillment, cross-channel, elasticity

1. Introduction

Omnichannel retailing is a recent trend sweeping companies across the retailing industry (Bell et al. 2014). An omnichannel strategy promises to revolutionize how companies engage with customers by creating a seamless customer shopping experience through an alignment of the retailer’s multiple sales channels. The following are just a few examples of the new capabilities enabled by omnichannel retail. A customer can purchase a product from the online store while she is in a brick-and-mortar store after finding through her mobile phone that it is offered at a cheaper price in the retailer’s online store. A customer who purchased a product online might choose a “buy online, pick up in store” option since he prefers the convenience of receiving the product sooner than having to wait

for the package to be shipped to his address. The package that an e-commerce customer receives might have been fulfilled from a nearby store since the e-commerce fulfillment center is out of stock.

From a customer’s perspective, an omnichannel environment makes it easier to compare prices between stores and online, to purchase a product from any channel, and to receive the product conveniently through any of the retailer’s multiple cross-channel fulfillment options. These capabilities allow the retailer to remain competitive in a crowded e-commerce market. With the fast pace of online sales growth ([US Census Bureau 2017](#) reported that online sales grew 14–16% compared to the previous year), retailers that are unable to quickly adapt to the changing retail landscape could be left behind, as seen through the recent wave of store closures in the US ([Rupp et al. 2017](#)).

Amazon.com, the largest U.S. e-commerce retailer, is poised to benefit from a sustained trend of online sales growth due its efficiency in fulfilling online orders from its multiple fulfillment centers. In contrast, primarily brick-and-mortar retailers only have a few fulfillment centers due to the significant fixed cost of each. Hence, an operational benefit from omnichannel integration is to enable the use of the brick-and-mortar store network to fulfill online sales via ship-from-store (SFS) fulfillment or buy-online-pickup-in-store (BOPS). Other benefits of cross-channel fulfillment include faster delivery times and the flexibility to better utilize capacity (e.g., e-commerce fulfillment can be assigned to stores with slow-moving inventory). An operations manager for a books/media retailer stated that enabling stores to fulfill e-commerce sales has driven cost down by 18% and revenue up by 20% ([Forrester Consulting 2014](#)).

Despite the many benefits of an omnichannel environment, it introduces many new challenges for price optimization, specifically to retailers whose existing pricing systems are based on traditional revenue management models. Traditional retail pricing models assume separate pools of inventory for each channel, and optimizes channel prices considering the available channel inventory. However, with store fulfillment, it is not clear which portion of store inventory is used for the online channel versus the store. Another challenge is due to potential demand substitution between the online store and brick-and-mortar store, which is affected by the prices offered on the two channels. Channel substitution is ignored in the traditional revenue management model which assumes price only affects demand in the same channel. Despite these new challenges due to the omnichannel environment, many omnichannel retailers utilize legacy price optimization systems that do not account for any channel interdependencies.

In the presence of cross-channel effects, using legacy pricing systems can have negative consequences which we demonstrate using a real-world example from a major U.S. retailer. [Fig. 1](#) shows weekly channel prices (e-commerce price and average brick-and-mortar price) and weekly channel sales for a Tablet computer that has been marked for clearance in both channels starting from Week 40. At the start of the clearance period, the legacy pricing software sets steep initial

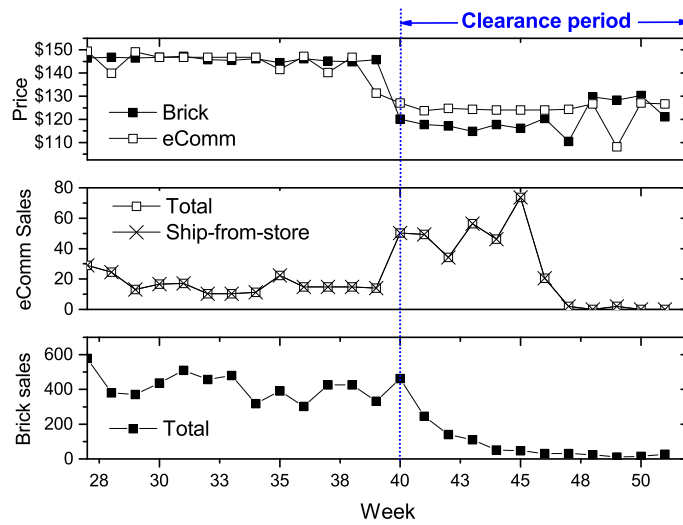


Figure 1 Time series data for the channel prices and the channel sales of a Tablet SKUs.

markdowns in the brick-and-mortar stores to clear off all store inventory over the next 12 weeks from walk-in store customers (see topmost panel). However, not all store inventory is at risk of becoming unsold since all e-commerce sales are fulfilled through ship-from-store during the clearance period (see middle panel). Because of the steep store markdowns, after an initial sales spike, total brick sales steadily decline due to an increasing number of stores stocking out (see bottom panel). In contrast to the steep brick markdowns, the legacy system initially sets the regular price for the online channel (since inventory in the e-commerce warehouse is depleted by Week 40, hence the system assumes there is no inventory left for online sales). However, due to the substitution effect, setting the online price significantly higher than brick price causes channel cannibalization. Merchandise managers would override the price outputs from the system to reduce this difference by making ad-hoc adjustments to the system inputs. A common adjustment used to set the online price (also applied to the Tablet example of Fig. 1) is to inject artificial inventory into the online channel. In discussions with merchandise managers, this manual adjustment process was revealed to be unmanageable, labor-intensive, and time consuming, and the bigger concern was the steep brick markdowns resulting in systemwide margin erosion.

As part of a joint partnership with IBM Commerce, a leading provider of merchandising solutions, we engaged with the omnichannel retailer in the example. The retailer generates over \$40 billion sales annually through its more than 1,000 brick-and-mortar stores in the U.S. and an online store. This paper describes the advanced analytics model we developed to address the novel challenges in omnichannel revenue management. The following are the contributions of our paper:

1. **Model formulation.** For a short lifecycle item without inventory replenishments, we propose a novel multistage stochastic programming formulation for the dynamic pricing problem in a

network of stores and an online channel. In the model, online sales are fulfilled using store inventories and demand exhibits channel substitution. The key idea is that we approximate the cross-channel fulfillment flows, which might be exogenously determined by an independent fulfillment engine, as endogenous model variables.

2. **Pricing policies based on store inventory partitioning.** We introduce “inventory partition” variables to approximate the optimal future use of omnichannel inventory. These variables ration the amount of store inventory to be used for online fulfillment, similar to booking limits in capacity control. In capacity control problems, the booking limits are met by accepting or rejecting customers. In contrast, the inventory partitions are met by directly influencing omnichannel demand through price management. Since the networkwide optimal inventory partitions depends on future prices, we propose two policies which jointly optimize multiperiod omnichannel prices (OCP) and inventory partitions in the cross-fulfillment network (X). We refer to this class of policies as OCPX.

- (a) The *Deterministic OCPX* policy optimizes prices and inventory partitions by solving a deterministic multiperiod non-convex optimization model under the expected demand.
- (b) The *Robust OCPX* policy chooses prices and partitions that have acceptable revenue for a set of demand sample paths (specified by a budget of uncertainty) via solving a robust max-min-max problem.

Given any demand model, if the feasible price set is discrete, then the Deterministic OCPX model can be reformulated exactly as a mixed integer program (MIP), while the Robust OCPX model can be approximated by a MIP. Moreover, a variety of business constraints can be added to both models. These policies are computationally tractable and can be solved efficiently using commercial off-the-shelf solvers such as CPLEX, making these policies suitable for a production environment. The key benefits of these approaches arise from inventory rebalancing using omnichannel prices (more SFS inventories from stores with low sell-through rates) and through better management of channel demands.

3. **Commercialization of OCPX and partnership with large U.S. retailer.** A proprietary implementation of the proposed OCPX analytics was commercialized in May 2016 as part of the IBM Commerce Markdown Price Solution. The new system was piloted at a major U.S. omnichannel retailer for clearance pricing. We demonstrate through computational experiments with historical data, independent tests of the production team, and analysis of the pilot data, that the OCPX analytics increases markdown revenue over the incumbent approaches. The OCPX system also results in shallower brick markdowns due to inventory partitions. A causal model analysis of the live pilot data shows that the improvement on clearance period revenue by OCPX is over 12%.

1.1. Literature review

Omnichannel retailing is a relatively new area, hence there are few academic papers on optimization of omnichannel operations, including pricing and/or cross-channel fulfillment. But broadly our work is in the area of revenue management, where most papers either adopt price controls (use prices to control demand(s)) or capacity control (allocate capacity of resource(s) to classes of demand). A novel aspect of this paper is that it employs both: cross-channel demand substitution that is managed through pricing, and allocation (partition) of capacities (inventories in multiple locations) optimally across channels. In addition, the capacities themselves are substitutable, as online demand can be fulfilled using inventory from any store. To better emphasize this distinguishing characteristic of the paper, we organize the literature review around these three groups.

For comprehensive surveys of dynamic pricing, see [Bitran and Caldentey \(2003\)](#), [Elmaghraby and Keskinocak \(2003\)](#), [Chen and Simchi-Levi \(2012\)](#). Single-product dynamic pricing models consider a finite amount of perishable inventory being sold to price-sensitive customers over a finite horizon. Typically, there is no inventory replenishment, an assumption applicable for hotel rooms, airline flights, and products with short selling periods and long lead times. Assuming that the demand follows a Poisson process with known time-dependent intensity as a function of price, [Gallego and van Ryzin \(1994\)](#), [Bitran and Mondschein \(1997\)](#), [Zhao and Zheng \(2000\)](#) characterize the optimal policies as functions of time and inventory. [Bitran and Mondschein \(1998\)](#) extend this basic model to a network of brick-and-mortar which coordinate prices and allow inventory transshipment between stores. However, while the model allows for inventory flows between locations (similar to cross-channel fulfillment), it only assumes the existence of a single channel and does not model substitution. [Caro and Gallien \(2012\)](#) discuss the development and implementation of a clearance pricing model for the fast-fashion retailer Zara.

Dynamic pricing models for multiple substitutable products assume that the price of one product affects the demand rates of multiple products. A popular model of incorporating customer choice in operation models is using discrete choice models. A common technique adopted by many papers is to convert the resultant nonconvex pricing problem into an equivalent convex problem in the market share or sales probability space ([Aydin and Porteus 2008](#), [Song and Xue 2007](#), [Dong et al. 2009](#), [Li and Huh 2011](#)). A recent paper by [Harsha et al. \(2015\)](#) develops a price optimization model for omnichannel demand, by modeling channel substitution using discrete choice models. Their model however assumes there are no inventory constraints or cross-channel fulfillment flows.

Our work is also related to papers where multiple products share multiple resources of limited capacity. [Gallego and van Ryzin \(1997\)](#) consider the problem of dynamic pricing for products, where products consume a fixed capacity of shared resources. [Maglaras and Meissner \(2006\)](#) study the problem where the products share a single resource. In omnichannel pricing, the channel-location

can be viewed as a “product” and the inventories in different locations as “resources”, however, consumption is not known *a priori* since it is determined by cross-channel fulfillment decisions.

Most of the above referenced papers, assume that the demands can be characterized completely using probability distribution function. But some recent papers have also considered formulations and models that avoid specifying complete distributional information for the unknown parameters. One such approach is using robust optimization formulations, based on either maximizing the minimum possible revenue (Adida and Perakis 2006, Thiele 2009), minimizing the worst-case regret (Perakis and Roels 2008), or maximizing the competitive ratio (Lan et al. 2008).

A related stream of literature is the topic of network capacity control, wherein unlike price-based controls, the goal is to allocate resources across products at fixed prices by deciding which subset of products to offer. In fact, the store inventory partitions introduced in our paper are similar in spirit to booking limits or protection limits in capacity control problems. Early papers in this topic assumed product independence (Talluri and van Ryzin 1999, Topaloglu 2009) while more recent papers model consumer substitution across products (Liu and van Ryzin 2008, Bront et al. 2009). Cross-channel fulfillment is related to supply-side substitution in capacity control literature. Gallego and Phillips (2004) introduce the idea of a “flexible product” which is a menu of alternative products. The customer purchases a flexible product knowing that the seller will assign the final product later. Gallego and Phillips (2004) models the capacity control problem of a single flexible product composed of two specific products. In the omnichannel setting, products purchased online are flexible products, since the retailer can later fulfill the purchase from a variety of store locations. Product upgrades also allow sellers to better utilize capacities. For example, Shumsky and Zhang (2009), Yu et al. (2015) study the dynamic capacity allocation problem of multiple classes, where customers purchasing a demand class can be upgraded to a higher class.

We briefly mention a few papers on pricing and/or fulfillment for a pure e-commerce retailer. Acimovic and Graves (2014) discuss the development of a dynamic fulfillment system to optimize outbound shipping costs (affected by shipping distances and split shipments). Jasin and Sinha (2015) propose heuristics based on solving a deterministic linear program approximation. A paper by Lei et al. (2016) study dynamic pricing and fulfillment of pure e-commerce retailer and similarly propose solving a deterministic approximation.

Finally, there is significant literature on multichannel pricing in marketing. Surveys by Zhang et al. (2010), Grewal et al. (2011) provide a good overview. Our work differs in that our goal is to develop a computationally tractable models and a decision support system for omnichannel pricing.

2. The omnichannel revenue management model

In this section, we introduce a multiperiod stochastic pricing model for omnichannel retailing. At the beginning of each time period, the retailer can set chain-wide prices based on the current

inventory levels across the entire retail network. Our model includes supply-side and demand-side cross-channel interactions. The former is from cross-channel inventory flows that allow store inventory to be used to fulfill online demand in any location. The latter is from modeling channel substitution by allowing the channel demand model to depend on the cross-channel price.

We describe the basic model in this section, while deferring a discussion of possible extensions to [Section 6](#). Consider a retailer selling a product with finite inventory over T sales periods. Inventory that is not sold by period T can be salvaged at a per-unit value q . Customers can purchase the product either through the retailer's online store or through any brick-and-mortar store in its network of retail stores. The retailer has one e-fulfillment center (EFC) that can fulfill any online purchase. Suppose that the retailer has stores in Z geographical zones, where we assume there is one retail store for each zone. The model can be trivially extended to the case with multiple EFC locations, multiple stores per zone, or no stores in some zones. We assume that customers in region z can only purchase from the online store or from the store in region z , and are not willing to travel to a store in another region z' . The retailer can only change prices at the start of each period.

At the start of period t , let x_e^t be the inventory level in the EFC, and let x_{bz}^t be the inventory available in store $z \in Z = (z_1, z_2, \dots, z_n)$. After observing the period t inventory levels in the retail network $x^t = (x_e^t, x_{bz_1}^t, \dots, x_{bz_n}^t)$, the retailer then determines the network-wide prices $p^t = (p_e^t, p_{bz_1}^t, \dots, p_{bz_n}^t)$. We assume that the retailer charges the same e-commerce price p_e^t for all online customers but can set retail store prices according to the store's geographic zone. Depending on the retailer's pricing strategy, alternative pricing constraints are easy to incorporate, such as setting e-commerce prices by geography or setting a uniform price offerings across channel. After setting the price, the channel-zone-level demands are then realized. For zone $z \in Z$ and channel $m \in \{e, b\}$, we model the stochastic channel-zone demand as $D_{mz}^t(p_e^t, p_{bz}^t, \xi_{mz}^t) = \xi_{mz}^t \times d_{mz}^t(p_e^t, p_{bz}^t)$, where $d_{mz}^t(\cdot)$ is a deterministic demand function and ξ_{mz}^t is a nonnegative random variable with mean 1. Hence, note that $d_{mz}^t(p_e^t, p_{bz}^t)$ is the expected demand value in channel m , zone z . Note that the expected channel demands in zone z are functions of the online price and the store price in zone z , but is independent of store prices in other zones. We denote by $D^t = (D_{ez}^t, D_{bz}^t)_{z \in Z}$ the random vector of all channel-zone demands. We refer to this as the omnichannel demand model.

The retailer also determines how to fulfill the realized online demand, either with EFC inventory or with store inventory through cross-fulfillment. In practice, price management systems are independent from the order management system (OMS) which automates the fulfillment of online sales. Therefore, fulfillment is exogenous to the price optimization model. Online orders are fulfilled by OMS based on operational cost minimization models that include a variety of costs incurred by a retailer to serve orders along with several business rules. We denote by y_{ez}^t the zone z online demand that has been fulfilled by EFC inventory, which we assume to incur a per-unit fulfillment

cost c_{ez} . We denote by $y_{z'z}^t$ the zone z online demand that has been fulfilled by store inventory from zone z' , which has a per-unit fulfillment cost $c_{z'z}$. Zone z' can either be the same zone as z or a different zone. The fulfillment costs include shipping costs and handling costs. We denote the vector of fulfillment variables at time t by $y^t = (y_{ez}^t, y_{z'z}^t)_{z, z' \in Z}$. The channel sales in zone $z \in Z$ are denoted by s_{ez}^t and s_{bz}^t . The vector of sales variables is $s^t = (s_{ez}^t, s_{bz}^t)_{z \in Z}$.

The starting inventory for the next period, $x^{t+1} = F(x^t, s^t)$, is determined from a transition function F of the current starting inventory, the sales variables, and implicitly on the realized fulfillment of online sales. Motivated by our setting of clearance products, we assume a short selling period, hence inventory is not replenished in the EFC or in the stores. This assumption is suitable for seasonal, fast-fashion items or for clearance products. Due to the short horizon, we also assume that there are no inventory holding costs and shortage costs, although the model can be easily modified to include these costs.

If Π is the set of all non-anticipatory pricing and fulfillment policies, then the optimal policy π^* maximizes the total expected profit:

$$J^* = \max_{\pi \in \Pi} E^\pi \sum_{t=1}^T r_t(x^t, p^{\pi, t}), \quad (2.1)$$

where r_t is the period t revenue. The optimal policy should satisfy the following Bellman equations:

$$V_*^t(x^t) = \max_{p^t \in \Omega} E_{\xi^t} \left(\max_{s^t \in \mathcal{S}^t(x^t, p^t, \xi^t)} p^{t\top} s^t + V_*^{t+1}(F(x^t, s^t)) \right), \quad t = 1, \dots, T, \quad (2.2)$$

$$V_*^{T+1}(x^{T+1}) = q(e^\top x^{T+1}), \quad (2.3)$$

where e is the vector of all ones, and \mathcal{S}^t is the set of all feasible sales variables (the online sales are feasible if they can be fulfilled). Since there is no explicit form of the state transition function F due to its dependence on fulfillment decisions determined by an independent OMS, solving the Bellman equations directly is challenging.

We introduce endogenous fulfillment variables y^t to approximate the exogenous fulfillment determined by the OMS. The state variables for $t+1$ then become $x_e^{t+1} = x_e^t - \sum_{z \in Z} y_{ez}^t$ and $x_{bz}^{t+1} = x_{bz}^t - s_{bz}^t - \sum_{z' \in Z} y_{z'z}^t$ for all z . We denote this linear mapping as $x^{t+1} = f(x^t, y^t, s^t)$. Since OMS determines the fulfillment decisions by minimizing operational costs, we introduce a penalty into the objective function proportional to the fulfillment costs. This results in the Bellman equations:

$$V^t(x^t) = \max_{p^t \in \Omega} E_t \left(\max_{(s^t, y^t) \in \mathcal{X}^t(x^t, p^t, \xi^t)} p^{t\top} s^t - c^\top y^t + V^{t+1}(f(x^t, y^t, s^t)) \right), \quad t = 1, \dots, T, \quad (2.4)$$

$$V^{T+1}(x^{T+1}) = q(e^\top x^{T+1}). \quad (2.5)$$

$\mathcal{X}^t(x^t, p^t, \xi^t)$ is the feasible set of period t fulfillment variables and sales variables. The optimality of the pricing policy solution to (2.4)–(2.5) under the original system depends on how accurately

the fulfillment variables y^t approximate the actual fulfillment vectors from the exogenous OMS. In [Section 5.3](#), we report the accuracy of this prediction in the commercial pilot implementation.

Feasible set \mathcal{X}^t depends on the current inventory levels and the demand realizations:

$$\mathcal{X}^t(x^t, p^t, \xi^t) = \left\{ s \geq 0, y \geq 0 \left| \begin{array}{ll} s_{mz}^t \leq \xi_{mz}^t \times d_{mz}^t(p_e^t, p_{bz}^t), & m = e, b, \forall z \in Z \\ \sum_{z \in Z} y_{ez}^t \leq x_e^t, & \\ s_{bz}^t + \sum_{z' \in Z} y_{zz'}^t \leq x_{bz}^t, & \forall z \in Z, \\ s_{ez}^t = y_{ez}^t + \sum_{z' \in Z} y_{z'z}^t, & \forall z \in Z, \end{array} \right. \right\}. \quad (2.6)$$

The first set of constraints ensure that the sales variables do not exceed the realized demand. The second constraint is that the EFC inventory used to fulfill online sales should not exceed the x_e^t . The third constraint is that the store z sales plus store inventory used for cross-fulfillment should not exceed x_{bz}^t . The last constraint imposes that all online sales in zone z must be fulfilled.

If we assume that $d_{mz}^t(p_e^t, p_{bz}^t)$ only takes nonnegative values, and since the random demand has multiplicative uncertainty, the realizations of channel-zone demands are nonnegative. Thus, starting from nonnegative inventory levels x^1 at $t = 1$, it follows that for any sample path of $\xi = (\xi^1, \dots, \xi^T)$, the sequence of feasible sets $\mathcal{X}^t(x^t, p^t, \xi^t)$ are nonempty, where $t = 1, \dots, T$.

In [\(2.4\)](#), Ω is the set of all feasible price vectors $p = (p_e, p_{bz_1}, \dots, p_{bz_n})$. This set may include any linear constraints that couple prices across time, zones or channels. For example, a business rule may be that the prices are nonincreasing over time (as often the case in clearance pricing), a group of neighboring zones offer the same store price or that the e-commerce price always be the lowest price offered and possibly a maximum markdown budget. Other linear price constraints that the retailer may want to impose are upper or lower bounds on the prices. For example, to ensure competitiveness, the retailer may want to set price bounds which are percentages of historical competitor prices. We also allow Ω to be discrete by allowing a finite number of feasible prices. For example, retailers often restrict markdowns to fixed discount levels (e.g. 20%, 30%, 40%) or prices with magic number endings (e.g. those ending with \$0.99) to exploit customer psychology.

3. Pricing policies based on inventory partitions

Solving Bellman equations [\(2.4\)–\(2.5\)](#) suffers from the curse of dimensionality since the state is a $(n + 1)$ -dimensional vector. In this section, we introduce heuristic pricing policies by approximating the value-to-go function $V^t(x^t)$ with the optimal value $\hat{V}^t(x^t)$ of a computationally tractable optimization model.

[Algorithm 1](#) provides a description of the pricing policy from an approximation \hat{V}^t of the value-to-go function. At the beginning of each time period t when we observe the current inventory levels x^t , we set the prices according to \hat{p}^t , which we define as the maximizer of function $\hat{U}(p^t; x^t)$, where:

$$\hat{U}(p^t; x^t) = \max_{(y, s) \in \mathcal{X}^t(x^t, p^t, e)} \left(p^{\top} s - c^{\top} y + \hat{V}^{t+1}(f(x^t, y, s)) \right). \quad (3.1)$$

Algorithm 1 Heuristic pricing policy**Require:** A vector of initial inventory levels x^1 , and feasible price set Ω **Ensure:** Determine a sequence of online prices and zone-level store prices

- 1: $t \leftarrow 1$
- 2: **while** $t \leq T$ **do**
- 3: Set price \hat{p}^t , where $\hat{p}^t \in \arg \max_{p^t \in \Omega} \hat{U}^t(p^t; x^t)$
- 4: Demands $D^t(p^t, \xi^t)$, sales s^t , and fulfillments y^t are realized exogenously
- 5: $x^{t+1} \leftarrow f(x^t, y^t, s^t)$
- 6: $t \leftarrow t + 1$

After setting period t prices, online demands and store demands are realized. The fulfillment and sales variables are exogenously determined by a standalone OMS. (The optimal decision variables (\hat{y}, \hat{s}) resulting from the optimization model of $\hat{U}^t(\hat{p}^t, x^t)$ are approximations for how inventory will be used at time t .) At the end of time t , the new state variables x^{t+1} are observed, and the algorithm proceeds to determine the prices for the next time period $t + 1$.

We next propose two approximations for V^t based on introducing *inventory partition* variables. These variables ration the amount of store inventory to be used for online fulfillment, similar to booking limits in capacity control. The inventory partitions are met by directly influencing omnichannel demand through price management. To motivate the partitions, let us revisit the legacy pricing system which we described in the introduction. Recall that the legacy system utilizes a pricing model that assumes a single-store environment. Therefore, the store z prices outputted by the legacy system aim to approximate the solution to the following Bellman equations:

$$V_{bz}^t(x_{bz}^t) = \max_p E_{\xi^t} \left(\max_{s \leq \min\{x_{bz}^t, D_{bz}^t(p, \xi^t)\}} ps + V_{bz}^{t+1}(x_{bz}^t - s) \right), \quad t = 1, \dots, T, \quad (3.2)$$

$$V_{bz}^{T+1}(x_{bz}^{T+1}) = qx_{bz}^{T+1}. \quad (3.3)$$

Let us denote by π^L the optimal policy which solves (3.2)–(3.3). Even if the zone z demand model in (2.4)–(2.5) does not have channel substitution (i.e., $d_{bz}^t(p_e^t, p_{bz}^t) = d_{bz}^t(p_{bz}^t)$) for all z , π^L is still be suboptimal in the omnichannel environment. This is because the single-store model assumes that store z inventory is only depleted by store z sales, yet store inventory is also used for fulfilling online sales through SFS. Hence, policy π^L would result in low store prices leading to significant margin erosion and stockouts in stores that are used for cross-channel fulfillment.

This example demonstrates that optimizing period t omnichannel prices can be aided by identifying “partitions” of store inventory units by their usage (e.g. store sales or online fulfillment). However, the true optimal partitions are unknown at time t , since they depend on future demand realizations. To see this, suppose that π^* is the optimal pricing policy that solves Bellman equations (2.4)–(2.5). The optimal portion of store z inventory to use for online fulfillment is $Y_z^{t, \pi^*} \triangleq \sum_{k=t}^T \sum_{l \in Z} y_{zl}^{k, \pi^*}$, and the optimal portion to leave in the store is $x_{bz}^t - Y_z^{t, \pi^*}$. However, since π^* is a

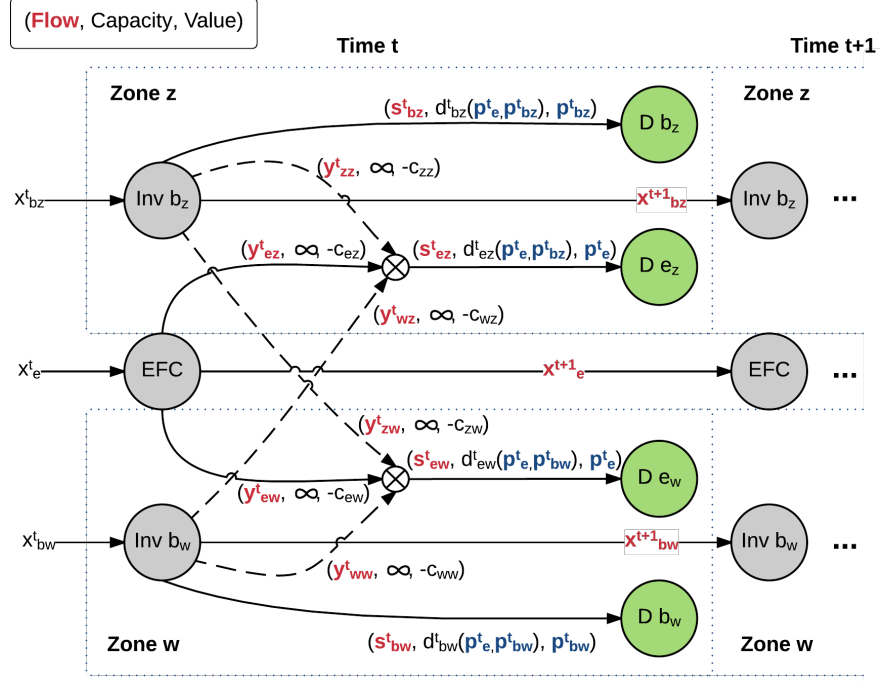


Figure 2 Time-space network for deterministic omnichannel model with two zones. The label on an arc is its (flow, capacity, value). Dashed arcs are cross-channel fulfillment flows.

non-anticipatory policy, then the value of y_{zl}^{k,π^*} is only known at time k . Hence, the value of Y_z^{t,π^*} is unknown when determining prices for time t . In this section, we propose pricing policies which approximate the optimal inventory partitions with values that can be computed in time t .

3.1. Deterministic inventory partitions

Because the optimal partition of store inventory not only depends on the fulfillment costs and salvage cost, but also on the future prices in the whole retail network, we approximate the value-to-go function V^t by the optimal value \hat{V}_D^t of a deterministic multiperiod optimization model that jointly optimizes prices $(p^t, p^{t+1}, \dots, p^T)$ and optimal inventory partitions. In particular,

$$\hat{V}_D^t(x^t) = \underset{p, s, Y, \theta}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \quad (\text{D-OC PX.1})$$

$$\text{subject to} \quad s_{ez}^k \leq d_{ez}^k(p_e^k, p_{bz}^k), \quad k = t, \dots, T, \quad \forall z \in Z, \quad (\text{D-OC PX.2})$$

$$s_{bz}^k \leq d_{bz}^k(p_e^k, p_{bz}^k), \quad k = t, \dots, T, \quad \forall z \in Z, \quad (\text{D-OC PX.3})$$

$$\sum_{z \in Z} Y_{ez} + \theta_e = x_e^t, \quad (\text{D-OC PX.4})$$

$$\sum_{k=t}^T s_{bz}^k + \sum_{z' \in Z} Y_{zz'} + \theta_{bz} = x_{bz}^t, \quad \forall z \in Z, \quad (\text{D-OC PX.5})$$

$$\sum_{k=t}^T s_{ez}^k = Y_{ez} + \sum_{z' \in Z} Y_{z'z}, \quad \forall z \in Z,$$

$$s \geq 0, Y \geq 0, \theta \geq 0, \quad (\text{D-OCPX.6})$$

$$p^k \in \Omega, \quad k = t, \dots, T. \quad (\text{D-OCPX.7})$$

The variables Y_{ez} and $Y_{zz'}$ determine the channel-zone use of EFC inventory and store inventory (i.e., inventory partitions). We refer to the optimization model above as **D-OCPX**.

Note that when there is no demand uncertainty and $\xi_{mz}^t = 1$, then V^t can be reformulated as a deterministic multiperiod optimization model. By setting $Y_{ez} = \sum_{k=t}^T y_{ez}^k$ and $Y_{zz'} = \sum_{k=t}^T y_{zz'}^k$, it is easy to check that \hat{V}_D^t is equal to V^t if demand is deterministic. Therefore, optimization model **D-OCPX** can be visualized through the time-space network illustrated in [Fig. 2](#). In the diagram, each time-zone has a brick inventory node and two demand nodes (e-commerce and brick). The dashed arcs are cross-channel fulfillment flows. The label on each arc is the flow, capacity, and value of the arc. If prices (in blue) are fixed, then the problem is a simple network flow model. This figure demonstrates that a challenge for solving **D-OCPX** is that the prices are also decision variables which determine the arc capacities and arc values. Model **D-OCPX** is nonconvex, since the objective function is bilinear, and the feasible set could be nonconvex (e.g., if the demand function is a nonconcave MNL function). However, as we will discuss in [Section 3.1.1](#), this optimization model can be reformulated as a mixed integer linear program if price set Ω is discrete (a common practical requirement).

We refer to as *Deterministic OCPX* the [Algorithm 1](#) pricing policy that results from setting $\hat{V}^{t+1} = \hat{V}_D^{t+1}$ to define \hat{U}^t in [\(3.1\)](#). Note that the period t prices which maximize [\(3.1\)](#) can be found by solving model **D-OCPX**.

3.1.1. Linearization of model D-OCPX. We are now going to assume that the feasible price set Ω is discrete, which in practice is a common constraint imposed on prices due to the requirement of magic number endings. Under a general demand model, we next discuss how we can reformulate model **D-OCPX** as a linear optimization model that can be solved to near-optimality in practice by commercial solvers such as CPLEX.

ASSUMPTION 1 (Discrete prices). $\Omega = \Omega_e \times \Omega_{bz_1} \times \dots \times \Omega_{bz_n}$, where $\Omega_e = \{\omega_i\}_{i \in I_e}$ and $\Omega_{bz} = \{\omega_i\}_{i \in I_{bz}}$ for $z \in Z$.

For ease of discussion in this section, [Assumption 1](#) allows no constraints that couple prices across zones or channels. However, these can be readily incorporated with additional linear constraints to the model. In fact, this is a major benefit of a network approach to optimizing omnichannel prices – the ability to include interchannel or interzone business constraints (such as an aggregate markdown budget). The legacy system is only able to meet such constraints heuristically, since it solves multiple separate single-store or single-channel price optimization models.

If Ω satisfies [Assumption 1](#), we can introduce the binary decision variables:

$$\mu_i^{ke} \in \{0, 1\}, \quad k = t, \dots, T, \forall i \in I_e, \quad (3.4)$$

$$\mu_j^{kz} \in \{0, 1\}, \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \quad (3.5)$$

$$\lambda_{ij}^{kz} \in \{0, 1\}, \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \forall j \in I_{bz}, \quad (3.6)$$

where $\mu_i^{ke} = 1$ if and only if $p_e^k = \omega_i$, where $\mu_j^{kz} = 1$ if and only if $p_{bz}^k = \omega_j$, and where $\lambda_{ij}^{kz} = 1$ if and only if both $p_e^k = \omega_i$ and $p_{bz}^k = \omega_j$. Hence, we impose the following linear constraints:

$$\sum_{i \in I_e} \mu_i^{ke} = 1, \quad (3.7)$$

$$\sum_{j \in I_{bz}} \mu_j^{kz} = 1, \quad \forall z \in Z, \quad (3.8)$$

$$\lambda_{ij}^{kz} \leq \mu_i^{ke}, \quad \forall z \in Z, \forall i \in I_e, \forall j \in I_{bz}, \quad (3.9)$$

$$\lambda_{ij}^{kz} \leq \mu_j^{kz}, \quad \forall z \in Z, \forall i \in I_e, \forall j \in I_{bz}, \quad (3.10)$$

We can rewrite constraint [\(D-OCPX.1\)](#)–[\(D-OCPX.2\)](#) as the following linear constraints:

$$s_{ez}^k \leq \sum_{i \in I_e} \sum_{j \in I_{bz}} \lambda_{ij}^{kz} d_{ez}^k(\omega_i, \omega_j), \quad k = t, \dots, T, \forall z \in Z, \quad (3.11)$$

$$s_{bz}^k \leq \sum_{i \in I_e} \sum_{j \in I_{bz}} \lambda_{ij}^{kz} d_{bz}^k(\omega_i, \omega_j), \quad k = t, \dots, T, \forall z \in Z. \quad (3.12)$$

And we can write the objective of [D-OCPX](#) as:

$$\sum_{k=t}^T \sum_{z \in Z} \left(\sum_{i \in I_e} \omega_i \mu_i^{ke} s_{ez}^k + \sum_{j \in I_{bz}} \omega_j \mu_j^{kz} s_{bz}^k \right) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right). \quad (3.13)$$

Note that while all the constraints are now linear, the objective function [\(3.13\)](#) is bilinear. We can simulate the nonlinear model using a linear model by “lifting” the problem to a higher dimensional space. This is done by introducing additional variables $W_{ezi}^k = \mu_i^{ke} s_{ez}^k$ and $W_{bzj}^k = \mu_j^{kz} s_{bz}^k$, for all i, j, k, z . The constraints of the original problem can be used to derive linear constraints for the higher dimensional problem. From any feasible solution of the linear model, we can obtain a feasible solution to [D-OCPX](#) by projecting it onto the original solution space. Note that any feasible solution to the bilinear model [D-OCPX](#) maps to a feasible solution of the higher-dimensional linear model, and achieves the same objective value in both models. Hence, if \bar{V}^* is the optimal value of the linear model, then \bar{V}^* is an upper bound to the optimal value of [D-OCPX](#). The following theorem shows that this bound is tight, thus the linear model solves [D-OCPX](#).

THEOREM 1. *Suppose the feasible price Ω satisfies [Assumption 1](#). Let \bar{V}^* be the optimal value of the following mixed integer program:*

$$\underset{s, Y, \theta, \mu, \lambda, W}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} \left(\sum_{i \in I_e} \omega_i W_{ezi}^k + \sum_{j \in I_{bz}} \omega_j W_{bzj}^k \right) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right)$$

subject to Constraints (3.4), (3.5), (3.6), (3.7), (3.8), (3.9), (3.10), (3.11), (3.12)

Constraints (D-OCPX.3), (D-OCPX.4), (D-OCPX.5), (D-OCPX.6)

$$\sum_{i \in I_e} W_{ezi}^k = s_{ez}^k, \quad k = t, \dots, T, \forall z \in Z, \quad (3.14)$$

$$\sum_{j \in I_{bz}} W_{bzj}^k = s_{bz}^k, \quad k = t, \dots, T, \forall z \in Z, \quad (3.15)$$

$$W_{ezi}^k \leq \sum_{j \in I_{bz}} \lambda_{ij}^{kz} d_{ez}^k(\omega_i, \omega_j), \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \quad (3.16)$$

$$W_{bzj}^k \leq \sum_{i \in I_e} \lambda_{ij}^{kz} d_{bz}^k(\omega_i, \omega_j), \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \quad (3.17)$$

$$W \geq 0. \quad (3.18)$$

Then, $\hat{V}_D^t(x^t) = \bar{V}^*$.

The proof of the theorem can be found in the electronic companion (Section EC.2). The proof constructs a feasible solution to D-OCPX from the optimal solution to the MIP which achieves an objective value \bar{V}^* , implying that $\hat{V}_D^t(x^t) = \bar{V}^*$.

Theorem 1 states that the price optimization problem D-OCPX can be solved using a MIP. When the demands follow discrete choice attraction models (see Section EC.3 in the e-companion), then it is possible to reduce the number of binary variables of the linear model from $O(I^2)$ to $O(I)$, where I refers to the maximum number of discrete prices used in any channel. Interested readers can find this result and its proof in the e-companion (Section EC.4). In our computational experiments later in the paper, when implementing this model in CPLEX for realistic problem sizes, the solver with its out-of-box parameter settings solves most instances in less than 40 seconds (see Section EC.1).

3.2. Robust inventory partitions

The heuristic in the previous section approximates the value function by assuming the stochastic demand would equal its expected value. Hence, for given price variables, the fulfillment variables y_{ez}^k and $y_{zz'}^k$, will minimize the fulfillment cost of the deterministic demand. However, when demand is uncertain, a strategy that minimizes fulfillment cost of the *expected* demand might result in significant lost in-store sales and fulfillment costs (higher than expected demand) under some sample paths or low revenue (lower than expected demand) in other some other sample paths.

As we discussed in the previous section, finding the optimal fulfillment under stochastic demand is challenging even if we can enumerate the demand realizations for each time period and use a scenario-tree approach. Thus, in this section, we will introduce a pricing policy which is based on *robust inventory partitions* determined from assuming that an adversary chooses the demand

realization that results in the worst-case firm profit for the given inventory positions. Hence, the performance of the robust pricing policy is guaranteed for many of the possible sample paths.

Applying standard robust optimization approaches to OCPX results in the adversary choosing *low demand* realizations systemwide. This policy is too conservative to implement in practice. We additionally want the adversary to consider certain high demand sample paths that result in worst case out-of-stock and fulfillment cost scenarios. We achieve this by ensuring that the budget of uncertainty set permits local demand fluctuations, yet preserves chain level demand. Optionally, we can incorporate a penalty for lost sales (i.e., denial-of-service to in-store shoppers due to stockouts) to quantify a cost for inventory imbalances in sample paths.

Suppose that we can bound the demand multiplicative factor: $\xi_{mz}^t \in [1 - \hat{\delta}_{mz}^t, 1 + \hat{\delta}_{mz}^t]$, for some $\hat{\delta}_{mz}^t \leq 1$. That is, we can express any realization of the factor as $\xi_{mz}^t = 1 + \hat{\delta}_{mz}^t w_{mz}^t$ for some $w_{mz}^t \in [-1, 1]$. It is improbable that all w factors take on extreme values. Hence, we introduce the following uncertainty set for w :

$$W(\Gamma, \Delta) = \left\{ w_{mz}^k \in [-1, 1], \forall k, m, z : \sum_{m,z} |w_{mz}^k| \leq \Gamma^k, \left| \sum_{m,z} a_{mz}^k w_{mz}^k \right| \leq \Delta^k \right\}, \quad (3.19)$$

where $\Gamma, \Delta \geq 0$ are *budgets of uncertainty*. Γ limits the number of zones and channels that have large deviations from the expected demand. Δ limits the aggregate chain level error in demand and is motivated from commonly observed behavior that aggregated demand tends to have lower error rates than its disaggregated counterparts. Here, scalars a_{mz}^k are normalization constants. For example, a statistic of the absolute percentage weekly error in estimated demand in channel m and zone z can be used as a proxy for $\hat{\delta}_{mz}^k$, and similarly at the chain level for Δ^k . The constants a_{mz}^k can be set to $\hat{\delta}_{mz}^k \bar{d}_{mz}^k / \sum_{m,z} \bar{d}_{mz}^k$ (i.e., $\hat{\delta}_{mz}^k$ weighted normalized demands), where \bar{d}_{mz}^k are the calibrated demands at observed prices. A choice of small Γ , shrinks the uncertainty set, with the deterministic problem being the extreme solution, while a choice of Δ that is close to zero retains the same level of nominal demand at chain level, yet permitting local demand fluctuations.

Using a robust framework, we assume that after the retailer decides on prices $p = (p^1, \dots, p^T)$, an “adversary” chooses demand realizations $w \in W(\Gamma, \Delta)$ which result in the *worst-case* profit. If p is the chosen price vector, then the worst-case profit is $\Pi_{\Gamma, \Delta}(p, x^t) = \min_{w \in W(\Gamma, \Delta)} \Pi(w, p; x^t)$, where $\Pi(w, p; x^t)$ is the profit realization given by:

$$\begin{aligned} \Pi(w, p; x^t) = & \underset{s, Y, \theta}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & \text{subject to} \quad s_{ez}^k \leq d_{ez}^t(p_e^t, p_{bz}^t)(1 + \hat{\delta}_{ez}^t w_{ez}^t), \quad k = t, \dots, T, \forall z \in Z, \\ & \quad s_{bz}^k \leq d_{bz}^t(p_e^t, p_{bz}^t)(1 + \hat{\delta}_{bz}^t w_{bz}^t), \quad k = t, \dots, T, \forall z \in Z, \\ & \text{Constraints (D-OCPX.3), (D-OCPX.4), (D-OCPX.5), (D-OCPX.6)} \end{aligned}$$

Because the cross-channel fulfillment and sales variables are optimally chosen by the retailer once the uncertain demand is realized, $\Pi(w, p; x^t)$ is itself an optimization problem. Faced with an adversary, the retailer's best strategy is to choose the price vector p that maximizes $\Pi_{\Gamma, \Delta}(p; x^t)$. Thus, the complexity of the robust approach depends on finding the solution to $\Pi_{\Gamma, \Delta}(p; x^t)$ efficiently.

Note that since w and p are given, $\Pi(w, p; x^t)$ is the optimal value of a linear program. Hence, by LP strong duality, we have that $\Pi(w, p; x^t)$ is equivalent to its dual (a minimization LP), and consequently $\Pi_{\Gamma, \Delta}(p; x^t)$ is the optimal value of a minimization problem with decision variables w and the dual variables of $\Pi(w, p; x^t)$. However, this minimization problem is bilinear, which in general is NP-hard to solve. However, if the feasible prices have finite upper bounds, we can formulate a linear program whose optimal value $\tilde{\Pi}_{\Gamma, \Delta}(p; x^t)$ is a lower bound on $\Pi_{\Gamma, \Delta}(p; x^t)$:

$$\begin{aligned} \tilde{\Pi}_{\Gamma, \Delta}(p; x^t) = \underset{s, Y, \theta, \chi, \lambda, \psi, \vartheta, f}{\text{maximize}} \quad & \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_z \left(c_{ez} Y_{ez} + \sum_{z'} c_{z'z} Y_{z'z} \right) + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & - \sum_k (\Gamma^k f^k + \Delta^k |\vartheta^k|) - \sum_k \sum_z \sum_{m=e, b} (A_{mz}^k |\lambda_{mz}^k| + |\chi_{mz}^k|) \\ \text{subject to} \quad & s_{mz}^k \leq d_{mz}^k(p) + |\lambda_{mz}^k| - |\psi_{mz}^k|, \quad \forall m, k, z, \end{aligned} \quad (3.20)$$

$$\lambda_{mz}^k + \psi_{mz}^k = \hat{\delta}_{mz}^k d_{mz}^k(p), \quad \forall m, k, z, \quad (3.21)$$

$$|A_{mz}^k \lambda_{mz}^k - \chi_{mz}^k - a_{mz}^k \vartheta^k| \leq f^k \quad \forall m, k, z, \quad (3.22)$$

Constraints (D-OCPX.3), (D-OCPX.4), (D-OCPX.5), (D-OCPX.6)

where $\{A_{mz}^k\}_{kmz}$ are constants that can be computed from upper bounds on prices, fulfillment costs, and the salvage value. Particularly, if P_e^k is an upper bound on the online price during period k , then set $A_{ez}^k = P_e^k - q - \min(c_{ez}, \min_{z'} c_{z'z})$. If P_{bz}^k is an upper bound on the store price during period k at zone z , then set $A_{bz}^k = P_{bz}^k - q$. The proof that $\tilde{\Pi}_{\Gamma, \Delta}(p; x^t) \leq \Pi_{\Gamma, \Delta}(p; x^t)$ can be found in the e-companion (Section EC.5). This framework can be extended for other types of polyhedral uncertainty sets, although the lower bound formulation will differ. For instance, we can consider correlated stochastic factors by setting $\xi_{mz}^k = 1 + (w_{mz}^k + \sigma^k) \delta_{mz}^k$, where w_{mz}^k is an independent uncertainty factor and σ^k is a systemic factor.

Note that the original worst-case profit $\Pi_{\Gamma, \Delta}$ is a min-max problem, with sales and inventory partition variables optimized by the innermost maximization problem. On the other hand, the lower bound $\tilde{\Pi}_{\Gamma, \Delta}$ is a maximization problem with these same variables, plus some auxiliary decision variables. From constraint (3.21), it is easy to deduce that $|\lambda_{mz}^k| - |\psi_{mz}^k|$ takes values between the possible demand deviation values in $[-\hat{\delta}_{mz}^k d_{mz}^k(p), \hat{\delta}_{mz}^k d_{mz}^k(p)]$. Therefore, $\tilde{\Pi}_{\Gamma, \Delta}$ allows sales, partition variables, and demand realizations to be determined jointly, but a penalty is incurred based on the demand realization.

Hence, we propose the following robust counterpart of the omnichannel pricing problem:

$$\hat{V}_R^t(x^t) = \max_{p=(p^1, \dots, p^T)} \tilde{\Pi}_{\Gamma, \Delta}(p; x^t), \quad (\text{R-OCPX})$$

which we refer to as **R-OCPX**. We refer to as *Robust OCPX* the [Algorithm 1](#) pricing policy when setting $\hat{V}^{t+1} = \hat{V}_R^{t+1}$ to define \hat{U}^t in (3.1). If the feasible price set satisfies [Assumption 1](#), then it is possible to linearize $\hat{V}_R^t(x^t)$ for general demand functions, similar to [Theorem 1](#), after the absolute terms in constraints (3.20) and (3.22) are linearized. This can further be reduced for attraction demand with significantly fewer binary variables, similar to [Theorem EC.1](#) in the e-companion.

4. Simulation experiments

In this section, we simulate a stochastic setting under which we compare the performance of our proposed omnichannel pricing policies to alternative pricing policies. In particular, we test the policies on 100 model instances, where each model instance has randomly chosen parameters for the attraction demand functions, the demand variability, and the fulfillment costs. On each model instance, we compare the policies by the resulting expected revenue, expected channel sales, and expected channel prices, which are calculated by applying the policies on the same 100 sample paths of demand uncertainty factors.

4.1. Data

In each of the 100 model instances, we set $T = 8$ weeks. Since the inventory in the e-fulfillment center of our partner retailer is usually depleted at the start of the clearance period, we assume that $x_e = 0$ in our experiments. For the simulations to finish in a reasonable amount of time while still introducing regional heterogeneity, we assume that there are $Z = 20$ pricing zones (compare with the commercial pilot, when $Z = 50$). Each zone has one retail store with $x_{bz} = 60$ units of inventory at the beginning of the 8 weeks. Within a zone, the expected online demand and the expected store demand follows a multinomial demand model (MNL). Zones are heterogenous in their demand parameters. We discuss how we generate the remaining parameters for a model instance and the sample paths in the e-companion ([Section EC.6](#)).

4.2. Methodology

For each of the 100 model instances in the simulation experiments, we test different pricing policies on the 100 sample paths of the random demand factors. We next describe the simulation experiment for a specific sample path n of a model instance. For each week t starting from week 1, the pricing policies (described later) output channel prices $p_e^t, (p_{bz}^t)_{z \in Z}$ based on the current inventory levels. Based on these prices and on the demand factors $(\xi_{ez}^{tn}, \xi_{bz}^{tn})_{z \in Z}$ of the sample path n , the actual demand $(D_{ez}^t, D_{bz}^t)_{z \in Z}$ is realized. The sales and fulfillment are then determined from a fulfillment

engine, which returns the inventory levels for week $t + 1$. The simulation then proceeds to the next week. Any unsold units leftover after week 8 is sold at the salvage value.

We next describe the details of the fulfillment engine used in the simulations. The engine takes in as input the week t realized zone-level demand, the week t prices, the week t store inventory levels, and the ship-from-store fulfillment cost matrix. Our aim is to design the fulfillment engine as an approximation for how fulfillment is done at the partner retailer: store customers arrive throughout the day, and fulfillment of online orders is determined periodically (e.g. after one day). Therefore, if week t is divided into K periods (e.g., 7 days), we assume that the realized week t zone-level demand is uniformly distributed over the K periods. At the end of the k^{th} period, starting from period 1, the k^{th} period zone-level store demands are met to the maximum extent possible using the same-zone store inventory. The remaining inventories in the Z stores are then used to fulfill the k^{th} period online demands through an optimization model which determines online sales and fulfillment with the objective of maximizing the myopic revenue minus fulfillment costs.

4.2.1. Pricing policies. We next describe the pricing policies that we have tested in the simulation experiments. Aside from the pricing policies proposed in this paper, we also test policies based on an adjustment to the legacy pricing model. The adjustment is meant to mimic the manual changes applied by pricing managers to outputs of the legacy model. Namely, set the store price as is (i.e., the output from optimizing store prices assuming that store inventory is only used for store demand), and set the online price using some manual adjustment to the model (i.e., inject artificial EFC inventory). The legacy demand model also assumes that that zone z online demand does not depend on the zone z store price, and that zone z store demand does not depend on the online price. The legacy demand model for zone z is the following logit demand model:

$$\begin{aligned} d_{ez}^t(p_e^t) &= N_z^t \times \frac{\exp(\alpha_{ez} - \beta_{ez}p_e^t)}{C_{ez} + \exp(\alpha_{ez} - \beta_{ez}p_e^t)}, \\ d_{bz}^t(p_{bz}^t) &= N_z^t \times \frac{\exp(\alpha_{bz} - \beta_{bz}p_{bz}^t)}{C_{bz} + \exp(\alpha_{bz} - \beta_{bz}p_{bz}^t)}, \end{aligned}$$

where we set parameters $C_{ez} = 1 + \exp(\alpha_{bz} - \beta_{bz}p^{\text{nom}})$ and $C_{bz} = 1 + \exp(\alpha_{ez} - \beta_{ez}p^{\text{nom}})$ for some parameter p^{nom} . In practice, even if cross-channel price is used as an independent variable for channel demand prediction, since the legacy model optimizes prices separately by channel, a nominal or target cross-channel price p^{nom} will need to be used to predict demand. We set $p^{\text{nom}} = \$245$ (a 30% discount from the full price). In practice, the parameters of the demand model may be recalibrated as more data is observed. However, in our simulations, because we want to compare the optimized results without the effect of demand learning, we assume that the demand parameters are static.

- **Legacy unadjusted.** Sets the week t store z prices independently by zone. It determines the price p_{bz}^t by solving a single-store multiperiod deterministic optimization model with the objective of maximizing revenue over the remaining $T - t$ weeks, and which assumes that current store z inventory x_{bz}^t is only used for store z sales. Sets the online price to be the full price of \$350.
- **Legacy with artificial EFC inventory.** This policy sets store prices in the same way as the “Legacy - Unadjusted” policy. To determine the week t online price, the policy injects artificial inventory into the online channel (a ‘virtual’ inventory pool) by setting $x_e^t = \sum_{z \in Z} 0.3x_{bz}^t$ (i.e., 30% of the total inventory).
- **OCPX Deterministic.** This policy sets week t prices based on solving optimization model **D-OCPX** for each period t , given the current store inventory levels $\{x_{bz}^t\}_{z \in Z}$.
- **OCPX Robust.** Sets week t prices based on solving the robust optimization model **R-OCPX**, given current store inventory levels store inventory levels $\{x_{bz}^t\}_{z \in Z}$. We chose uncertainty set parameters $\frac{\Gamma^k}{2Z} = 0.05$, $a_{mz}^t = \delta_{mz}^t$, and $\Delta^t = 0$. Note that this uncertainty set includes all sample paths where the mean absolute percentage error of zone-level demand is less than 5%, and the total deviation of chain-wide demand from its mean is zero (motivated by Central Limit Theorem). We set parameters A in model **R-OCPX** according to $A_{bz}^k = p_{bz}^{t-1} - q$ and $A_{bz}^k = p_e^{t-1} - q - \min(c_{ez}, \min_{z' \in Z} c_{z'z})$, for all $k = t, \dots, T$ (assume $p^0 = \$350$).
- **Legacy with OCPX store inventory partitions.** This policy determines prices by adjusting both the store inventory and online inventory using the store inventory partitions from optimal variables $\{Y_{zz'}\}_{z, z'}$ of model **D-OCPX**. It determines price p_{bz}^t by solving the multiperiod deterministic optimization model which maximizes revenue over the remaining $T - t$ weeks, assuming that the portion of store inventory used for store z sales is $x_{bz}^t - \sum_{z'} Y_{zz'}$. It then determines the week t online price by injecting artificial inventory into the online channel by setting $x_e^t = \sum_{z, z'} Y_{zz'}$.

4.3. Results

For each of the pricing policies, **Fig. 3** summarizes the distribution of price or sales statistics over 100 model instances (each model instance provides one data point for the box plots). **Fig. 3a** shows box plots representing the distribution (over 100 model instances) of the the average channel price, where the average is taken over the 100 sample paths, the 8 weeks, and the 20 zones. **Fig. 3b** shows the distribution (over 100 model instances) of standard deviation of store prices in a given week, averaged over the 8 weeks and 100 sample paths. **Fig. 3c** summarizes the distribution of average channel sales over the 100 model instances. The average sales of a model instance is taken over the 100 sample paths.

We are also interested in determining the lost revenue from using the policies to set prices. We can find a bound for this by finding the “perfect foresight” prices for each sample paths. The “perfect




Figure 3 (a) Average channel price, as percentage of regular price, (b) standard deviation in weekly store mark-downs, and (c) average channel sales, as percentage of initial inventory. The box plots show median values (solid horizontal line), 25-75th percentile values (boundary of boxes), and 5-95th percentile values (whiskers).

foresight” policy sets prices using model [D-OCPX](#) where the demands are equal to the actual demand realizations of the sample path. Averaging the optimal perfect foresight profit over the sample paths is an upper bound to the expected profit of the optimal pricing policy. This is because [D-OCPX](#) with the actual demand realizations (i) sets the optimal prices for the sample path, and (ii) sets fulfillment flows for each period after all demand is known, resulting in fulfillment decisions that are also forward-looking. We use this upper bound to compute the expected revenue loss of a nonanticipatory policy, shown in [Fig. 4](#), which is the difference between the perfect hindsight revenue and the policy revenue, averaged over 100 sample paths.




Figure 4 Expected revenue loss relative to perfect foresight. Box plots show median value (solid horizontal line), 25-75th percentile values (boundary of boxes), and 5-95th percentile values (labeled whiskers).

We first analyze the results of the legacy policies which optimize prices separately by channel-zone, and models channel demand as independent from cross-channel price. From Fig. 3a, “Legacy unadjusted” sets zero discount on the online price since there is no EFC inventory, and for majority of the instances sets store prices at deep discount (25%–55% off) to clear store inventory with in-store demand. Due to the significant difference between store and online price, this results in negligible online sales (Fig. 3c) since the true demand process is generated from a MNL model. Since almost all sales have small profit margin since they occur through the deeply discounted store channel, the average revenue loss of the unadjusted legacy policy is significant (Fig. 4); in half of the instances, it is between 11% to 16%.

The “Legacy with artificial EFC inventory” policy sets a nonzero online discount by injecting artificial inventory into the legacy optimization model to determine online price. From Fig. 3a, this results in an average online discount of around 20%. Since online prices are now closer to store prices, a small proportion of store sales (around 10%, based on Fig. 3c) are converted into online sales. Consequently, the average revenue loss is smaller (Fig. 4) because of the increase in online sales with higher profit margin.

An interesting observation from Fig. 3b is that the legacy model introduces significant variation across store prices in a given week, even if all zones start the simulation with the same store inventory levels. This is a result of the heterogeneity in demand characteristics across regions. Stores with more demand require less store price discounts than stores with less demand.

Compared to both legacy policies discussed above, OCPX Deterministic sets store prices with smaller discounts (10%–20% discount in most instances, based on Fig. 3a). This is due to the OCPX model “reducing” store inventory available for store walk-in customers through the optimized inventory partitions. These inventory partitions are also used by the model to determine the online price, resulting in online markdowns of 15%–30% for most instances. From Fig. 3b, observe the

significantly smaller variation in store prices by OCPX Deterministic. This is due to two reasons: (i) OCPX models demand substitution between channels, and prevents channel cannibalization by setting store prices close to the online price, and (ii) the store inventory partitions “re-balance” inventory according to the zone-level demand rate. Compared to the two legacy policies, the average online price is lower and the average store price is higher, resulting in higher online sales (Fig. 3c). Also, since the store profit margin of OCPX Deterministic is higher, the result is a significantly smaller revenue loss of 2.5%–6%.

From Fig. 3a, we observe that the online price set by OCPX Deterministic is lower than the average store price. This was initially surprising since we expected the online price to be higher to account for the fulfillment cost. We can understand this result by comparing the perfect foresight prices against the OCPX Deterministic prices for a specific sample path (Fig. EC.1 in e-companion). If there was no demand uncertainty, then the optimal online price is higher than the average store price. OCPX Deterministic starts the clearance period by setting a higher online price. However, since the actual demand is different from the expected demand, the inaccurate sales and inventory partitions predicted by OCPX Deterministic introduce inventory imbalances across the zones which accumulate week-to-week. As the last clearance week approaches, there is less opportunity to clear excess store inventory using store demand due to the decreasing market size; online demand then becomes a better option since it can be fulfilled with store inventory from a zone with excess store inventory regardless of the customer’s location. Hence, OCPX Deterministic increases online sales through deep online discounts towards the end of the horizon.

Unlike OCPX Deterministic, OCPX Robust explicitly accounts for demand uncertainty (through the uncertainty set) in computing inventory partitions and prices. OCPX Robust has a higher average revenue loss compared with OCPX Deterministic (Fig. 4). This is because OCPX Robust is conservative by design, since the prices should result in acceptable profit values for a variety of sample paths, which degrades its expected profit. One of the ways it accomplishes this is by setting online prices higher than store prices (Fig. 3a). Intuitively, this is because a lower online price would result in sample paths with low profit (i.e., paths where online demand is higher than expected, resulting in low profit margins due to the low price and the cost to fulfill online sales).

A final test that we wanted to conduct in these simulations is whether the optimal OCPX Deterministic inventory partitions can be used to improve the performance of the legacy markdown optimization model (Legacy with OCPX partitions). Since the OCPX Deterministic partitions are used to reduce the store inventory fed into the legacy store markdown model, we observe in Fig. 3a that the average store prices are higher than the other legacy policy variants. However, since the demand model is inaccurate, the margins from this policy are lower compared to OCPX Deterministic, resulting in a higher revenue loss under the legacy model with OCPX Deterministic

partitions. This implies that the accuracy of the legacy demand model is critical when utilizing OCPX partitions for pricing. Nevertheless, amongst the legacy markdown optimization models, the variant with OCPX partitions results in the lowest expected profit loss.

5. Partnership with a large U.S. retailer

In this section, we describe the project with the partner retailer. The retailer has annual sales of more than \$40 billion and operates more than 1,000 brick-and-mortar stores and one e-fulfillment warehouse. The objective of the project was to rollout an omnichannel pricing analytics as a part of the IBM Commerce markdown optimization solution, which the retailer can then use for clearance pricing. The project proceeded in three phases: (i) a business value assessment on historical data, (ii) an independent test by the production team, and (iii) a commercial release and pilot of the new omnichannel pricing analytics.

5.1. Phase 1: Business value assessment

We engaged with the retailer for 6 months to conduct the business value assessment (BVA) to understand the impact of the OCPX pricing system on representative product categories based on historical data. The first few weeks were spent on working with all stakeholders to define the business problem, select the categories to be analyzed, and collect and process the historical data. Thereafter, we performed the demand model calibration and the value assessment.

The retailer provided transaction log data and inventory data for three product categories (Notebooks, Tablets, and Tablet Accessories) between January 1, 2014 to December 31, 2014 (see [Table 1](#) for the data summary). In the dataset, only 195 of the SKUs sold in both channels were on clearance. These SKUs contribute to a total of 7.5 million annual clearance sales units, and total annual clearance revenues of \$107 million. On average, about 89% of lifecycle sales made in these categories is through store purchase, and about 11% is through the online channel. Almost 94% of online lifecycle sales are fulfilled using inventory shipped from a store. During clearance alone, the online sales share steeply increases to about 24% and the ship-from-store (SFS) fraction is close to 100%. We were additionally provided with price information for 18 online competitors. This competitor

Table 1 Data summary of SKUs used in the business value assessment.

Categories	# SKUs	Clearance Sales (million units)	Clearance Revenue (\$ millions)	% Sales From E-commerce	% E-commerce Ship From Store
Notebooks	100	4.1	84	13.4%	94.6%
Tablets	45	1.4	17	14.4%	93.7%
Tablet Accessories	50	2	6	5.4%	90.4%
Total	195	7.5	107	11.4%	93.8%

Table 2 Impact of OCPX on average channel prices and category-wide sales in business value assessment.

Category	Store price		Online price		% Change in Sales			Change in
	Actual	OCPX	Actual	OCPX	Online	Store	Total	Clearance Revenue
Notebooks	81%	93%	90%	85%	+16%	−1%	+3%	+10% (+\$9.0 million)
Tablets	81%	90%	79%	77%	+15%	+10%	+11%	+12% (+\$2.3 million)
Accessories	81%	93%	90%	85%	+11%	+5%	+6%	+6% (+\$1.2 million)

price data overlapped with 59 SKUs in our study and among those that had an overlap, on average, there were 6 competitors per SKU. We use this data to calibrate the SKU-zone level demand models using a procedure we describe in [Section EC.7](#) of the e-companion. The value assessment presented in this section are predictions based on this calibrated demand model.

For each of the 195 clearance SKUs, using the estimated zone-level demand models, we optimize the 10–12 week clearance prices through the MIP formulation of model [D-OCPX](#) (without reoptimization, since we do not know the “actual” demand values for the chosen prices). The retailer’s fulfillment strategy was to fulfill the item from the warehouse, if it had any inventory, and if not, the nearest in-stock store was chosen. Fulfillment cost data was not available, hence we set distance-based shipping costs, with additional cost for SFS fulfillment to adhere to priority ordering. For each SKU, the retailer provided the salvage values and the pricing business rules, which typically included a minimum time between markdowns, minimum and maximum markdown price percentages, and price bounds. We included the constraints in OCPX to simulate the revenue that OCPX would achieve when implemented with business rules applied by the retailer.

The OCPX model was developed as a Java API which was evaluated on an OS X computer with an Intel Core i7 processor. CPLEX 12.6.2 with its out-of-box parameter settings was used to solve the MIPs to optimality. The optimization model was solved at the start of the clearance period and used up to 20 price discretizations based on the initial price. Problem instances had up to 10K binary variables, 50K flow variables and 100K constraints. In comparison, a legacy approach solves a single store problems, which are significantly smaller in size. [Fig. EC.2](#) in the e-companion provides the distribution of the computational run times of the MIPs, the percentage gap of the MIP optimal solution to the LP root node, and the total number of branch-and-bound nodes in solving the MIP problem. We note that 80% of the instances solved in less than 40 seconds, with a gap of at most 1% with the root node LP with little or no branching required. Our MIP-based approach allows us solve the large scale integrated pricing problem within the required business cycle.

[Table 2](#) shows the projected effect of the OCPX model on average channel prices (normalized by regular price), on category-wide sales, and on category-wide clearance revenues. The actual store

prices are based on the outputs of the legacy markdown system using the initial store inventory level. The actual online prices are from manual adjustments by pricing managers, typically from an injection of artificial EFC inventory, on the legacy markdown system. We make the following general observations. (1) The retailer is currently setting lower store prices than OCPX store prices, agreeing with our observations on the legacy model from the simulation experiments. (2) The OCPX online prices are lower than the actual online prices, resulting in an increase in online sales activity. (3) The OCPX prices result in higher total sales (7% on average across categories with more than 23% reduction in the unsold clearance inventory). While the cause for the increase in online sales activity is obvious, the increase in store sales for two of the categories is less obvious since OCPX store prices are higher. The root cause of this increase is the conversion of lost store sales into actual sales through: directing online sales away from stores that are likely to sell out in the future (and to stores with stagnant inventories) and better demand management with omnichannel prices. In fact, the omnichannel demand model predicts more than 21% reduction in lost store sales opportunities at the recommended prices across categories. The combined effect of the optimal prices and fulfillment inventories yields an increase in the clearance revenue ranging between 6% and 12%. This translates to a combined annual increase of \$12.5 million in the total revenue from the clearance pricing of the 195 SKUs alone.

5.2. Phase 2: Independent test of the incumbent system using OCPX partitions

The BVA was presented to the retailer and reviewed by their team of pricing analysts and senior executives who were well-versed with the incumbent MDO system. The pricing analysts have long recognized that the incumbent system fails for SKUs with significant SFS fulfillment due to its inability to determine how to partition inventory between store demand and online fulfillment. Their feedback was overwhelmingly positive along with an immediate request for productizing OCPX.

A proprietary version of the analytics described in the paper was approved for commercial deployment at IBM. To productize the new OCPX solution (omnichannel demand forecasting and OCPX optimization with a multitude of business rules), we collaborated with the IBM Commerce team who developed the necessary IT infrastructure, data processing, analytics integration, and user interface. The team decided that as an intermediate step before fully transitioning to the OCPX system, the incumbent system should be modified using store partitions identified by the OCPX solution, so as to gauge benefits by SKU (i.e., with and without OCPX). To have an unbiased assessment of the effect from this adjustment, it was determined that the projected KPIs should be estimated from the incumbent demand forecasting engine. This is because the retailer had years of experience with the incumbent forecasting engine and its accuracy for various categories.




Figure 5 Gain in markdown revenue plotted against store inventory partitions for SFS fulfillment. Each point represents a product category.

There were 43 clearance SKUs (across 10 consumer electronics categories) selected to be priced with the modified incumbent system during 2015 Q4 through 2016 Q1. To mitigate the operational risk of using the modified system, the selected categories had a low use of ship-from-store fulfillment or a low online share. The projected revenue and other KPIs were estimated using the legacy forecast engine at the beginning of the clearance season under two settings: without partitions (OFF) and with OCPX partitions (ON). Since the exact inventory injections (if any) that pricing managers applied on the unadjusted legacy solution was unavailable, we cannot project the revenue of the incumbent with artificial EFC injections.

Fig. 5 presents the projected percentage improvement in category-level markdown revenue plotted against the percentage store inventory partitions for SFS fulfillment. The horizontal axis represents the degree of channel interdependency based on inventory, where 0% implies that no store inventory is used for online sale fulfillment. The figure plots the results from this independent test, and compares it with results from the business value assessment. We observe that, consistent across the tests, the revenue gain varies by category and is positively correlated with the online SFS activity. We also tracked impact on other KPIs (e.g., markdown depth, sales dollars, sell-through), which were consistent with the impact observed from the the business value assessment. All the KPI values are normalized (weighted average) by the store inventory partition units due to its dependency and the achieved benefits are reported.

The independent test showed that the modification led to a 6% increase on the average markdown revenue in the categories analyzed. The sales dollars improved by 8% while the sell-through increased by 5 percentage points. Utilizing the OCPX partitions also led to an average 2 percentage point reduction in brick-and-mortar markdown by avoiding steep and early markdowns.

We include an additional observation about the optimized prices in the 10 test categories. The price dispersion across the brick stores, empirically measured using the coefficient of variation of the highest markdown, tended to decrease due to OCPX partitions by up to 6.45%, suggesting a higher degree of price parity through the retail network. This agrees with our observations from the simulations in [Section 4](#).

5.3. Phase 3: Commercial release and pilot

The successful results from the independent test convinced both IBM Commerce and the partner retailer of the benefits achievable by the OCPX system and led to the full-scale commercialization of the proprietary version of the solution, with the partner retailer as its pilot client. The limited commercial release of the OCPX system went live in March 2016, followed with its general availability in May 2016. The capabilities of the solution were announced at the IBM Amplify 2016 conference.

Subsequent to the commercial release in Q2 2016, the retailer has gradually transitioned SKUs marked for clearance to OCPX. For the transitioned SKUs, the parameters of the omnichannel demand model are periodically recalibrated and the OCPX models are re-optimized every period. The SKUs were monitored in real-time by IBM Commerce as part of continued performance testing. For these SKUs, the quality of the recommended prices depend on how accurately the OCPX inventory partitions predict its realized values (actual store inventory used for SFS fulfillment). Hence, an inventory partition accuracy metric is additionally reported to the retailer. This accuracy metric compares the OCPX’s Week 1 optimal inventory partition solution against the actual fulfillments during the same weeks. Unlike typical forecast metrics that measure the prediction accuracy of a statistical forecast, this metric measures the prediction accuracy of an optimized decision variable. For 22 SKUs across 5 consumer electronics categories (where ON prices were executed in the independent testing phase above), the mean absolute percentage error at the chain-level was about 20%. Closing this gap requires the modeling of the market-basket fulfillment and complex operational factors and is a topic for future research.

5.3.1. Causal model analysis. A controlled experiment to assess the benefits of OCPX could not be executed for several reasons. Firstly, OCPX optimizes across the complete retail network and uses as input the inventory levels across the 1000+ stores. Hence, it is difficult to estimate the treatment effect through finding close substitutes (similar to [Caro and Gallien 2012](#)) since not only should the identified substitutes have similar demand characteristics, but the distribution of inventory at the start of clearance should also be identical. Secondly, a SKU is only in clearance for one season before being retired. Hence, we cannot apply pre-treatment (legacy) and post-treatment (OCPX) on the same SKU in different clearance seasons, precluding a Difference-in-Differences

test similar to Fisher et al. (2016). Instead, we evaluated the benefits of OCPX from the pilot data through causal model analysis via statistical adjustment using appropriate pre-treatment predictors and confounding variables (Gelman and Hill 2006, Chapter 9). We obtained a one year sample of historical data from the IBM Commerce production system for a sample of SKUs across 34 categories, where each SKU either had the OCPX treatment or not. All SKUs had online presence, and had a markdown *end date* (i.e., last date of the clearance season) in Q1 of 2017. We selected SKUs where we observed the full clearance season. For the average treatment SKU, the online sales share is 12.6% during the regular season and this nearly doubles to 24.1% during clearance season. In addition, nearly 98% of online sales during clearance season is fulfilled using SFS with approximately 20% of store inventory.

We use the following regression model to estimate the average causal effect of OCPX, where each SKU is an observation:

$$\begin{aligned} \ln(\text{Avg-Weekly-MD-Rev}_i) \sim & \alpha_0 + \alpha_1 \ln(\text{Avg-Weekly-Reg-Rev}_i) + \alpha_2 \text{Avg-Weekly-MD-Inventory}_i \\ & + \alpha_3 \text{Treatment}_i * \text{Reg-Online-Share}_i + \alpha_4 \text{Control}_i * \text{Reg-Online-Share}_i \\ & + \alpha_5 \text{Avg-Online-MD-Depth}_i + \alpha_6 \text{Avg-Store-MD-Depth}_i \quad \forall i \in \text{Items}. \end{aligned} \quad (5.1)$$

The terms Avg, MD, Reg and Rev correspond to average, markdown, regular and revenue respectively. Here, $\text{Treatment}_i = 1$ if the SKU i used OCPX and 0 otherwise, while $\text{Control}_i = 1 - \text{Treatment}_i$. The average weekly markdown and regular season revenues are obtained by dividing the total revenue in each season by their respective durations, and the average weekly markdown inventory is the initial clearance inventory divided by markdown duration. This normalization was done since the durations can vary by season for the same item and across items.

The dependent variable in the regression model is the log markdown revenue rate. Motivated by the observation in Fig. 5 that the gain in OCPX markdown revenue is proportional to use of SFS fulfillment, we interact the regular season online share with the treatment effect (For this retailer, the regular season online share is highly correlated with the SFS fulfillment during clearance). Moreover, for both treatment and control SKUs, we expect that higher the online presence, the retailer can clear the inventory better using SFS fulfillment. However, due to collinearity of variable $\text{Reg-Online-Share}_i$ with $\text{Treatment}_i * \text{Reg-Online-Share}_i$, we instead use variable $\text{Control}_i * \text{Reg-Online-Share}_i$. We also control for pre-treatment predictors, which are the log regular season revenue rate and the normalized initial clearance inventory.

The coefficients of interest are α_3 and α_4 . The difference, $\alpha_3 - \alpha_4$, can be interpreted as the percentage markdown revenue change (per percentage online share) due to OCPX. Table 3 shows the estimation results. We conduct analyses on two regression models, with and without the channel-specific markdown depth, and obtain consistent estimates of the treatment effect for both. We

Table 3 Revenue Impact of OCPX based on causal model estimated from pilot data.

$\ln(\text{Avg-Weekly-MD-Rev})$	w/o markdown depth	with markdown depth
Constant	1.820*** (0.178)	1.486*** (0.230)
$\ln(\text{Avg-Weekly-Reg-Rev})$	0.597*** (0.021)	0.596*** (0.021)
Avg-Weekly-MD-Inventory	0.004*** (0.001)	0.004*** (0.001)
Treatment * Reg-Online-Share	1.481** (0.469)	1.483*** (0.436)
Control * Reg-Online-Share	0.565** (0.263)	0.394 [‡] (0.248)
Avg-Online-MD-Depth		-0.865*** (0.220)
Avg-Store-MD-Depth		2.070*** (0.318)
# obs	275	275
# treatment	57	57
R-sq	0.832	0.855
Improvement due to OCPX	12.3%	14.8%

*** $p < .001$, ** $p < .05$, * $p < .1$, [‡] $p = .11$. SKUs with very small durations and rate of sales were eliminated.

use the treatment group’s average regular season online share ($\sim 12.6\%$) to compute the average improvement due to OCPX reported in the last row of [Table 3](#). Depending on the control variables included in the model, the markdown season revenue increase due to OCPX ranges between 12.3% to 14.8%.

6. Conclusion

The integration of a retailer’s online channel and network of retail stores have introduced new demand-side and supply-side cross-channel interactions that are often ignored in legacy pricing systems. A challenge in capturing supply-side interactions in price optimization systems is that cross-channel fulfillment is determined exogenously by a standalone order management system (OMS). In this paper, we proposed two pricing policies (Deterministic OCPX and Robust OCPX) that estimate the future use of omnichannel inventory through inventory partition variables. Both policies determine the current prices by solving a mixed integer program (MIP). Commercial off-the-shelf optimization solvers such as CPLEX can effectively solve many practical large-scaled MIPs, making the OCPX solution ideal for implementation in a production environment. A proprietary version of the OCPX system has been implemented as part of the IBM Commerce Markdown Price Solution and has been piloted by a major U.S. retailer. Causal analysis on the data from the pilot test conducted by the retailer in rolling out the new system revealed 12% improvement in

clearance period revenue. The key benefits of the OCPX policies are two-fold: Inventory rebalancing through pricing by directing online sales away from stores that are likely to sell-out towards stores that have slow moving inventories, and better management of channel demands using omni-channel prices. The proposed methods can be used, more generally, to price resources in shared inventory systems.

6.1. Possible extensions to the OCPX policies

The model discussed in the paper can be modified to account for different network setups. For instance, multiple fulfillment centers can be introduced by adding more EFC inventory nodes. Multiple stores in one zone can be modeled by multiple store inventory nodes per zone and multiple store demand functions per zone. Warehouse-to-store allocations, which allow mid-season replenishment of store inventory, can be modeled by introducing new flow variables between warehouse and stores. Practically motivated business constraints, besides those on prices, such as those on ship-from-store capacities can also be easily modeled in OCPX.

Finally, some aspects not considered in the model (but which are frequently associated with omnichannel retail) are buy-online-pickup-in-store (BOPS) fulfillment, shipping cost charged to the customer, and product returns. To model BOPS, we can introduce a zone-specific proportion β_z of online customers in zone z that will choose to pick-up from the store in the same zone. The last two constraints specifying feasible sales and fulfillment variables in (2.6) should be modified accordingly as $s_{bz}^t + \beta_z s_{ez}^t + \sum_{z' \in Z} y_{z'z}^t \leq x_{bz}^t$ and $(1 - \beta_z) s_{ez}^t = y_{ez}^t + \sum_{z' \in Z} y_{z'z}^t$. Shipping cost charged to the customer can be modeled as an additional decision variable (adding to the retailer's profit) that also affects omnichannel demand, potentially having a different elasticity compared to the prices. Product returns can also be modeled as a parameter r_z of online customers in zone z who will return to store z , with appropriate adjustments to the inventory flows in (2.6).

Acknowledgments

Firstly, we thank the *Omni-channel Retailer* who wishes to remain anonymous and more specifically, their *pricing team*, who provided the key business motivation and insights that helped in shaping the solution design. The commercialization of this work is a result of the combined efforts that includes the following team members from the IBM Commerce business unit: *Saibal Bhattacharya, Irina Fedulova, Vitaly Grechko, Nikolay Murzin, Manjunath Pandit, Emily Port, Erich Schellhas, Oleg Sidorkin, Mahesh Virupakshaiah, Maloney Whitfield, Jin Jing Xie, Emrah Zarifoglu and Pavel Zelinsky*. This list includes the data scientists, product managers and client executives. We also thank Markus Ettl and Suzanne Valentine who provided us executive support through the project.

References

- Acimovic, J., S.C. Graves. 2014. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Operations Management* **17**(1) 34–51.
- Adida, E., G. Perakis. 2006. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming* **107**(1) 97–129.
- Aydin, G., E.L. Porteus. 2008. Joint inventory and pricing decisions for an assortment. *Operations Research* **56**(5) 1247–1255.
- Bell, D. R., S. Gallino, A. Moreno. 2014. How to win in an omnichannel world. *MIT Sloan Management Review* .
- Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* **5**(3) 203–229.
- Bitran, G., S.V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Science* **43**(1) 64–79.
- Bitran, R., G. and Caldentey, S. Mondschein. 1998. Coordinating clearance markdown sales of seasonal products in retail chains. *Operations research* **46**(5) 609–624.
- Bront, J.J.M., I. Méndez-Díaz, G. Vulcano. 2009. A column generation algorithm for choice-based network revenue management. *Operations Research* **57**(3) 769–784.
- Caro, F., J. Gallien. 2012. Clearance pricing optimization for a fast-fashion retailer. *Operations Research* **60**(6) 1404–1422.
- Chen, X., D. Simchi-Levi. 2012. *The Oxford handbook of pricing management*, chap. Pricing and inventory management. Oxford University Press, 784–822.
- Dong, L., P. Kouvelis, Z. Tian. 2009. Dynamic pricing and inventory control of substitute products. *Manufacturing & Service Operations Management* **11**(2) 317–339.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* **49**(10) 1287–1309.
- Fisher, M., S. Gallino, J. Li. 2016. Competition-based dynamic pricing in online retailing: A methodology validated with field experiments .
- Forrester Consulting. 2014. *Customer desires vs. retailer capabilities: minding the omni-channel gap*. Thought leadership paper commissioned by Accenture and hybris, an SAP company.
- Gallego, G., R. Phillips. 2004. Revenue management of flexible products. *Manufacturing & Service Operations Management* **6**(4) 321–337.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40**(8) 999–1020.

-
- Gallego, G., G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* **45**(1) 24–41.
- Gelman, A., J. Hill. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.
- Grewal, Dhruv, Kusum L Ailawadi, Dinesh Gauri, Kevin Hall, Praveen Kopalle, Jane R Robertson. 2011. Innovations in retail pricing and promotions. *Journal of Retailing* **87** S43–S52.
- Harsha, P., S. Subramanian, M. Ettli. 2015. A practical price optimization approach for omni-channel retailing. Submitted.
- Jasin, S., A. Sinha. 2015. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Operations Research* **63**(6) 1336–1351.
- Lan, Y., H. Gao, M.O. Ball, I. Karaesmen. 2008. Revenue management with limited demand information. *Management Science* **54**(9) 1594–1609.
- Lei, Y., S. Jasin, A. Sinha. 2016. Dynamic joint pricing and order fulfillment for e-commerce retailers. *Available at SSRN* .
- Li, H., W.T. Huh. 2011. Pricing multiple products with the multinomial logit and nested logit models: Concavity and implications. *Manufacturing & Service Operations Management* **13**(4) 549–563.
- Liu, Q., G. van Ryzin. 2008. On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* **10**(2) 288–310.
- Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management* **8**(2) 136–148.
- Perakis, G., G. Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research* **56**(1) 188–203.
- Rupp, L., L. Coleman-Lochner, N. Turner. 2017. America’s retailers are closing stores faster than ever. URL <https://www.bloomberg.com/news/articles/2017-04-07/stores-are-closing-at-a-record-pace-as-amazon-chews-up-retailers>.
- Shumsky, R.A., F. Zhang. 2009. Dynamic capacity management with substitution. *Operations research* **57**(3) 671–684.
- Smith, S., D. Achabal. 1998. Clearance pricing and inventory policies for retail chains. *Management Science* **44**(3) 285–300.
- Song, J.-S., Z. Xue. 2007. Demand management and inventory control for substitutable products. *Working paper, Duke University* .
- Subramanian, S., P. Harsha. 2017. Demand modeling in the presence of unobserved lost sales. Submitted.
- Subramanian, S., H.D. Sherali. 2010. A fractional programming approach for retail category price optimization. *Journal of Global Optimization* **48**(2) 263–277.

-
- Talluri, K., G. van Ryzin. 1999. A randomized linear programming method for computing network bid prices. *Transportation Science* **33**(2) 207–216.
- Thiele, A. 2009. Multi-product pricing via robust optimisation. *Journal of Revenue and Pricing Management* **8**(1) 67–80.
- Topaloglu, H. 2009. Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. *Operations Research* **57**(3) 637–649.
- US Census Bureau. 2017. Quarterly retail e-commerce sales. URL https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf.
- Yu, Y., X. Chen, F. Zhang. 2015. Dynamic capacity management with general upgrading. *Operations Research* **63**(6) 1372–1389.
- Zhang, Jie, Paul W Farris, John W Irvin, Tarun Kushwaha, Thomas J Steenburgh, Barton A Weitz. 2010. Crafting integrated multichannel retailing strategies. *Journal of Interactive Marketing* **24**(2) 168–180.
- Zhao, W., Y.-S. Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management science* **46**(3) 375–388.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Electronic companion

EC.1. Additional Figures

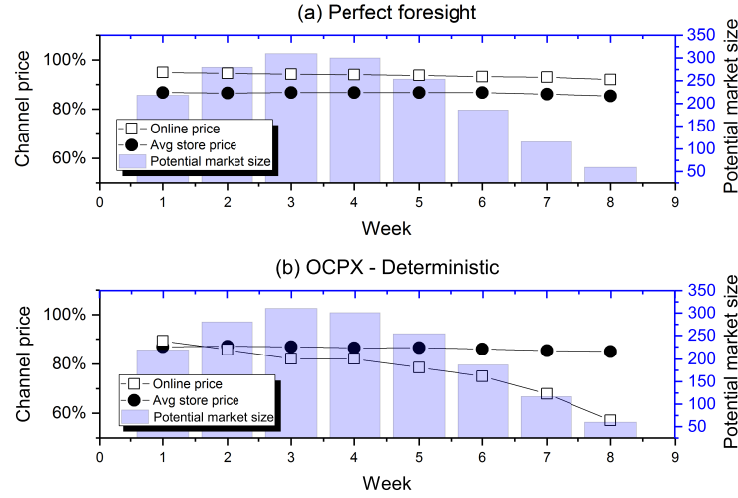


Figure EC.1 For a specific model instance and sample path, the potential market size (secondary axis) and the weekly channel prices (primary axis) of (a) perfect foresight model and (b) OCPX - Deterministic policy.




Figure EC.2 Distribution of the computational runtimes, the percentage gap of the MIP optimal solution with respect to its LP root node, and number of branch and bound nodes.

EC.2. Proof of Theorem 1

Proof. Since the feasible price set is discrete, then D-OCPX is equivalent to the optimization model with (3.13) as an objective, and with constraints (3.11)–(3.12), (D-OCPX.3)–(D-OCPX.6),

and (3.4)–(3.10). Let us denote by $(\hat{s}, \hat{Y}, \hat{\theta}, \hat{\lambda}, \hat{\mu})$ the optimal solution to this model, which has an objective value of $\hat{V}_D^t(x^t)$. Defining

$$\begin{aligned}\hat{W}_{ezi}^k &= \hat{\mu}_i^{ke} \hat{s}_{ez}^k, & k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \\ \hat{W}_{bzj}^k &= \hat{\mu}_j^{kz} \hat{s}_{bz}^k, & k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz},\end{aligned}$$

we can easily check that $(\hat{s}, \hat{Y}, \hat{\theta}, \hat{\lambda}, \hat{\mu}, \hat{W})$ is a feasible solution to the mixed integer program (MIP) given in Theorem 1 and achieves the objective value $\hat{V}_D^t(x^t)$. Hence, $\hat{V}_D^t(x^t) \leq \bar{V}^*$.

We next show that $\hat{V}_D^t(x^t) \geq \bar{V}^*$. Let us denote the maximizer of the MIP as $(\bar{s}, \bar{Y}, \bar{\theta}, \bar{\lambda}, \bar{\mu}, \bar{W})$, which achieves the optimal value \bar{V}^* . From the binary constraints of the MIP, and from constraints (3.7)–(3.8), it follows that for time k , there exist price indices $(i_{ke}, j_{kz_1}, \dots, j_{kz_n})$ such that:

$$\bar{\mu}_{i_{ke}}^{ke} = 1, \quad \bar{\mu}_i^{ke} = 0, \quad \forall i \neq i_{ke}, \quad (\text{B.1})$$

$$\bar{\mu}_{j_{kz}}^{kz} = 1, \quad \forall z \in Z, \quad \bar{\mu}_j^{kz} = 0, \quad \forall j \neq j_{kz}, \forall z \in Z. \quad (\text{B.2})$$

These then imply from constraints (3.9)–(3.10) that $\bar{\lambda}_{ij}^{kz} = 0$ for all i, j where $i \neq i_{ke}$ or $j \neq j_{kz}$, for all $z \in Z$. Therefore, from constraint (3.16), it follows that $\bar{W}_{ezi}^k = 0$ for all $i \neq i_{ke}$, $z \in Z$. Similarly, from constraint (3.17), it follows that $\bar{W}_{bzj}^k = 0$ for all $j \neq j_{kz}$, $z \in Z$. Thus, this implies from constraints (3.14)–(3.15) that $\bar{W}_{ezi_{ke}}^k = \bar{s}_{ez}^k$ for all $z \in Z$, and that $\bar{W}_{bzj_{kz}}^k = \bar{s}_{bz}^k$ for all $z \in Z$. Thus, it is easy to check based on the definitions (B.1)–(B.2) that we have the following relationships:

$$\begin{aligned}\bar{W}_{ezi}^k &= \bar{\mu}_i^{ke} \bar{s}_{ez}^k, & k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \\ \bar{W}_{bzj}^k &= \bar{\mu}_j^{kz} \bar{s}_{bz}^k, & k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz},\end{aligned}$$

Therefore $(\bar{s}, \bar{Y}, \bar{\theta}, \bar{\lambda}, \bar{\mu})$ is a feasible solution to D-OCPX with an objective value \bar{V}^* . Hence, $\bar{V}^* \leq \hat{V}_D^t(x^t)$, which proves the theorem. \square

EC.3. Discrete choice attraction demand models

Common demand functions used for substitutes are attraction demand models, for which the multinomial logit (MNL) model is a special case. For the class of attraction demand models:

$$d_{ez}^t(p_e^t, p_{bz}^t) = N_z^t \times \frac{f_{ez}^t(p_{ez}^t)}{1 + f_{bz}^t(p_{bz}^t) + f_{ez}^t(p_{ez}^t)}, \quad (\text{C.1})$$

$$d_{bz}^t(p_e^t, p_{bz}^t) = N_z^t \times \frac{f_{bz}^t(p_{bz}^t)}{1 + f_{bz}^t(p_{bz}^t) + f_{ez}^t(p_{ez}^t)}, \quad (\text{C.2})$$

where N_z^t is the market size at time t in zone z . We refer to $f_{ez}^t(\cdot), f_{bz}^t(\cdot)$ as attraction functions, which are positive and strictly decreasing functions. Note that these attraction functions are indexed by t and z , since the attractiveness of a channel may differ by zone or even by time

(due to holiday effects or seasonality effects). For the MNL model, the attraction functions are exponentially decreasing in price, i.e., $f_{mz}^t(p) = \exp(\alpha_{mz}^t - \beta_{mz}p)$ where $m = e, b$. The attraction model has been used successfully in estimating demand in econometric studies, in marketing, and in operations. Aside from MNL, other special cases of the attraction model are the Multiplicative Competitive Interaction (MCI) model, and the linear attraction model.

EC.4. Linear reformulation of **D-OCPX** under discrete prices and attraction demand models

Theorem 1 provides a linear reformulation for model **D-OCPX** for general demand models and discrete prices. Here, we provide a different linear reformulation with significantly fewer variables (**Theorem EC.1**) if the demand follows an attraction model. The reformulation uses a variable transformation similar to that in **Subramanian and Sherali (2010)**.

Suppose that the feasible price set Ω is discrete and satisfies **Assumption 1**. Suppose that the demand follows the attraction model (C.1)–(C.2). Let us introduce the following constants:

$$\phi_{ezi}^k = f_{ez}^k(\omega_i), \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \quad (\text{D.1})$$

$$\phi_{bzj}^k = f_{bz}^k(\omega_j), \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}. \quad (\text{D.2})$$

Therefore, after introducing the binary decision variables (3.4)–(3.6), we can reformulate optimization model **D-OCPX** as:

$$\begin{aligned} & \underset{s, Y, \theta, \mu}{\text{maximize}} && \sum_{k=t}^T \sum_{z \in Z} \left(\sum_{i \in I_e} \omega_i \mu_i^{ke} s_{ez}^k + \sum_{j \in I_{bz}} \omega_j \mu_j^{kz} s_{bz}^k \right) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & \text{subject to} && s_{ez}^k \leq \frac{N_z^t \sum_{i \in I_e} \phi_{ezi}^k \mu_i^{ke}}{1 + \sum_{i \in I_e} \phi_{ezi}^k \mu_i^{ke} + \sum_{j \in I_{bz}} \phi_{bzj}^k \mu_j^{kz}}, \quad k = t, \dots, T, \forall z \in Z, \end{aligned} \quad (\text{D.3})$$

$$s_{bz}^k \leq \frac{N_z^t \sum_{j \in I_{bz}} \phi_{bzj}^k \mu_j^{kz}}{1 + \sum_{i \in I_e} \phi_{ezi}^k \mu_i^{ke} + \sum_{j \in I_{bz}} \phi_{bzj}^k \mu_j^{kz}}, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{D.4})$$

Constraints (**D-OCPX.3**), (**D-OCPX.4**), (**D-OCPX.5**), (**D-OCPX.6**)

Constraints (3.4), (3.5), (3.7), (3.8)

(D.5)

As before, we denote $\hat{V}_D^t(x^t)$ as its optimal value.

THEOREM EC.1. Suppose Ω satisfies [Assumption 1](#). Let \bar{V}_2^* be the optimal value of the following mixed integer program:

$$\begin{aligned} & \underset{s, Y, \theta, \mu, W, R, G}{\text{maximize}} && \sum_{k=t}^T \sum_{z \in Z} \left(\sum_{i \in I_e} \omega_i W_{ezi}^k + \sum_{j \in I_{bz}} \omega_j W_{bjz}^k \right) - \sum_{z \in Z} c_{ez} Y_{ez} - \sum_{z \in Z} \sum_{z' \in Z} c_{zz'} Y_{zz'} + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & \text{subject to} && s_{ez}^k \leq N_z^t \sum_{i \in I_e} \phi_{ezi}^k R_{ezi}^k, \quad k = t, \dots, T, \forall z \in Z, \end{aligned} \quad (\text{P2.1})$$

$$s_{bz}^k \leq N_z^t \sum_{j \in I_{bz}} \phi_{bjz}^k R_{bjz}^k, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.2})$$

Constraints [\(D-OCPX.3\)](#), [\(D-OCPX.4\)](#), [\(D-OCPX.5\)](#), [\(D-OCPX.6\)](#)

Constraints [\(3.4\)](#), [\(3.5\)](#), [\(3.7\)](#), [\(3.8\)](#)

$$\sum_{i \in I_e} W_{ezi}^k = s_{ez}^k, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.3})$$

$$\sum_{j \in I_{bz}} W_{bjz}^k = s_{bz}^k, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.4})$$

$$W_{ezi}^k \leq N_z^k \phi_{ezi}^k R_{ezi}^k, \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \quad (\text{P2.5})$$

$$W_{bjz}^k \leq N_z^k \phi_{bjz}^k R_{bjz}^k, \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \quad (\text{P2.6})$$

$$\sum_{i \in I_e} R_{ezi}^k = G_z^k, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.7})$$

$$\sum_{j \in I_{bz}} R_{bjz}^k = G_z^k, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.8})$$

$$R_{ezi}^k \leq \mu_i^{ke}, \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e \quad (\text{P2.9})$$

$$R_{bjz}^k \leq \mu_j^{kz}, \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz} \quad (\text{P2.10})$$

$$G_z^k + \sum_{i \in I_e} \phi_{ezi}^k R_{ezi}^k + \sum_{j \in I_{bz}} \phi_{bjz}^k R_{bjz}^k = 1, \quad k = t, \dots, T, \forall z \in Z, \quad (\text{P2.11})$$

$$W \geq 0, R \geq 0, G \geq 0. \quad (\text{P2.12})$$

Then, $\hat{V}_D^t(x^t) = \bar{V}_2^*$.

Proof. Let us denote by $(\hat{s}, \hat{Y}, \hat{\theta}, \hat{\mu})$ as the optimal solution to [D-OCPX](#) under an attraction demand model, which has an objective value of $\hat{V}_D^t(x^t)$. Defining

$$\begin{aligned} \hat{W}_{ezi}^k &= \hat{\mu}_i^{ke} \hat{s}_{ez}^k, & k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \\ \hat{W}_{bjz}^k &= \hat{\mu}_j^{kz} \hat{s}_{bz}^k, & k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \\ \hat{G}_z^k &= \frac{1}{1 + \sum_{i \in I_e} \phi_{ezi}^k \hat{\mu}_i^{ke} + \sum_{j \in I_{bz}} \phi_{bjz}^k \hat{\mu}_j^{kz}}, & k = t, \dots, T, \forall z \in Z, \\ \hat{R}_{ezi}^k &= \hat{\mu}_i^{ke} \hat{G}_z^k, & k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \\ \hat{R}_{bjz}^k &= \hat{\mu}_j^{kz} \hat{G}_z^k, & k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \end{aligned}$$

we can easily check that $(\hat{s}, \hat{Y}, \hat{\theta}, \hat{\mu}, \hat{W}, \hat{R}, \hat{G})$ is a feasible solution to the mixed integer program [P2](#) and achieves the objective value $\hat{V}_D^t(x^t)$. Hence, $\hat{V}_D^t(x^t) \leq \bar{V}_2^*$.

We next show that $\hat{V}_D^t(x^t) \geq \bar{V}_2^*$. Let us denote the maximizer of model **P2** as $(\bar{s}, \bar{Y}, \bar{\theta}, \bar{\mu}, \bar{W}, \bar{R}, \bar{G})$, which achieves the optimal value \bar{V}_2^* . From the binary constraints of **P2**, and from constraints (3.7)–(3.8), it follows that for time k , there exist price indices $(i_{ke}, j_{kz_1}, \dots, j_{kz_n})$ such that:

$$\bar{\mu}_{i_{ke}}^{ke} = 1, \quad \bar{\mu}_i^{ke} = 0, \quad \forall i \neq i_{ke}, \quad (\text{D.6})$$

$$\bar{\mu}_{j_{kz}}^{kz} = 1, \quad \forall z \in Z, \quad \bar{\mu}_j^{kz} = 0, \quad \forall j \neq j_{kz}, \forall z \in Z. \quad (\text{D.7})$$

From constraints (D.6) and (P2.5), it follows that $\bar{W}_{ezi}^k = 0$ for all $i \neq i_{ke}, z \in Z$. Similarly, from constraints (D.7) and (P2.6), it follows that $\bar{W}_{bzj}^k = 0$ for all $j \neq j_{kz}, z \in Z$. Thus, these imply from constraints (P2.3)–(P2.4) that $\bar{W}_{ezi_{ke}}^k = \bar{s}_{ez}^k$ for all $z \in Z$, and that $\bar{W}_{bzj_{kz}}^k = \bar{s}_{bz}^k$ for all $z \in Z$. Thus, it is easy to check based on the definitions (D.6)–(D.7) that we have the following relationships:

$$\bar{W}_{ezi}^k = \bar{\mu}_i^{ke} \bar{s}_{ez}^k, \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \quad (\text{D.8})$$

$$\bar{W}_{bzj}^k = \bar{\mu}_j^{kz} \bar{s}_{bz}^k, \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \quad (\text{D.9})$$

From constraints (D.6) and (P2.9), it follows that $\bar{R}_{ezi}^k = 0$ for all $i \neq i_{ke}, z \in Z$. Similarly, from constraints (D.7) and (P2.10), it follows that $\bar{R}_{bzj}^k = 0$ for all $j \neq j_{kz}, z \in Z$. Thus, these imply from constraints (P2.7)–(P2.8) that $\bar{R}_{ezi_{ke}}^k = \bar{G}_z^k$ for all $z \in Z$, and that $\bar{R}_{bzj_{kz}}^k = \bar{G}_z^k$ for all $z \in Z$. Thus, it is easy to check based on the definitions (D.6)–(D.7) that we have the following relationships:

$$\bar{R}_{ezi}^k = \bar{\mu}_i^{ke} \bar{G}_z^k, \quad k = t, \dots, T, \forall z \in Z, \forall i \in I_e, \quad (\text{D.10})$$

$$\bar{R}_{bzj}^k = \bar{\mu}_j^{kz} \bar{G}_z^k, \quad k = t, \dots, T, \forall z \in Z, \forall j \in I_{bz}, \quad (\text{D.11})$$

From (D.10)–(D.11), and from (P2.11), it follows that $1 = \bar{G}_z^k + \phi_{ezi_{ke}}^k \bar{G}_z^k + \phi_{bzj_{kz}}^k \bar{G}_z^k$. Thus,

$$\bar{G}_z^k = \frac{1}{1 + \phi_{ezi_{ke}}^k + \phi_{bzj_{kz}}^k} = \frac{1}{1 + \sum_{i \in I_e} \phi_{ezi}^k \bar{\mu}_i^{ke} + \sum_{j \in I_{bz}} \phi_{bzj}^k \bar{\mu}_j^{kz}}. \quad (\text{D.12})$$

Substituting (D.10)–(D.11) into (P2.1)–(P2.2), and using the relationship (D.12), we have

$$\bar{s}_{ez}^t \leq N_z^k \phi_{ezi_{ke}}^k \bar{G}_z^k = \frac{N_z^k \sum_{i \in I_e} \phi_{ezi}^k \bar{\mu}_i^{ke}}{1 + \sum_{i \in I_e} \phi_{ezi}^k \bar{\mu}_i^{ke} + \sum_{j \in I_{bz}} \phi_{bzj}^k \bar{\mu}_j^{kz}}, \quad (\text{D.13})$$

$$\bar{s}_{bz}^t \leq N_z^k \phi_{bzj_{kz}}^k \bar{G}_z^k = \frac{N_z^k \sum_{j \in I_{bz}} \phi_{bzj}^k \bar{\mu}_j^{kz}}{1 + \sum_{i \in I_e} \phi_{ezi}^k \bar{\mu}_i^{ke} + \sum_{j \in I_{bz}} \phi_{bzj}^k \bar{\mu}_j^{kz}}. \quad (\text{D.14})$$

Thus, (D.13)–(D.14) proves that $(\bar{s}, \bar{Y}, \bar{\theta}, \bar{\mu})$ is feasible for model **D-OCPPX**. Moreover, since we have (D.8)–(D.9), then this solution achieves the objective value \bar{V}_2^* . Thus, $\bar{V}_2^* \leq \hat{V}_D^t(x^t)$, which proves the theorem. \square

EC.5. Robust reformulation

Note that if w and p are fixed, then $\Pi(w, p; x^t)$ is the optimal value of a linear program. Hence, by strong duality of LP, $\Pi(w, p; x^t)$ is equivalent to a minimization LP. Thus, the worst-case retailer profit is equivalent to:

$$\begin{aligned} \Pi_{\Gamma, \Delta}(p; x^t) = & \underset{\alpha, \beta, w}{\text{minimize}} && \sum_{k=t}^T \sum_{z \in Z} \left[d_{ez}^k(p) \left(1 + \hat{\delta}_{ez}^k w_{ez}^k \right) \alpha_{ez}^k + d_{bz}^k(p) \left(1 + \hat{\delta}_{bz}^k w_{bz}^k \right) \alpha_{bz}^k \right] + \beta_e x_e^t + \sum_{z \in Z} \beta_{bz} x_{bz}^t \\ & \text{subject to} && \alpha_{bz}^k + \beta_{bz} \geq p_{bz}^k, \quad \forall k \forall z, \\ & && \alpha_{ez}^k + \beta_e \geq p_e^k - c_{ez}, \quad \forall k \forall z, \\ & && \alpha_{ez}^k + \beta_{bz'} \geq p_e^k - c_{z'z}, \quad \forall k, \forall z, \forall z', \\ & && \beta_e \geq q, \beta_{bz} \geq q, \\ & && \alpha \geq 0, \\ & && w \in W(\Gamma, \Delta), \end{aligned}$$

where α, β are the variables in the dual of $\Pi(w, p; x^t)$. Note that the above optimization problem is bilinear. Since w is bounded, then if we can find finite bounds for α , then we can use a “lifting” technique to find a LP lower bound for $\Pi_{\Gamma, \Delta}(p; x^t)$.

LEMMA EC.1. *Let $(\bar{\alpha}, \bar{\beta}, \bar{w})$ be the optimal solution for $\Pi_{\Gamma, \Delta}(p; x^t)$. Then, $\bar{\alpha}_{bz}^k \leq p_{bz}^k - q$ and $\bar{\alpha}_{ez}^k \leq p_e^k - \underline{c}_z - q$ for all $z \in Z$, where $\underline{c}_z = \min(c_{ez}, \min_{z' \in Z} c_{z'z})$*

Proof. Since the demands are nonnegative for all realizations, then given $\bar{\beta}$, $\bar{\alpha}$ must take its smallest feasible value. Thus, $\bar{\alpha}_{bz}^k = \max(0, p_{bz}^k - \bar{\beta}_{bz})$, and $\bar{\alpha}_{ez}^k = \max(0, p_e^k - c_{ez} - \bar{\beta}_e, \max_{z' \in Z} (p_e^k - c_{z'z} - \bar{\beta}_{bz'}))$. Therefore, since the $\bar{\beta}$ variables are bounded below by q , then in fact we have the following upper bounds for the $\bar{\alpha}$ variables: $\bar{\alpha}_{bz}^k \leq p_{bz}^k - q$ and $\bar{\alpha}_{ez}^k \leq p_e^k - \underline{c}_z - q$ for all $z \in Z$. \square

The upper bounds for $\bar{\alpha}$ in the lemma are nonnegative if we impose regularity conditions on Ω :

ASSUMPTION EC.1. *Consider any feasible price vector $\bar{p} \in \Omega$. For all $z \in Z$, $\bar{p}_{bz} \geq q$, and $\bar{p}_e - \min\{c_{ez}, \min_{z' \in Z} c_{z'z}\} \geq q$.*

If the first condition of [Assumption EC.1](#) is not met by \bar{p} , then it is more profitable to hold store inventory than to sell it to a store customer. If the second condition is not met by \bar{p} , then the firm will not sell to an online customer from a zone z , since regardless of the fulfillment location, it is more profitable to hold on to the inventory and sell it at salvage.

Note that α variables are the shadow prices for the demand constraints of $\Pi(w, p; x^t)$, while the β variables are the shadow prices for its inventory constraints. Thus the upper bounds make sense because α_{bz}^k is the marginal increase in value with an additional unit of store demand, which cannot

exceed the marginal value of a store sale. Similarly, α_{ez}^k is the marginal increase in value with an additional unit of online demand, which cannot exceed the maximum marginal value of an online sale (i.e., using the cheapest fulfillment).

From the lemma, we can add constraints $\alpha_{mz}^k \leq A_{mz}^k$ (for some finite A_{mz}^k) that do not change the optimal value of $\Pi_\Gamma(p; x^t)$. Since the w and α variables are bounded, we can use a “lifting” technique to find a linear program whose optimal value $\tilde{\Pi}_{\Gamma, \Delta}(p; x^t)$ is a lower bound to $\Pi_{\Gamma, \Delta}(p; x^t)$:

$$\begin{aligned} \tilde{\Pi}_{\Gamma, \Delta}(p; x^t) = & \underset{\alpha, \beta, w, \nu, \eta}{\text{minimize}} \quad \sum_{k=t}^T \sum_{z \in Z} \left[d_{ez}^k(p) \left(\alpha_{ez}^k + \hat{\delta}_{ez}^k \eta_{ez}^k \right) + d_{bz}^k(p) \left(\alpha_{bz}^k + \hat{\delta}_{bz}^k \eta_{bz}^k \right) \right] + \beta_e x_e^t + \sum_{z \in Z} \beta_{bz} x_{bz}^t \\ & \text{subject to} \quad \alpha_{bz}^k + \beta_{bz} \geq p_{bz}^k, \quad \forall k \forall z, \end{aligned} \quad (\text{D1.1})$$

$$\alpha_{ez}^k + \beta_e \geq p_e^k - c_{ez}, \quad \forall k \forall z, \quad (\text{D1.2})$$

$$\alpha_{ez}^k + \beta_{bz'} \geq p_e^k - c_{z'z}, \quad \forall k, \forall z, \forall z', \quad (\text{D1.3})$$

$$\beta_e \geq q, \quad \beta_{bz} \geq q, \quad (\text{D1.4})$$

$$\alpha \geq 0, \quad (\text{D1.5})$$

$$w \in W(\Gamma, \Delta), \quad (\text{D1.6})$$

$$\eta_{mz}^k + \alpha_{mz}^k \geq 0, \quad (\text{D1.7})$$

$$-\eta_{mz}^k + \alpha_{mz}^k \geq 0, \quad (\text{D1.8})$$

$$A_{mz}^k - A_{mz}^k w_{mz}^k - \alpha_{mz}^k + \eta_{mz}^k \geq 0, \quad (\text{D1.9})$$

$$A_{mz}^k + A_{mz}^k w_{mz}^k - \alpha_{mz}^k - \eta_{mz}^k \geq 0. \quad (\text{D1.10})$$

We “lifted” the nonlinear problem to a higher dimensional space by introducing variables $\eta_{mz}^k = \alpha_{mz}^k w_{mz}^k$, which linearizes the objective. The new constraints (D1.7)–(D1.10) are valid inequalities that are satisfied by any feasible solution to the linearized problem. Note that due to strong duality, the linear program D1 achieves an optimal value equal to the optimal value of its dual:

$$\begin{aligned} \tilde{\Pi}_{\Gamma, \Delta}(p; x^t) = & \underset{s, y, \theta, r, R, g, G, h, H, u, U, l, L, f}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} \left(p_{bz}^k s_{bz}^k + (p_e^k - c_{ez}) y_{ez}^k + \sum_{z'} (p_e^k - c_{z'z}) y_{z'z}^k \right) + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & - \sum_k \Gamma^k f^k - \sum_k \sum_z \sum_{m=e, b} (u_{mz}^k + U_{mz}^k) - \sum_k \sum_z \sum_{m=e, b} A_{mz}^k (g_{mz}^k + G_{mz}^k) \\ & - \sum_k \Delta^k (l^k + L^k) \\ & \text{subject to} \quad s_{bz}^k + r_{bz}^k + R_{bz}^k - g_{bz}^k - G_{bz}^k \leq d_{bz}^k(p), \\ & \quad y_{ez}^k + \sum_{z'} y_{z'z}^k + r_{ez}^k + R_{ez}^k - g_{ez}^k - G_{ez}^k \leq d_{ez}^k(p), \\ & \quad \sum_k \sum_z y_{ez}^k + \theta_e = x_e^t, \end{aligned}$$

$$\begin{aligned}
\sum_k s_{bz}^k + \sum_k \sum_{z'} y_{zz'}^k + \theta_{bz} &= x_{bz}^t, \\
u_{mz}^k - U_{mz}^k + h_{mz}^k - H_{mz}^k - A_{mz}^k (g_{mz}^k - G_{mz}^k) + a_{mz}^t (l^k - L^k) &= 0, \\
-f^k + h_{mz}^k + H_{mz}^k &= 0, \\
r_{mz}^k - R_{mz}^k + g_{mz}^k - G_{mz}^k &= \hat{\delta}_{mz}^k d_{mz}^k(p), \\
s, y, \theta, r, R, g, G, h, H, u, U, l, L &\geq 0.
\end{aligned}$$

Define the following variables: $s_{ez}^k = y_{ez}^k + \sum_{z'} y_{z'z}^k$, $Y_{ez} = \sum_{k=t}^T y_{ez}^k$, and $Y_{z'z} = \sum_{k=t}^T y_{z'z}^k$. Thus, the dual program is equivalent to the following:

$$\begin{aligned}
\tilde{\Pi}_{\Gamma, \Delta}(p; x^t) = \underset{s, Y, \theta, r, R, g, G, h, H, u, U, l, L, f}{\text{maximize}} \quad & \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_z \left(c_{ez} Y_{ez} + \sum_{z'} c_{z'z} Y_{z'z} \right) + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\
& - \sum_k \Gamma^k f^k - \sum_k \sum_z \sum_{m=e, b} (u_{mz}^k + U_{mz}^k) - \sum_k \sum_z \sum_{m=e, b} A_{mz}^k (g_{mz}^k + G_{mz}^k) \\
& - \sum_k \Delta^k (l^k + L^k) \\
\text{subject to} \quad & s_{mz}^k + r_{mz}^k + R_{mz}^k - g_{mz}^k - G_{mz}^k \leq d_{mz}^k(p), \quad m = e, b
\end{aligned} \tag{D2.1}$$

Constraints (D-OCPX.3), (D-OCPX.4), (D-OCPX.5)

$$u_{mz}^k - U_{mz}^k + h_{mz}^k - H_{mz}^k - A_{mz}^k (g_{mz}^k - G_{mz}^k) + a_{mz}^t (l^k - L^k) = 0, \tag{D2.2}$$

$$-f^k + h_{mz}^k + H_{mz}^k = 0, \tag{D2.3}$$

$$r_{mz}^k - R_{mz}^k + g_{mz}^k - G_{mz}^k = \hat{\delta}_{mz}^k d_{mz}^k(p), \tag{D2.4}$$

$$s, Y, \theta, r, R, g, G, h, H, u, U, l, L, f \geq 0. \tag{D2.5}$$

Note that the form of $\tilde{\Pi}_{\Gamma}(p; x^t)$ in model D2 is already a linear program. However, we wish to simplify it further in order to interpret the model.

EC.5.1. Variable transformations

Let us introduce the following variable transformations:

$$\begin{aligned}
\chi_{mz}^k &= u_{mz}^k - U_{mz}^k, & X_{mz}^k &= u_{mz}^k + U_{mz}^k, \\
\lambda_{mz}^k &= g_{mz}^k - G_{mz}^k, & \Lambda_{mz}^k &= g_{mz}^k + G_{mz}^k, \\
\psi_{mz}^k &= r_{mz}^k - R_{mz}^k, & \Psi_{mz}^k &= r_{mz}^k + R_{mz}^k, \\
v_{mz}^k &= h_{mz}^k - H_{mz}^k, & \Upsilon_{mz}^k &= h_{mz}^k + H_{mz}^k, \\
\vartheta^k &= l^k - L^k, & \Theta^k &= l^k + L^k.
\end{aligned}$$

Therefore, by replacing these new variables into model [D2](#), we have its equivalent formulation:

$$\begin{aligned} \tilde{\Pi}_{\Gamma,\Delta}(p; x^t) = & \underset{s,Y,\theta,\chi,X,\lambda,\Lambda,\psi,\Psi,v,\Upsilon,\vartheta,\Theta,f}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_z \left(c_{ez} Y_{ez} + \sum_{z'} c_{z'z} Y_{z'z} \right) + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & - \sum_k (\Gamma^k f^k + \Delta^k \Theta^k) - \sum_k \sum_z \sum_{m=e,b} (X_{mz}^k + A_{mz}^k \Lambda_{mz}^k) \\ \text{subject to} \quad & s_{mz}^k \leq d_{mz}^k(p) + \Lambda_{mz}^k - \Psi_{mz}^k, \quad m = e, b \end{aligned} \quad (\text{D3.1})$$

Constraints [\(D-OCPX.3\)](#), [\(D-OCPX.4\)](#), [\(D-OCPX.5\)](#), [\(D-OCPX.6\)](#)

$$\chi_{mz}^k + v_{mz}^k - A_{mz}^k \lambda_{mz}^k + a_{mz}^t \vartheta^k = 0, \quad (\text{D3.2})$$

$$\Upsilon_{mz}^k = f^k, \quad (\text{D3.3})$$

$$\lambda_{mz}^k + \psi_{mz}^k = \hat{\delta}_{mz}^k d_{mz}^k(p), \quad (\text{D3.4})$$

$$X_{mz}^k \geq |\chi_{mz}^k|, \quad (\text{D3.5})$$

$$\Lambda_{mz}^k \geq |\lambda_{mz}^k|, \quad (\text{D3.6})$$

$$\Psi_{mz}^k \geq |\psi_{mz}^k|, \quad (\text{D3.7})$$

$$\Upsilon_{mz}^k \geq |v_{mz}^k|, \quad (\text{D3.8})$$

$$\Theta^k \geq |\vartheta^k|. \quad (\text{D3.9})$$

Note that since $\Upsilon_{mz}^k = f^k$, then we can eliminate the Υ_{mz}^k variables by replacing the constraint [\(D3.8\)](#) by $f^k \geq |v_{mz}^k|$. By equation [\(D3.2\)](#), we know $v_{mz}^k = A_{mz}^k \lambda_{mz}^k - \chi_{mz}^k - a_{mz}^t \vartheta^k$. Due to the maximizing objective, in the optimal solution, we have $X_{mz}^k = |\chi_{mz}^k|$ and $\Theta^k = |\vartheta^k|$. Next, because the coefficient of s_{mz}^k in the objective is positive and Ψ_{mz}^k reduces the upper bound for s_{mz}^k without impacting the objective, an alternative optimal solution is obtained when Ψ_{mz}^k is reduced and set equal to $|\psi_{mz}^k|$. Lastly, note that the coefficient of Λ_{mz}^k is negative, and it increases the bound on s_{mz}^k through equation [\(D3.1\)](#). However, since A_{mz}^k is an upper bound on the marginal profit for every unit of sale in channel m and zone z , it is optimal to decrease Λ_{mz}^k to the smallest feasible value and incur the least penalty A_{mz}^k . Hence, in the optimal solution $\Lambda_{mz}^k = |\lambda_{mz}^k|$. Hence, we have that model [D1](#) is equivalent to

$$\begin{aligned} \tilde{\Pi}_{\Gamma,\Delta}(p; x^t) = & \underset{s,Y,\theta,\chi,\lambda,\psi,\vartheta,f}{\text{maximize}} \quad \sum_{k=t}^T \sum_{z \in Z} (p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k) - \sum_z \left(c_{ez} Y_{ez} + \sum_{z'} c_{z'z} Y_{z'z} \right) + q \left(\theta_e + \sum_{z \in Z} \theta_{bz} \right) \\ & - \sum_k (\Gamma^k f^k + \Delta^k |\vartheta^k|) - \sum_k \sum_z \sum_{m=e,b} (A_{mz}^k |\lambda_{mz}^k| + |\chi_{mz}^k|) \\ \text{subject to} \quad & s_{mz}^k \leq d_{mz}^k(p) + |\lambda_{mz}^k| - |\psi_{mz}^k|, \quad m = e, b \end{aligned} \quad (\text{D4.1})$$

$$\lambda_{mz}^k + \psi_{mz}^k = \hat{\delta}_{mz}^k d_{mz}^k(p), \quad (\text{D4.2})$$

$$|A_{mz}^k \lambda_{mz}^k - \chi_{mz}^k - a_{mz}^t \vartheta^k| \leq f^k \quad (\text{D4.3})$$

Constraints [\(D-OCPX.3\)](#), [\(D-OCPX.4\)](#), [\(D-OCPX.5\)](#), [\(D-OCPX.6\)](#).

EC.6. Generating parameters for simulations in [Section 4](#)

Below, we describe how we generate parameters and sample paths for a model instance.

EC.6.1. Generating a model instance.

The average week t online and store demand in zone z follows a multinomial logit (MNL) demand model:

$$d_{ez}^t(p_e^t, p_{bz}^t) = N_z^t \times \frac{\exp(\alpha_{ez} - \beta_{ez} p_e^t)}{1 + \exp(\alpha_{ez} - \beta_{ez} p_e^t) + \exp(\alpha_{bz} - \beta_{bz} p_{bz}^t)},$$

$$d_{bz}^t(p_e^t, p_{bz}^t) = N_z^t \times \frac{\exp(\alpha_{bz} - \beta_{bz} p_{bz}^t)}{1 + \exp(\alpha_{ez} - \beta_{ez} p_e^t) + \exp(\alpha_{bz} - \beta_{bz} p_{bz}^t)}.$$

While the starting inventory is the same across zones, the zones are heterogeneous in the MNL demand parameters, including the market size. The time series of zone z weekly average market size parameters, $(N_z^1, N_z^2, \dots, N_z^8)$, is a scaled beta distribution with shape parameters a, b that are randomly chosen such that $1 \leq a \leq b \leq 5$. After generating the market size parameters for all 100 instances and all zones, we observe that the resulting average 8-week market size in a zone is in the range $[67, 90]$. Thus, in all model instances, there is inventory scarcity since the store inventory of 60 units is less than the zone-level market size. To model stochasticity of demand, we assume that the random demand scaling factor ξ_{mz}^t is uniformly distributed in $[1 - \delta_{mz}^t, 1 + \delta_{mz}^t]$ for $m = e, b$. We let δ_{bz}^t (the store demand factors) to be randomly chosen values between 0 to 0.25, while we randomly chose δ_{ez}^t to be between 0 and 1, since we observed from the retailer's data that e-commerce sales are more variable.

While the market size in each zone is nonstationary, we assume that the parameters $(\alpha_{ez}, \alpha_{bz}, \beta_{ez}, \beta_{bz})$ are stationary. In particular, parameters α_{ez} and α_{bz} are randomly chosen values between 9 and 15. The online price sensitivity parameter β_{ez} is randomly chosen in $[0.0375, 0.0625]$, and the store price sensitivity parameter β_{bz} is randomly chosen in $[0.03, 0.05]$. The range of parameters were chosen to represent typical price elasticities that we have observed from the data of the retailer in electronics categories. We let the full price be \$350, and the set of feasible prices for both online and in stores is $\Omega_e = \Omega_b = [\$87.5, \$100, \dots, \$350]$, where the feasible prices are in \$12.50 increments. Note that the smallest feasible price represents a 75% markdown from the full price. We assume that the salvage value is $q = \$35$, which represents a 90% write-off.

To determine the fulfillment costs for the model instance, we generate a 20×20 random distance matrix, where distances are randomly drawn from a uniform distribution over $[0, 2500]$ miles. To ensure that the randomly generated distance matrix satisfies properties such as triangle inequality, we use the Python package `skbio.stats`. We then calculate the ship-from-store fulfillment cost c_{ij} from zone i to zone j using the single-item shipping cost function $c_{ij} = 9.182 + 0.000541 * d_{ij}$ where d_{ij} is the distance (in miles) between i and j . This function is from Section EC.3 of [Jasin](#)

and Sinha (2015) which uses UPS shipping cost data to estimate shipping cost functions through linear regression. The resulting fulfillment cost values are between \$9.18 to \$10.53.

EC.6.2. Generating the sample paths.

For each model instance, we randomly generate 100 sample paths, $\{(\xi_{mz}^{1n}, \xi_{mz}^{2n}, \dots, \xi_{mz}^{8n})\}_{n=1}^{100}$, of the random demand factors for all channels $m = e, b$ and all zones $z \in [Z]$. Factors $\{\xi_{mz}^{tn}\}_{n=1}^{100}$ are drawn from the uniform distribution over $[1 - \delta_{mz}^t, 1 + \delta_{mz}^t]$, with parameter δ_{mz}^t different for each model instance. Given the prices, the demand factor determines the realizations of demand according to the function $D_{mz}^t(p_e^t, p_{bz}^t, \xi_{mz}^t) = \xi_{mz}^t \times d_{mz}^t(p_e^t, p_{bz}^t)$.

EC.7. Demand estimation for Business Value Assessment

We geo-spatially clustered the retailer's stores into 50 zones using a k-means (k=50) algorithm on the store coordinates. Fig. EC.3 shows the 50 zones used for the experiments. We geo-tag all transaction data using zones based on the origin of the demand and the fulfillment location. We ignored the buy-online-pickup-in-store option since we observed few such transactions in the data for the items analyzed. Fig. EC.3 also shows the zonal distribution of sales (the volume is proportional to the pie size) for one of the product category in our data for the retailer. The pie in each zone illustrates the relative frequency of brick-and-mortar sales and e-commerce sales in the zone. Note the heterogeneity of the e-commerce channel share across zones (e.g., 4% to 11%), which can result in certain zones having relatively more ship-from-store activity. Aside from the ability to model geographic-based heterogeneity, another advantage of zone tagging is the ability to tractably capture cross-channel effect (Harsha et al. 2015). We use the zone-tagged data to estimate SKU-zone level demand models described in this section.

We use discrete choice attraction demand function described in Eqs. (C.1–C.2), specifically the MNL function, to model the aggregate consumer behavior across channels. The time series sales data exhibit a distinct product lifecycle (PLC) representing the baseline popularity of a product over its selling season that begins at time t_{start} and has a pre-planned exit date of t_{end} . We estimate the PLC curve by fitting a beta distribution which encompasses a variety of curve shapes, as well as other prediction coefficients using the procedure described later in this section. Model selection and cross-validation on a variety of training instances yielded the following market-size model that predicts the customer arrival rate for any week t in the selling season:

$$\begin{aligned} \log(\text{Market Size}_t) = & \gamma_0 + \gamma_1 \log(1 + t - t_{\text{start}}) + \gamma_2 \log(1 + t_{\text{end}} - t) \\ & + \sum_k \gamma_{3,k} \text{HOLIDAY-VARIABLES}_{k,t}, \end{aligned} \quad (\text{G.1})$$

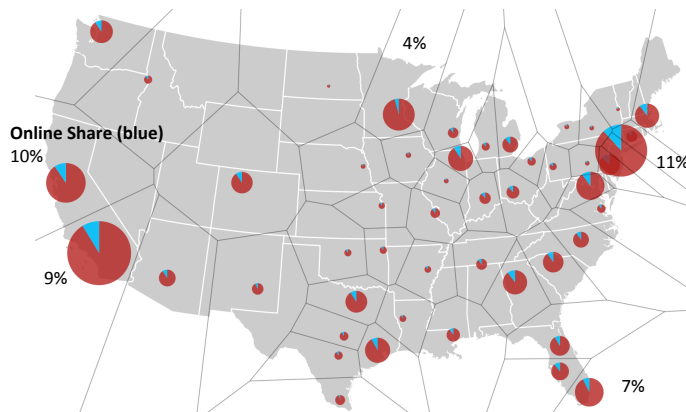


Figure EC.3 Distribution of sales over 50 zones for a product category. Sales volume is proportional to the pie size. The pie in each zone shows the relative frequency of brick-and-mortar sales and e-commerce sales.

and the following market-share model to predict the channel shares in week t :

$$\begin{aligned} \log(\text{Channel Attraction}_t) = & \beta_0 + \beta_1 \text{PRICE}_t + \sum_k \beta_{2,k} \text{PROMOTION-VARIABLES}_{k,t} \quad (\text{G.2}) \\ & + \sum_j \beta_{3,j} \text{COMPETITOR-PRICE-VARIABLES (optional)}_{j,t}. \end{aligned}$$

Holiday spikes, if any, are addressed using holiday indicator variables. It was also beneficial to add channel-specific temporal lag effects prior to the holiday weeks in order to model the spike in online gift orders placed earlier due to the lead time of delivery. Promotional indicators, which include whether the product was advertised that week, were also useful. Competitor prices are introduced as channel-specific attributes, whenever they are available. Future competitor price data are generally not available, but we can use time series methods to forecast competitor prices based on historical trends.

Since the clearance period occurs during the final 10-12 weeks of the product lifecycle, the end-of-life sales decay (measured by decay coefficient γ_2) is a key prediction component for clearance period demand. This decay can occur due to factors such as the waning popularity of the product towards the end of life. Due to the broken assortment effect, this decay may be amplified by inventory depletion, since the item is less visible to store customers. To gauge the incremental impact of low inventory levels on store sales during the markdown period, we experimented with several threshold-based inventory-effect models (Smith and Achabal 1998, Caro and Gallien 2012). However, we did not observe any significant improvement in prediction quality after incorporating such inventory effects, and a PLC-based market-size prediction model was adequate for our application. A review of the in-store display procedures followed by sales associates indicated that the categories we

analyzed were unlikely to be influenced by the broken assortment effect and the ‘store-presentation’ effects. Nevertheless, incorporating inventory effects in an omnichannel environment can be useful for relevant product categories such as fashion apparel.

EC.7.1. Estimation procedure.

The γ and β coefficients in Eqs. (G.1–G.2) are estimated using real historical sales and price data. The goal is to predict future end-of-life sales by channel and location using partial (early and mid-season) TLOG data from the current selling season. There are two challenges.

First, the standard methods to estimate discrete choice models require historical information about every choice, which in our setting, includes censored lost sales. We employ an integrated mixed-integer programming (MIP) approach that jointly estimates market size and the market share parameters in the presence of censored lost sales data proposed by [Subramanian and Harsha \(2017\)](#). Their method performs imputations endogenously in the MIP by estimating optimal values for the probabilities of the unobserved censored choice. Under mild assumptions, they show the method is asymptotically consistent. Besides being a computationally fast single step method, this estimation approach is capable of calibrating market-size covariates (e.g., $\gamma_1, \gamma_2, \gamma_3$), a critical feature with real data. We incorporated model enhancements such as regularization using lasso and ridge penalties and sign constraints on price coefficients to enable an automated demand estimation environment.

The second challenge is in estimating the decay coefficient, γ_2 , for the PLC curve without the full sales history of an item (e.g., future end-of-life sales trajectory can be convex, concave or affine). To overcome this problem, we employed the following two-phase procedure to estimate the parameters of the demand model. In the first phase, the average end-of-life sales decay coefficient γ_2 for a representative SKU from the category was estimated using a learning procedure, and employed as a ‘prior’ desired value in the second phase of estimation that is done at a SKU-zone level for all SKUs. Such priors can also be estimated using historical values of like-SKUs in the same category.

The training sales data was used to estimate parameters of the channel attraction models and the market size model. As we move closer toward the end of the season, and more end-of-life sales data becomes available, the prediction model is recalibrated on a weekly basis using the most recent data, updating all coefficients including the decay coefficient γ_2 with its previous estimate used as a prior, to produce improved sales forecasts for the remaining weeks.

EC.7.2. Prediction accuracy.

We next discuss the achieved model fit and prediction quality using the retailer’s actual sales data and prices. The prediction results presented here is the 12-week look-ahead forecast for the entire clearance period as opposed to rolling horizon weekly sales predictions. We present the




Figure EC.4 Predicted and actual e-commerce sales, and brick-and-mortar sales for a specific SKU-zone.

look-ahead forecast because the OCPX model at the start of the clearance period requires an estimate of demand for all future periods until the planned end date. As time progresses, the demand predictions for the remaining weeks will need to be revised each period. The forecast quality was measured in terms of the volume weighted mean absolute percentage error (WMAPE). We observed this to be largely dependent on the sales rates and hence, the level of disaggregation at which the model calibration was performed, which is consistent with the observations in [Caro and Gallien \(2012\)](#). The achieved out-of-sample WMAPE at the category-chain level was about 22%. This WMAPE value is in close proximity to that observed by [Caro and Gallien \(2012\)](#) who report a 23.8% WMAPE at the category-chain level. At the lowest level of aggregation (SKU-zone level), the average sales rate across the SKUs analyzed during the clearance period (10 per week for brick and 2 for online) was more than 10 times lower than the mid-season sales rate, resulting in the predicted weekly sales deviating from actual sales by ± 5 units for stores and ± 1.2 units for online. [Fig. EC.4](#) is a sample graphical plot of the model fit (for training data) and the look-ahead predictions (for test, the final 12 weeks of sales).

To measure the impact of cross-channel causals to prediction accuracy, we compared the estimated model to a baseline which uses the data to estimate channel demand models without cross-channel causals. We observed the omnichannel demand model reduces SKU-zone level WMAPEs by 1.5 percentage points for the brick channel, and 5 percentage points for the digital channel over the baseline. The incremental gain in prediction accuracy was higher when compared to the partner retailers incumbent single-channel demand forecasting system. Although the benefit of incorporating of cross-channel effects varies by product category, channel price differential, and selling season, for the categories we tested, we observed that the online sales prediction (in general, across categories and across multiple retailers) tends to improve after incorporating cross-channel price and promotion effects. Overall, our demand model and estimated parameters result in prediction qualities consistent with the goals set by the retailer, and was embedded within our proposed optimization framework to calculate optimal prices and inventory partitions.

Table EC.1 Average same-channel and cross-channel price elasticities for Tablets Category.

Channel Sales	Elasticity to brick-and-mortar price	Elasticity to e-commerce price
Brick-and-mortar sales	−1.3	0.7
E-commerce sales	2.8	−3.9

EC.7.3. Estimated price elasticities.

We present the average price same-channel and cross-channel elasticity values evaluated at the average channel price for the Tablets category in [Table EC.1](#). These relatively high elasticity values are typical of markdown settings. Note that the cross-elasticities are asymmetric in that the impact of brick prices on the online sales is different from (and tends to be higher than) the impact of the online prices on brick sales. It is indicative of the heterogeneity of the customers shopping in the different channels as well as the volume share of these channels (the absolute change in volume of brick sales is much higher than that for the online channel).