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 an e-commerce channel and a network of brick-and-mortar stores. An example is cross-channel fulfillment which allows a store to fulfill online orders in any location. Another is price transparency, which allows customers to compare the online price with store prices. This paper studies a new and widespread problem due to omnichannel retail: price optimization in the presence of cross-channel interactions in demand and supply, where cross-channel fulfillment is exogenous. We model the omnichannel pricing problem as a dynamic stochastic program. We propose two pricing policies that are based on the idea of "partitions" to the store inventory that approximate how this shared resource will be utilized. These policies are practical, since they rely on solving computationally tractable mixed integer programs that can accept various business and pricing rules. Moreover, in extensive simulation experiments, they achieve a small optimality gap relative to theoretical upper bounds on the optimal expected profit. The good observed performance of our pricing policies results from managing substitutive channel demands in accordance with partitions that rebalance inventory in the network. A proprietary implementation of the analytics is commercially available as part of the IBM Commerce Markdown Price Solution. The system results in an estimated 13.7% increase in clearance period revenue based on causal model analysis of the data from a pilot implementation for clearance pricing at a large U.S. retailer.

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 shopping experience through an alignment of the retailer's multiple sales channels. The following are a few examples of the new capabilities enabled by omnichannel retail:

- A customer can buy 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 online.
- A customer who purchased a product online might choose a "buy online, pick up in store" option to receive the product sooner than 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. 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., online order 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 an omnichannel environment having many benefits, it also introduces many new challenges for price optimization. Traditional retail pricing models optimize channel prices under the assumption that there is no inventory sharing and coordination between channels. This assumption does not hold in omnichannel where store inventory is additionally used for fulfilling customer orders placed online. 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 traditional revenue management models which assume 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.

This paper studies a widespread and new problem due to omnichannel retail: price optimization on an omnichannel network in the presence of cross-channel interactions in demand and supply, where *cross-channel fulfillment is exogenous and unknown.* This is motivated from a business problem faced by a major U.S. omnichannel retailer, who we engaged with in a partnership with IBM Commerce, a leading provider of merchandising solutions. The partnership aimed to develop an omnichannel clearance pricing optimization system to replace the retailer's legacy system that ignores cross-channel interactions. For this retailer, and for many other client retailers of IBM Commerce, it is difficult to fully coordinate pricing with fulfillment because order fulfillment is managed by a separate and independent operational solution called the order management system (OMS) that fulfills e-commerce orders on-the-fly, while price changes have a slower cadence (e.g., weekly). This problem is different from other models in revenue management literature where prices have to be optimized, either for a dedicated resource, or for an endogenously allocated shared resource. We model the problem as a dynamic stochastic program, and propose two pricing policies that are based on the idea of "partitions" to the store inventory that approximate how this shared resource will be utilized:

- Policy D-OCPX determines inventory partitions by optimizing the profit of a deterministic model where all random variables are replaced with their expected values.
- Policy R-OCPX determines inventory partitions by optimizing the worst-case realized profit in an uncertainty set of the demand realizations.

Here OCPX refers to omnichannel pricing (OCP) with cross-channel (X) interactions. We demonstrate that these pricing policies are:

1. practical since they rely on solving computationally tractable mixed integer programs or MIPs (Lemma 2, Lemma 3) that can be solved by commercial off-the-shelf optimization solvers. In experiments on data from the retailer, the model to price an item across all stores and channels solves within 40 seconds for most problem instances, which is well within the business cycle requirement for markdown optimization. Moreover, various business rules and pricing rules of a retailer can be included as constraints to these models.
2. shown to perform well: (i) in extensive simulation experiments, and (ii) in a pilot implementation of the omnichannel markdown optimization system with the partner retailer.

Deriving a MIP formulation for D-OCPX required proving that a deterministic *nonlinear* optimization problem has the same optimal value as an optimization model with linear objectives and constraints. Deriving the MIP formulation to R-OCPX required finding a tractable lower bound to an NP-hard problem, a two-stage adjustable robust linear program. We prove a lower bound in Proposition 1 by exploiting a structural property of the omnichannel problem.

The good observed performance of our pricing policies results from managing demand substitution across channels in accordance with partitions that rebalance inventory in the network. In particular, the policy encourages online demand, which can be fulfilled from any store, such as stores with slow moving inventory. To demonstrate 2(i), we prove upper bounds on the optimal expected profit (Lemma 1, Lemma 4) with which we benchmark the expected profit of a pricing policy. In simulations, we find that the proposed pricing policies outperform 'deterministic' linear program (DLP) approaches and policies that simulate the legacy pricing system of the retailer.

To demonstrate 2(ii), we conduct a rigorous empirical analysis of data resulting from the pilot implementation of the omnichannel system. We estimate a 13.7% increase in clearance period revenues due to the omnichannel system. Because of the positive feedback from the retailer during the engagement and the widespread nature of the problem, a proprietary implementation of the proposed analytics was commercialized in May 2016 is now available as part of the IBM Commerce Markdown Price Solution.

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 et al. (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). In deriving a robust optimization counterpart to our deterministic problem, because cross-channel flows are recourse decisions after the uncertain demand is realized, the problem turns out to be a two-stage adaptive robust optimization problem, which is known to be NP hard (Zeng and Zhao 2013, Bertsimas and de Ruiter 2016), and for which we develop an approximation.

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. (2017) 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 of short lifecycle products.

2. The omnichannel markdown optimization problem

The partner retailer generates over $40 billion sales annually through an online store and its more than 1,000 brick-and-mortar stores in the U.S. as of 2014. While clearance sales is a small fraction of the retailer's total sales, it is significant in absolute terms (more than $1 billion per year). Table 1 summarizes data for a sample of 195 clearance SKUs in 2014 from the partner retailer. These SKUs contribute to a total of 7.5 million annual clearance sales units, and total annual clearance revenues of $107 million. About 89% of sales made in these categories is through store purchase, and about 11% is through the online channel. Almost 94% of online 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%.

Table 1 Data summary of a sample of 195 clearance SKUs in 2014 from partner retailer.

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%

With omnichannel operations evident from Table 1, using legacy markdown optimization systems that ignored cross-channel interactions result in negative consequences. 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 large initial 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 e-commerce

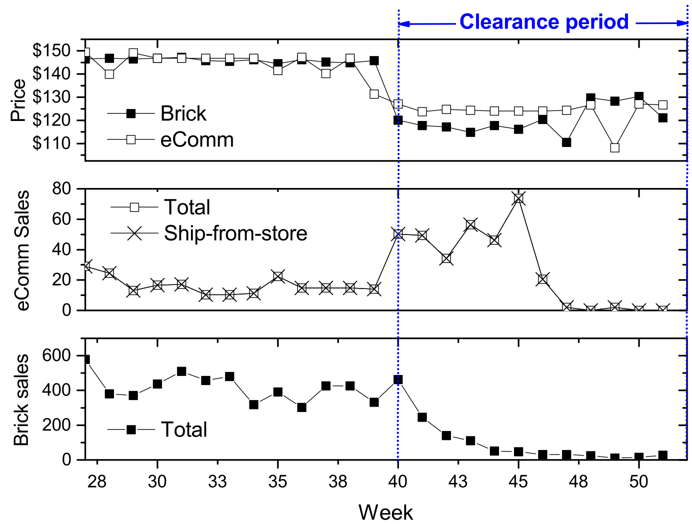

Figure 1 **Time series data for the channel prices and the channel sales of a Tablet SKUs.**

sales are primarily fulfilled by stores during the clearance period (see middle panel). Because of the large store markdowns, after an initial sales spike on Week 40, total brick sales steadily decline due to an increasing number of stores stocking out (see bottom panel). In contrast to the large brick markdowns, the legacy system initially sets no markdown 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 substitution, setting the online price significantly higher than the brick price will cause channel cannibalization. Merchandise managers would override the price outputs from the system to reduce this price 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) is to inject artificial inventory into the online channel. Merchandise managers we interviewed described this manual adjustment process as unmanageable, labor-intensive, and time consuming. Additionally, a manual adjustment process cannot be used for each of the more than 1,000 stores. Hence, a major concern was that since most sales occured through stores, the large brick markdowns would result in systemwide margin erosion.

2.1. Stochastic optimization model

We next model the omnichannel retailer's problem through a dynamic stochastic optimization model. We use the notation $[N]$ to denote the set $\{1, 2, \ldots, N\}$ for any integer $N \geq 1$. Any variables with tilde notation denote random variables. Subscripts e and b denote e-commerce and brick-and-mortar store variables, respectively.

Consider a retailer selling a product with finite inventory over T sales periods. 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 proposed 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 $z \in [Z]$ 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 $z' \neq z$. 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. Inventory that is not sold by period T can be salvaged at a per-unit value q.

The retailer can only change prices at the start of each period. At the start of period $t \in [T]$, let x_e^t be the inventory level in the EFC, and let x_{bz}^t be the inventory available in store $z \in [Z]$. After observing the period t inventory levels in the retail network $x^t = (x_e^t, x_{b1}^t, x_{b2}^t, \ldots, x_{bZ}^t)$, the retailer then determines the network-wide prices. A business rule of the partner retailer is that the online price must be uniform across customer locations, but each individual brick-and-mortar store can set its own price. We denote the vector of network prices as $p^t = (p_e^t, p_{b1}^t, p_{b2}^t \ldots, p_{bZ}^t)$, where p_e^t is the online price and p_{bz}^t is the store price in zone $z \in [Z]$. We let Ω be a discrete set of all 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.

After setting the price, the random channel demands and fulfillments are realized. We denote by $\tilde{D}^t = (\tilde{D}_{ez}^t, \tilde{D}_{bz}^t)_{z\in[Z]}$ the random vector of demands, where $\tilde{D}_{mz}^t = \tilde{D}_{mz}^t(p_e^t, p_{bz}^t)$ is the random channel m demand in zone z, which depends on the online price and the store price in zone z. Note that this demand model allows substitution between the online channel and store channel. We denote $d_{mz}^t(p_e^t, p_{bz}^t) := E[\tilde{D}_{mz}^t(p_e^t, p_{bz}^t)]$ as its expectation. Let $\tilde{y}_{ez}^t$ be a random variable denoting the portion of zone z online demand that has been fulfilled by EFC inventory and $\tilde{y}_{z'z}^t$ be the random portion fulfilled by store inventory from zone $z' \in [Z]$; any unfulfilled portion is lost. Zone z' can either be the same zone as z or a different zone. Suppose that zone z online orders fulfilled from the EFC incur a per-unit fulfillment cost c_{ez}, while those fulfilled from a store in zone z' incurs a per-unit fulfillment cost $c_{z'z}$. The fulfillment costs include shipping costs and handling costs. Define $c = (c_{ez}, c_{1z}, c_{2z}, \ldots, c_{Zz})_{z\in[Z]}$. The fulfillment vector $\tilde{y}^t = (\tilde{y}_{ez}^t, \tilde{y}_{1z}^t, \tilde{y}_{2z}^t, \ldots, \tilde{y}_{Zz}^t)_{z\in[Z]}$ is a random vector since it is realized exogenously (to the model) by the retailer's order management system (OMS), and cannot be jointly optimized with the prices.

Given the period t realization of demand $\tilde{D}^t$ and fulfillment $\tilde{y}^t$, the retailer realizes its period t profit by optimizing recourse sales variables. The channel sales in zone $z \in [Z]$ are denoted by s_{ez}^t and s_{bz}^t. We denote the vector of sales variables as $s^t = (s_{ez}^t, s_{bz}^t)_{z\in Z}$. The state variables for $t+1$ then become $\tilde{x}_e^{t+1} = x_e^t - \sum_{z\in[Z]} \tilde{y}_{ez}^t$ and $\tilde{x}_{bz}^{t+1} = x_{bz}^t - s_{bz}^t - \sum_{z'\in[Z]} \tilde{y}_{zz'}^t$ for all z. Hence, $\tilde{x}^{t+1} = f(x^t, s^t, \tilde{y}^t)$, where f is a linear mapping.

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

$$J^* = \max_{\pi\in\Pi} \quad E^\pi \sum_{t\in[T]} r_t\left(\tilde{x}^t, p^{\pi,t}\right), \tag{2.1}$$

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

$$V_*^t(x^t) = \max_{p^t\in\Omega^{Z+1}} E\left(\max_{s^t\in\mathcal{S}^t(x^t,\tilde{D}^t,\tilde{y}^t)} p^{t\top} s^t - c^\top \tilde{y}^t + V_*^{t+1}\Big(f(x^t, s^t, \tilde{y}^t)\Big)\right), \quad t\in[T], \tag{2.2}$$

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

where V_*^t is the period t value function, e is the vector of all ones, and $\mathcal{S}^t$ is the set of all feasible sales variables given the period t inventory, demand, and fulfillment.

There are several practical challenges for solving the Bellman equations. Firstly, the dynamic programming equations suffer from the well-known 'curse of dimensionality'. Secondly, it is difficult to determine a probability distribution for the random vector $\tilde{y}^t$ since it depends on a separate and independent order management system that determines on-the-fly the fulfilling store or EFC primarily based on operational costs, but may also include complex rules that depend on store performance, store traffic, store fulfillment capacity, minimum store inventory levels, local weather, and order splitting rules. Furthermore, these rules may be frequently adjusted by OMS users.

3. Omnichannel pricing policies

In this section, we propose omnichannel pricing policies that overcome the challenges of solving the omnichannel markdown optimization problem where pricing and fulfillment cannot be fully coordinated. The policies we propose use the idea of *inventory partition* variables. These variables ration the amount of store inventory available for online fulfillment, similar to booking limits in capacity control. Unlike booking limits, however, these partitions cannot be enforced using capacity control (i.e., by accepting or rejecting fulfillment requests). Rather, the policies utilize channel prices in order to adjust demand to meet the inventory partitions.

We were motivated to consider pricing policies based on inventory partitions due to observed deficiencies of the legacy pricing system in an omnichannel environment. Particularly, the legacy system outputs suboptimally large store markdowns because of its assumption that store inventory is only used for store sales, even though in reality it is also used for online fulfillment. Optimizing price markdowns of products used in omnichannel fulfillment require identifying partitions of store inventory units by their usage (e.g. store sales or online fulfillment). However, the true partitions are unknown since they depend on the exogenously determined fulfillment. Directly forecasting such partitions using historical data yields a 70–80% mean absolute percentage error.

To develop more accurate inventory partitions that would be used for the pricing policies, we introduce endogenous fulfillment variables y^t to approximate the exogenous fulfillment $\tilde{y}^t$ (see Section 4.3 later for the resultant lower errors during the pilot). That is, we consider another model in which pricing and fulfillment are coordinated This results in the Bellman equations:

$$V^t(x^t) = \max_{p^t \in \Omega^{Z+1}} E_{\tilde{D}^t} \left(\max_{(s^t, y^t) \in \mathcal{X}^t(x^t, p^t, \tilde{D}^t)} p^{t\top} s^t - c^\top y^t + V^{t+1}\left(f(x^t, s^t, y^t)\right) \right), \; t \in [T], \tag{3.1}$$

$$V^{T+1}(x^{T+1}) = q\left(e^\top x^{T+1}\right), \tag{3.2}$$

where $\mathcal{X}^t(x^t, p^t, \tilde{D}^t)$ is a polyhedral feasible set of period t fulfillment variables and sales variables. Note that if OMS fulfillment satisfies all constraints of $\mathcal{X}^t$, then $V_*^1(x^t) \le V^1(x^t)$ since jointly optimizing prices and fulfillment achieves a higher expected profit. Hence, since V_*^1 cannot be evaluated due to the unknown fulfillment model, we will use V^1 (or upper bounds to V^1) as an objective benchmark with which we compare the omnichannel pricing policies.

The feasible set $\mathcal{X}^t$ can be written as

$$\mathcal{X}^t(x^t, p^t, \tilde{D}^t) = \left\{ s \ge 0, y \ge 0 \;\middle|\; \begin{array}{ll} s^t_{mz} \le \tilde{D}^t_{mz}(p^t_e, p^t_{bz}), & m = e, b, \forall z \in [Z] \\ s^t_{ez} = y^t_{ez} + \sum_{z' \in [Z]} y^t_{z'z}, & \forall z \in [Z], \\ \sum_{z \in [Z]} y^t_{ez} \le x^t_e, & \\ s^t_{bz} + \sum_{z' \in [Z]} y^t_{zz'} \le x^t_{bz}, & \forall z \in [Z], \end{array} \right\} \tag{3.3}$$

The first set of constraints in (3.3) ensures that the sales variables do not exceed the realized demand. The second set of constraints ensures that all online sales are fulfilled. The remaining constraints ensure inventory levels in all periods are nonnegative. If $\tilde{D}^t_{mz}$ is a nonnegative random variable, then starting from nonnegative inventory levels x^1 at $t = 1$, the sequence of feasible sets $\mathcal{X}^t(\tilde{x}^t, \tilde{p}^t, \tilde{D}^t)$ for $t \in [T]$ are nonempty with probability 1.

In this paper, we do not make an assumption on the stochastic model of channel demand $\tilde{D}^t$, since the policies that we propose can be used under any demand model. However, it is worth discussing the multinomial logit (MNL) model of demand that is commonly used in revenue management literature to model substitution among choices. Under this model, if the price online and in-store are p_e and p_b, respectively, then the probability of a zone z arrival purchasing from channel m is:

$$\theta_{mz}(p_e, p_b) = \frac{\exp(\alpha_{mz} - \beta_{mz} p_m)}{1 + \exp(\alpha_{ez} - \beta_{ez} p_e) + \exp(\alpha_{bz} - \beta_{bz} p_b)}, \quad \text{for } m = e, b, \tag{3.4}$$

where $\alpha_{ez}, \alpha_{bz}, \beta_{ez}, \beta_{bz}$ are the MNL parameters.

Lemma 1. *Suppose period t is subdivided into N^t subperiods, where in each subperiod there is no more than one arrival. Let $\frac{N_z^t}{N^t}$ be the probability that there is an arrival in zone z during a subperiod in t, where $\sum_{z\in[Z]} N_z^t \le N^t$. Under a MNL demand model, $V^1(x) \le V_{LP}$, where*

$$
\begin{aligned}
V_{LP} := \underset{\lambda,\mu,y,s}{\text{maximize}} \quad & \sum_{t\in[T]}\sum_{z\in[Z]}\sum_{i\in[I+1]}\sum_{j\in[I+1]} \lambda_{ij}^{tz}\cdot N_z^t\cdot\Big[p_i\theta_{ezij}+p_j\theta_{bzij}\Big] + q\cdot\left(x_e+\sum_{z\in[Z]} x_{bz}\right) && \text{(3.5a)}\\
& -\sum_{t\in[T]}\sum_{z\in[Z]} N^t\Big[qs_{bz}^t+(c_{ez}+q)y_{ez}^t+\sum_{z'\in[Z]}(c_{z'z}+q)y_{z'z}^t\Big] && \text{(3.5b)}\\
\text{subject to} \quad & y_{ez}^t+\sum_{z'\in[Z]} y_{z'z}^t = \sum_{i\in[I+1]}\sum_{j\in[I+1]} \lambda_{ij}^{tz}\cdot\frac{N_z^t}{N^t}\cdot\theta_{ez}(p_i,p_j), \quad \forall t\in[T], z\in[Z], && \text{(3.5c)}\\
& s_{bz}^t = \sum_{i\in[I+1]}\sum_{j\in[I+1]} \lambda_{ij}^{tz}\cdot\frac{N_z^t}{N^t}\cdot\theta_{bz}(p_i,p_j), \quad \forall t\in[T], z\in[Z], && \text{(3.5d)}\\
& \sum_{t\in[T]}\sum_{z\in[Z]} N^t y_{ez}^t \le x_e, && \text{(3.5e)}\\
& \sum_{t\in[T]}\sum_{z'\in[Z]} N^t y_{zz'}^t + \sum_{t\in[T]} N^t y_{bz}^t \le x_{bz}, \quad \forall z\in[Z], && \text{(3.5f)}\\
& \sum_{j\in[I+1]} \lambda_{ij}^{tz} = \mu_i^{te}, \quad \forall t\in[T], z\in[Z], i\in[I+1], && \text{(3.5g)}\\
& \sum_{i\in[I+1]}\sum_{j\in[I+1]} \lambda_{ij}^{tz} = 1, \quad \forall t\in[T], z\in[Z], && \text{(3.5h)}\\
& \sum_{i\in[I+1]} \mu_i^{te} = 1, \quad \forall t\in[T], && \text{(3.5i)}\\
& \lambda\ge 0,\ \mu\ge 0,\ y\ge 0,\ s\ge 0 && \text{(3.5j)}
\end{aligned}
$$

where $p_{I+1}=\infty$, and $\theta_{mzij}=\theta_{mz}(p_i,p_j)$.

While introducing endogenous fulfillment variables aids in identifying inventory partitions, solving the Bellman equations (3.1)–(3.2) suffers from the curse of dimensionality so the optimal pricing policy cannot be solved in practice. Lemma 1 provides an upper bound to the optimal expected profit, which we will use in later experiments to evaluate the optimality gap of different pricing policies under MNL demand. The proof of the lemma is in the appendix. The idea behind the proof is that $V^1(x)$ is bounded by the optimal expected profit if the firm can dynamically price and fulfill orders for *each* arriving customer. For this latter problem, the optimal profit is bounded by the optimal value of the deterministic linear program (3.5), with the proof following similar lines as Lemma 1 of Lei et al. (2017), but with time-inhomogenous demand rates.

Before we introduce our pricing policies, we first discuss why new heuristics had to be developed to address the practical constraints existing in omnichannel revenue management. A popular approach in revenue management is the use of deterministic and fluid approximation models, in which stochastic quantities are replaced by their mean values and capacity and demand are assumed to be continuous (Gallego and van Ryzin 1997, Liu and van Ryzin 2008, Jasin and Kumar 2013). The 'deterministic' linear program (DLP) approach will take the fractional solution to (3.5) and use these as probabilities to randomly determine the prices *and* fulfillment for each new arrival. As the demand and capacity scales up, this approach is known to converge to the optimal profit under the setting where the firm can coordinate prices and fulfillment for each customer (Lei et al. 2017). However, as we demonstrate in computational experiments later, these policies are suboptimal under the omnichannel setting where the price has to be the same for all arrivals in period t *and* fulfillment is determined exogenously. Hence, we needed to develop a new heuristic that performs well under an exogenous fulfillment.

3.1. Deterministic omnichannel pricing policy (D-OCPX)

To develop our pricing policy, we approximate the value function V^t that is difficult to compute with the optimal value of a computationally tractable optimization model. Since the optimal inventory partitions not only depend on the fulfillment costs and salvage cost, but also on the future prices in the whole retail network, we approximate V^t by the optimal value V_D^t of a deterministic multiperiod optimization model that jointly optimizes prices $(p^t, p^{t+1}, \ldots, p^T)$ and optimal inventory partitions. In particular,

$$V_D^t(x^t) := \underset{p,s,y,u}{\text{maximize}} \quad \sum_{k=t}^{T} \sum_{z\in[Z]} \left(p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k\right) - \sum_{z\in[Z]} c_{ez} y_{ez} - \sum_{z\in[Z]} \sum_{z'\in[Z]} c_{zz'} y_{zz'} + q\left(u_e + \sum_{z\in[Z]} u_{bz}\right) \tag{3.6a}$$

$$\text{subject to} \quad s_{mz}^k \le d_{mz}^k(p_e^k, p_{bz}^k), \qquad k = t, \ldots, T,\ m = e, b,\ \forall z \in [Z], \tag{3.6b}$$

$$\sum_{z\in[Z]} y_{ez} + u_e = x_e^t, \tag{3.6c}$$

$$\sum_{k=t}^{T} s_{bz}^k + \sum_{z'\in[Z]} y_{zz'} + u_{bz} = x_{bz}^t, \quad \forall z \in [Z], \tag{3.6d}$$

$$\sum_{k=t}^{T} s_{ez}^k = y_{ez} + \sum_{z'\in[Z]} y_{z'z}, \qquad \forall z \in [Z], \tag{3.6e}$$

$$s \ge 0,\ y \ge 0,\ u \ge 0, \tag{3.6f}$$

$$p^k \in \Omega^{Z+1} \qquad k = t, \ldots, T. \tag{3.6g}$$

Variables y_{ez} and $y_{zz'}$ are the *inventory partitions* that determine the channel-zone use of EFC inventory and store inventory. Note that replacing all random demands in (3.1) with their expected values and by setting $y_{ez} = \sum_{k=t}^{T} y_{ez}^k$ and $y_{zz'} = \sum_{k=t}^{T} y_{zz'}^k$ results in (3.6). Hence, V_D^t is the "certainty equivalence" formulation of (3.1), a standard technique in stochastic control (Bertsekas 1995).

Algorithm 1 Pricing policy D-OCPX

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

$t \leftarrow 1$
while $t \leq T$ **do**
    Solve (3.6) for the optimal solution (p^*, s^*, y^*, u^*)
    $p_e^t \leftarrow p_e^{t*}$
    $p_{bz}^t \leftarrow p_{bz}^{t^*}$ for $z \in [Z]$
    Exogenously realize demands $D^t(p^t, \xi^t)$, sales s^t, and fulfillments y^t
    $x^{t+1} \leftarrow f(x^t, y^t, s^t)$
    $t \leftarrow t+1$

Algorithm 1 is a pricing policy, which we refer to as Deterministic OCPX (D-OCPX), based on these deterministic store inventory partitions. In each period, the pricing policy solves the optimization model (3.6) and sets the price according to its optimal solution. The optimality of D-OCPX under exogenous fulfillment depends on how accurately the fulfillment variables y^t approximate the actual fulfillment vectors $\tilde{y}^t$ from the OMS. In Section 4.3, we report the accuracy of this prediction in the commercial pilot implementation.

We next discuss the computational tractability of D-OCPX. In each period t, the algorithm solves model (3.6) which is nonconvex since the objective function is bilinear and the feasible set could be nonconvex (for instance, it is nonconvex if the demand function is MNL). However, since the feasible price set is discrete, the following lemma shows that this optimization model can be reformulated as a tractable mixed integer linear program.

Lemma 2. *$V_D^t(x^t)$ is equal to the optimal value of a mixed integer program with $O(I^2)$ binary variables. Under a multinomial logit demand model, the number of binary variables is $O(I)$.*

Lemma 2 states that the price optimization problem (3.6) can be solved as a MIP with $O(I^2)$ binary decision variables. The proof of the lemma can be found in the electronic companion. A key step in the reformulation is to introduce binary decision variables corresponding to the feasible prices in the discrete set Ω. The binary variable μ_i^{ke} is equal to 1 if and only if the online price at time $k \in [T]$ is p_i where $i \in [I]$. The binary variable μ_j^{kz} is 1 if and only if the store price in zone $z \in [Z]$ at time $k \in [T]$ is set to p_j where $j \in [I]$. Therefore, instead of solving (3.6), D-OCPX can instead solve a tractable MIP in each period t and set $p_e^t \leftarrow \sum_{i\in[I]} \mu_i^{te} p_i$ and $p_{bz}^t \leftarrow \sum_{i\in[I]} \mu_i^{tz} p_i$ for all $z \in [Z]$. The lemma also states that the number of binary variables of (A.11) reduces to $O(I)$ under the special case of a multinomial logit model of demand. The proof requires the use of the reformulation-linearization technique (Sherali and Adams 1999) and Charnes and Cooper (1962) transformations. In our computational experiments later in Section 4, when implementing this model in CPLEX for realistic problem sizes, the solver with its default termination criteria solves most instances in less than 40 seconds (see Fig. EC.1 in the e-companion).

We next demonstrate the performance of D-OCPX in simulations under the arrival process described in Lemma 1. The feasible price set is between \$125 to \$500 (with a stepsize of \$12.50), and the salvage value is \$50. We set $Z = 5$ and $T = 5$. The distance between zones is randomly chosen, which determines the fulfillment costs. To test the pricing policies, we chose high fulfillment costs (\$46.70 for in-zone, \$50.70–\$71.20 for cross-zone). We generated parameters for (3.4) and N_z^t according to the method we describe later in Section 5.1, and assume zero probability of no arrivals in a subperiod. We set $x_e = 0$ and, for all $z \in [Z]$, set $x_{bz} = 0.5 \times \sum_{t\in[T]} N_z^t$ so that the total inventory is less than the expected number of arrivals in the zone. The baseline total arrivals $N := \sum_{t\in[T]} N^t$ is 50. In different runs of the experiment, we report the performance of D-OCPX as the problem size increases by scaling the number of arrivals and the total inventory to σN and σx, respectively, where $\sigma \in \{1, 2, 4, 8, 16, 20, 24, 34\}$ is a scaling parameter.

At the beginning of period $t \in [T]$, the simulation sets the prices $p^t = (p_e^t, p_{b1}^t, \ldots, p_{b5}^t)$ based on the solution of the D-OCPX model (3.6). During period t, a total of $N^t = \sum_{z\in[Z]} N_z^t$ customers arrive one at a time. The probability of an arrival originating from zone z is $\frac{N_z^t}{N^t}$. Given a zone z arrival, the probability of this customer choosing to purchase from the store and from online are $\theta_{bz}(p_e^t, p_{bz}^t)$ and $\theta_{ez}(p_e^t, p_{bz}^t)$, respectively. If the customer chooses to purchase from the store, the sale is fulfilled using store z inventory, if available. Otherwise, the sale is lost. If the customer chooses to purchase online, the sale is fulfilled by the store with lowest fulfillment cost out of all stores with positive inventory. For each experiment run, we estimate the expected profit by averaging the profit of D-OCPX on 10^4 randomly generated sample paths, where each sample path is a different sequence of the N customer arrivals. We compare the D-OCPX sample average profit to the LP upper bound (3.5). We also compare D-OCPX to the following three randomized policies that use the deterministic linear program (DLP) solution $(\lambda^*, \mu^*, y^*, s^*)$ to (3.5):

- Policy DLP-FP-LF (Fixed Price - Low Fulfill) uses $(\mu_i^{te*})_{i\in[I]}$ as probabilities to draw a random online price to set for all customers arriving in time t. Given the online price p_i, the policy then uses the probabilities $(\lambda_{ij}^{tz*})_{j\in[I]}$ to draw the store price to be set for all arriving customers in zone z during time t. Fulfillment for each online sale uses the lowest fulfillment cost rule.
- Policy DLP-RP-LF (Random Price - Low Fulfill) uses (μ^{t*}, λ^{t*}) to draw a *new* random online price and random store price for *each* arriving customer in time t. Fulfillment for each online sale uses the lowest fulfillment cost rule.
- Policy DLP-RP-RF (Random Price - Random Fulfill) is a joint pricing and fulfillment policy for each new arriving customer, similar to Lei et al. (2017). It uses (μ^{t*}, λ^{t*}) to draw a *new* random online price and random store price for *each* arriving customer in time t. It chooses the fulfillment node based on y^{t*}. Given an online sale in zone z during time t, it randomly chooses the fulfilling zone based on the probabilities $(y_{z'z}^{t*})_{z'\in[Z]}$.

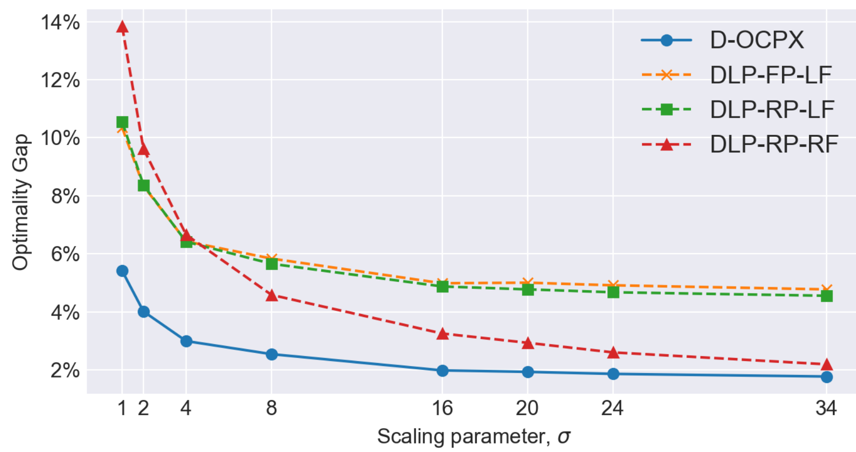

Figure 2 **Optimality gap of D-OCPX, DLP-FP-LF, DLP-RP-LF, and DLP-RP-RF plotted against the problem size scale. The optimality gap is the relative difference between the expected profit and V_{LP}.**

DLP-RP-LF and DLP-RP-RF are not feasible policies for (2.2) since stores cannot change prices for each new customer and, for the latter policy, fulfillment is not exogenous. Since their expected profits are also bounded above by V_{LP}, we test these policies in the simulation as benchmarks. Fig. 2 reports the optimality gap of the pricing policies plotted against the problem size scale parameter σ. A well-known property of a 'deterministic' linear program approach is that its solution is asymptotically optimal to stochastic revenue management problem as demand and capacity are scaled up (see for example Cooper 2002, Liu and van Ryzin 2008, Lei et al. 2017). We observe this property on DLP-RP-RF, with an empirical convergence closely resembling the $O(\frac{1}{\sqrt{\sigma}})$ bound by Lei et al. (2017). On the other hand, when $\sigma = 34$, DLP-FP-LF and DLP-RP-LF have optimality gap of 4.78% and 4.56%, respectively. In these two policies, fulfillment is determined exogenously, instead of using the LP solution to randomize fulfillment, so they do not achieve the desirable convergence properties of DLP-RP-RF. Of all the policies, D-OCPX has the smallest optimality gap even for small problem sizes. Hence, D-OCPX is suited for clearance pricing problems which are characterized by low inventory and low demand.

3.2. Robust omnichannel pricing (R-OCPX)

Pricing policy D-OCPX approximates the value function by assuming the stochastic demand would equal its expected value. That is, given the omnichannel prices, the inventory partition variables y_{ez} and $y_{zz'}$ in (3.6) are aggregate fulfillment quantities that minimize the fulfillment cost of the *deterministic* demand. In the experiments with multinomial demand, D-OCPX results in expected profits with small optimality gaps. In general, under demand distributions with higher variability, D-OCPX could result in a higher optimality gap. This is because the partition variables may be poor approximations of the optimal fulfillment quantities, since ignoring uncertainty results in fulfillment decisions that incur significant lost in-store sales and fulfillment costs under higher than expected demand realizations or low revenue under lower than expected demand realizations.

We will next enhance the proposed pricing policy D-OCPX with inventory partitions that incorporate uncertainty through a robust optimization approach. The robust formulation assumes that after the firm chooses the omnichannel prices, an "adversary" chooses demand realizations from a specified *uncertainty set* which results in the worst-case firm profit. Knowing the adversary's strategy, the firm then chooses the prices to maximize its worst-case profit. Note that the prices chosen by the firm can be interpreted as a Stackelberg-type game equilibrium solution, with the firm as the leader and the adversary as the follower. The advantage of the robust approach is that it does not suffer from the curse of dimensionality inherent in a stochastic dynamic programming approach. Moreover, the optimal value of the robust model is a lower bound on the firm's profit for any demand realization in the uncertainty set.

We next describe the demand uncertainty set. We assume a multiplicative uncertainty, where $\tilde{D}^t_{mz}(p^t_e, p^t_{bz}) = \tilde{\xi}^t_{mz} \times d^t_{mz}(p^t_e, p^t_{bz})$ for some nonnegative random variable $\tilde{\xi}^t_{mz}$ with mean 1. Suppose that $\tilde{\xi}^t_{mz} \in [1-\delta^t_{mz}, 1+\delta^t_{mz}]$ for some $\delta^t_{mz} \le 1$. We can then express any of its realizations as $\xi^t_{mz} = 1 + \delta^t_{mz} w^t_{mz}$ for some $w^t_{mz} \in [-1, 1]$. Hence w^t_{mz} determines the realization of the random demand $\tilde{D}^t_{mz}$, and it is sufficient to define an uncertainty set for w. If this uncertainty set allows each w^t_{mz} variable to take values independently, the adversary will choose $w^t_{mz} = -1$ for all periods, zones and channels, resulting in the robust solution setting large markdowns to reduce the number of unsold units. Thus we include several other coupling constraints in the uncertainty set of w:

$$W^t_{\Gamma,\Delta} := \left\{ w \;\middle|\; \begin{array}{ll} -1 \le w^k_{mz} \le 1, & k = t, \dots, T, \; \forall z \in [Z], m = e, b, \\ \sum_{m=e,b} \sum_{z \in [Z]} |w^k_{mz}| \le \Gamma^k, & k = 1, \dots, T \\ \left| \sum_{m=e,b} \sum_{z \in [Z]} a^k_{mz} w^k_{mz} \right| \le \Delta^k, & k = 1, \dots, T \end{array} \right\} \tag{3.7}$$

where $\Gamma, \Delta \ge 0$ are *budgets of uncertainty.* Γ limits the average absolute percentage deviation of channel-zone demand. Motivated from commonly observed behavior that aggregate chain-level demand has a lower forecast error than the zone-level demand, Δ limits the aggregate chain-level demand percentage deviation. Scalars a^t_{mz} are normalization constants. A choice of small Γ shrinks the uncertainty set; a choice of small Δ permits only a small forecast error in the aggregate chain-level demand while allowing larger zone-level demand errors. To set the parameters for (3.7), one can set δ^t_{mz} as the absolute percentage error of the channel m zone z demand forecast in previous periods, and Δ^t as the error of the chain-level demand forecast. The constants a^t_{mz} can be set as $\delta^t_{mz} \bar{d}^t_{mz} / \sum_{mz} \bar{d}^t_{mz}$ (i.e., demands normalized by δ^t_{mz}), where $\bar{d}^t_{mz}$ are the calibrated demands at observed prices.

Under a robust framework, after the retailer chooses prices $p = (p^t, \ldots, p^T)$, an adversary sets the demand realizations by choosing $w \in W^t_{\Gamma,\Delta}$ which results in the worst-case profit. We denote the worst-case profit during period t to T resulting from setting price vector p as

$$U_R^t(p, x^t) := \underset{w \in W^t_{\Gamma,\Delta}}{\text{minimize}} \quad U^t(w, p; x^t) \tag{3.8}$$

where $U^t(w, p; x^t)$ is the profit during period t to T under realization w given the initial inventory level x^t. Faced with an adversary, the retailer's best strategy is to choose the price vector p that maximizes $U_R^t(p; x^t)$. Thus, the computational complexity of the robust approach depends on finding the maximizer of $U_R^t(p; x^t)$ efficiently. However, (3.8) is NP-hard, as we discuss next.

In the sequence of the robust model, the fulfillment vector y and the sales vector s are chosen after the adversary chooses a demand realization. Hence, the profit realization $U^t(w, p; x^t)$ is the optimal value of following optimization model:

$$U^t(w, p; x^t) = \underset{s,y,u}{\text{maximize}} \quad \sum_{k=t}^{T} \sum_{z \in [Z]} \left(p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k\right) - \sum_{z \in [Z]} \left(c_{ez} y_{ez} + \sum_{z' \in [Z]} c_{zz'} y_{zz'} \right) + q \left(u_e + \sum_{z \in [Z]} u_{bz} \right) \tag{3.9a}$$

$$\begin{aligned}
\text{subject to} \quad & \text{Constraints (3.6c) – (3.6f)}, \\
& s_{ez}^k \le d_{ez}^t(p_e^t, p_{bz}^t)(1 + \delta_{ez}^t w_{ez}^t), \quad k = t, \ldots, T,\ z \in [Z], && (3.9b) \\
& s_{bz}^k \le d_{bz}^t(p_e^t, p_{bz}^t)(1 + \delta_{bz}^t w_{bz}^t), \quad k = t, \ldots, T,\ z \in [Z] && (3.9c)
\end{aligned}$$

Note that $U^t(w, p; x^t)$ is the optimal value of a linear program. Hence, by LP strong duality, we have that $U^t(w, p; x^t)$ is equivalent to its dual (a minimization LP), and consequently $U_R^t(p; x^t)$ is the optimal value of a minimization problem with decision variables w and the dual variables of (3.9). However, this minimization problem is bilinear, which in general is NP-hard.

To overcome the computational challenge of the robust model, we next develop a tractable linear program which gives a lower bound on the worst-case profit $U_R^t(p; x^t)$ under the following regularity condition on the feasible price set Ω.

Assumption 1. *For any $p \in \Omega$, $p \ge q + c_z^{\min}$ $\forall z \in [Z]$, where $c_z^{\min} := \min\left\{c_{ez}, \min_{z' \in [Z]} c_{z'z}\right\}$.*

If a price p violates the condition of Assumption 1 for some $z \in [Z]$, then the firm will not sell to any online customer from zone z since, regardless of the fulfillment location, it is more profitable to hold on to the inventory and sell it at salvage value. The assumption also implies that $p \ge q$, hence it is not profitable to hold store inventory than sell to a store customer.

Defining $p^{\max} := \max\{p : p \in \Omega\}$ as the maximum price, let us introduce the following parameters: $A_{bz}^k := p^{\max} - q$ and $A_{ez}^k := p^{\max} - c_z^{\min} - q$ where $k = t, \ldots, T$ and $z \in [Z]$. Then the following lemma gives a lower bound for U_R^t.

Algorithm 2 Pricing policy R-OCPX

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

$t \leftarrow 1$
while $t \le T$ **do**
    Solve (3.6) for the optimal solution (p^*, s^*, y^*, u^*)
    $\hat{p} \leftarrow \arg\max_{p \in \Omega_Z^{t,T}(p^*)} U_l^t(p; x^t)$
    $p_e^t \leftarrow \hat{p}_e^t$
    $p_{bz}^t \leftarrow \hat{p}_{bz}^t$ for $z \in [Z]$
    Exogenously realize demands $D^t(p^t, \xi^t)$, sales s^t, and fulfillments y^t
    $x^{t+1} \leftarrow f(x^t, y^t, s^t)$
    $t \leftarrow t+1$

Proposition 1. Under Assumption 1, $U_R^t(p; x^t) \ge U_l^t(p; x^t)$ for any $p = (p^t, \ldots, p^T)$, where

$$
\begin{aligned}
U_l^t(p;x^t) = \underset{s,y,u,\chi,\upsilon,\psi,\vartheta,f}{\text{maximize}} \quad & \sum_{k=t}^{T} \sum_{z\in[Z]} \left(p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k\right) - \sum_{z\in[Z]} \left(c_{ez} y_{ez} + \sum_{z'\in[Z]} c_{z'z} y_{z'z} \right) + q \left(u_e + \sum_{z\in[Z]} u_{bz} \right) && \text{(3.10a)} \\
& - \sum_{k=t}^{T} \left(\Gamma^k f^k + \Delta^k |\vartheta^k|\right) - \sum_{k=t}^{T} \sum_{z\in[Z]} \sum_{m=e,b} \left(A_{mz}^k |\upsilon_{mz}^k| + |\chi_{mz}^k|\right) && \text{(3.10b)} \\
\text{subject to} \quad & s_{mz}^k \le d_{mz}^k(p) + |\upsilon_{mz}^k| - |\psi_{mz}^k|, \quad m = e,b,\ k = t,\ldots,T,\ \forall z \in [Z], && \text{(3.10c)} \\
& \upsilon_{mz}^k + \psi_{mz}^k = \delta_{mz}^k d_{mz}^k(p_e^k, p_{bz}^k), \quad m = e,b,\ k = t,\ldots,T,\ \forall z \in [Z], && \text{(3.10d)} \\
& \left| A_{mz}^k \upsilon_{mz}^k - \chi_{mz}^k - a_{mz}^k \vartheta^k \right| \le f^k, \quad m = e,b,\ k = t,\ldots,T,\ \forall z \in [Z], && \text{(3.10e)} \\
& \text{Constraints (3.6c)–(3.6f)}
\end{aligned}
$$

The proof of Proposition 1 can be found in the e-companion. While (3.10) is not a linear program, the objective and the constraints (3.10c) and (3.10e) can be linearized with standard techniques.

Note that (3.9) maximizes profit under a given demand realization through the sales and inventory partition variables. On the other hand, (3.10) maximizes profit with penalty terms through these same variables and auxiliary decision variables. From constraint (3.10d), it is easy to deduce that $|\upsilon_{mz}^k| - |\psi_{mz}^k|$ takes values between the possible demand deviations from $-\delta_{mz}^k d_{mz}^k$ to $+\delta_{mz}^k d_{mz}^k$. Therefore, $U_l^t(p; x^t)$ allows sales, partition variables, and demand realizations to be determined jointly, but a penalty is incurred based on the demand realization.

We propose Algorithm 2 which we refer to as Robust OCPX (R-OCPX). In each period, the policy finds prices that maximize the lower bound on the worst-case firm profit $U_l^t(p; x^t)$, i.e.,

$$
V_R^t(x^t) := \max_{p=(p^t,\ldots,p^T)\in\Omega_Z^{t,T}} U_l^t(p; x^t), \tag{3.11}
$$

where $\Omega_Z^{t,T}$ is a feasible price set for $p = (p^t, \ldots, p^T)$. If $\Omega_Z^{t,T}$ includes all feasible prices allowed in Ω, we observe that the resulting pricing policy is too conservative. Hence, we restrict the allowable prices that R-OCPX can choose. In particular, at each period t, R-OCPX first solves (3.6) for the optimal prices $p^* = (p^{t*}, \ldots, p^{T^*})$. Then $\Omega_Z^{t,T}$ are all the feasible prices that are some neighborhood of p^*. A motivation for this is that in practical large-scale combinatorial optimization instances such as D-OCPX, there tend to be a large pool of (near) alternative optimal solutions. Hence R-OCPX can preserve the overall D-OCPX expected margin, while choosing the maximally robust price vector among these alternatives.

Note that even though U_l^t is the optimal value of a linear program, optimization model (3.11) is nonlinear. However, it is possible to linearize the problem for general demand functions by adding binary variables for each feasible price, the same technique used to derive (A.11). Moreover, under a multinomial logit demand model, the number of binary variables can be reduced. We omit the proof since it follows along the same lines as the proof of Lemma 2.

Lemma 3. *$V_R^t(x^t)$ is equal to the optimal value of a mixed integer program with $O(I^2)$ binary variables. Under a multinomial logit demand model, the number of binary variables is $O(I)$.*

3.3. Applicability to other omnichannel retailers

Markdown (or clearance pricing) optimization is a widely used pricing solution to boost profitability while making room for the next seasons inventory. In developing the pricing policies, we made several modeling assumptions to tailor the model to the partner retailer's specific business requirements. While some features are specific to the retailer, the business problem that we described in Section 2 (i.e., inventory is used for multiple channels, but existing pricing systems ignore cross-channel interactions) exists for many major omnichannel retailers. This can lead to reduced sales in one channel and margin erosion in the other, and, depending on the retailer, one can be more severe than the other. To offer new solutions for omnichannel retail, IBM Commerce commercialized the omnichannel markdown pricing solution that includes a proprietary implementation of the models proposed in this paper. Below we describe how the models described in this paper can be tailored to meet the idiosyncratic requirements of other retailers.

In our work with other major retailers, we have observed different implementations of omnichannel pricing. The partner retailer allows each store to set its own markdowns, but only allows a single e-commerce markdown. Other retailers, on the other hand, may choose to set a single price regardless of the channel or location. These omnichannel business constraints on prices can be easily modeled through inequalities on prices or binary variables.

Omnichannel retailers may also differ in their fulfillment methods. Fulfillment has to be accurately modeled by the inventory partition variables, which can be achieved through modifications to the model. For instance, a retailer could have multiple e-fulfillment centers, which can be modeled by adding new supply nodes. Practically motivated constraints on fulfillment, such as ship-from-store (SFS) capacities per store can also be easily encoded in the model to better approximate fulfillment quantities. Other retailers have a 'buy online pickup in store' option for their online customers. This can be modeled by assuming this option will be chosen by a known proportion of online customers in a zone, then modifying the constraints (3.6d)–(3.6e) appropriately.

Mid-season replenishment of store inventory can be modeled by introducing flow variables between warehouse and stores. Some retailers may also use store transshipment to rebalance inventory between stores during the sales horizon. This can be modeled with store-to-store continuous flow variables similar to y. Transshipment complements SFS and aims to avoid lost sales in the store channel whereas SFS avoids lost sales in the e-commerce channel. However, it is worth mentioning that unlike SFS, transshipments cannot satisfy immediate inventory shortages. This is due to the transshipment leadtime and that, unlike online sales that are realized as soon as a package is shipped, out-of-stock stores lose sales even when the transshipment is in the pipeline.

The model can also be modified for single channel retailers to price shared inventories (e.g., a retailer selling the same item across different websites that have independent markdowns). The model can also be used to generate displacement costs (shadow prices) of a unit inventory which can be used by OMS to optimize sourcing decisions and to coordinate pricing and fulfillment (Ettl et al. 2017).

4. Pilot implementation at a large U.S. retailer

We next describe the development of an omnichannel price analytics system (part of the IBM Commerce markdown optimization solution). The project proceeded in three phases: (i) a business value assessment on historical data, (ii) the development of the omnichannel system, and (iii) a commercial release and pilot implementation for clearance pricing at a major U.S. retailer.

4.1. Business value assessment

We engaged with the retailer for 6 months to conduct the business value assessment (BVA) to understand the impact of D-OCPX 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 195 clearance SKUs in three product categories (Notebooks, Tablets, and Tablet Accessories) between January 1, 2014 to December 31, 2014 (see Table 1 for the data summary). We were additionally provided with price information for 18 online competitors. This competitor 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

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

	Store price		Online price		% Change in Sales			Change in
Category	Actual	D-OCPX	Actual	D-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)

use this data to calibrate the zone-level demand models of each SKU using a procedure we describe in Section EC.2 of the e-companion.

Using the estimated zone-level demand models of each SKU, we optimize the 10–12 week clearance prices through the MIP formulation of (3.6) with $t = 1$. The retailer's fulfillment strategy was to use the EFC to fulfill the online orders if it had inventory; otherwise, the nearest in-stock store was used, unless there were other operational costs and business rules. Fulfillment cost data was not available, hence we set distance-based shipping costs, with additional cost for ship-from-store (SFS) fulfillment to adhere to priority ordering. For each SKU, the retailer provided the salvage values. We included in (3.6) pricing constraints based on SKU-specific rules provided by the retailer, such as a minimum time between markdowns, minimum and maximum markdown price percentages, and price bounds. Model (3.6) 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 default termination criteria was used to solve the MIPs. Model (3.6) was solved with up to 20 price discretizations. Problem instances had up to 10K binary variables, 50K flow variables, and 100K constraints. In comparison, the pre-pilot approach was to solve price optimization models for each channel and store(s) separately, which are significantly smaller in size. Fig. EC.1 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. Thus, our MIP-based approach allows us to solve the large scale problem within the required business cycle.

Since the "true" demand values and fulfillment quantities are unknown for the counterfactual prices, we use the revenue component of V_D^1 in (3.6) as the predicted clearance revenue of D-OCPX. Table 2 shows the predictions on D-OCPX average channel prices (normalized by regular price), category-wide sales, and category-wide clearance revenues. We make the following observations. First, D-OCPX results in lower online prices than the historical online prices, resulting in a projeceted increase in online sales from customers buying online instead of from the store. Higher online sales is desirable since, unlike store sales that can only be fulfilled through same-zone stores, online sales can be directed by OMS to any store with slow moving inventory. A benefit of D-OCPX is that it uses prices to manage channel demand in accordance with partitions that rebalance inventory in the network. In stores with slow moving inventory, SFS has a large partition, resulting in a low store markdown. In stores with strong sales, the store markdown is also low to maximize profitability from store sales. Thus our second observation: D-OCPX results in higher average store prices than the historical average store prices. Hence, D-OCPX is addressing the concern of margin erosion due to the low store prices recommended by the legacy markdown optimization (MDO) system. In fact, our demand models predict more than 21% reduction in lost store sales opportunities at the recommended prices across categories. The net effect of the D-OCPX prices and the predicted fulfillment yields an increase in clearance revenues ranging between 6% and 12%, a combined increase of $12.5 million in the total clearance revenue of the 195 SKUs.

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. Their feedback was overwhelmingly positive. The pricing analysts have long recognized that the incumbent system fails for SKUs with significant cross-channel (SFS) fulfillment due to its inability to partition inventory between store demand and online fulfillment, and hence they requested to productize the OCPX solutions immediately.

4.2. Development of a commercial solution

A proprietary version of an omnichannel price analytics (OCPX) solution was approved for commercial deployment at IBM. To productize the OCPX solution, we collaborated with the IBM Commerce team who developed the necessary IT infrastructure, data processing, analytics integration, and user interface.

The team conducted system integration testing of the OCPX system to validate that the system performs in accordance with user expectations. This was done by comparing, on the same SKU, the price outputs of the OCPX system and of the incumbent MDO system. There were 43 clearance SKUs (across 10 consumer electronics categories) during 2015 Q4 through 2016 Q1 selected for the system integration test. Unlike the tests conducted with historical clearance data, these SKUs were soon to enter their clearance period. The test revealed that compared to the incumbent system, the OCPX system lowered store markdowns by 2 percentage points on average without any drop in the system predicted sell-through rate. The lower OCPX store markdowns, an important metric addressing the retailer's main concern of margin erosion, were consistent with the observations from the BVA. For each SKU, we also constructed a vector whose elements are the maximum markdowns of each store. We observed that the price dispersion in stores, which is empirically measured as the coefficient of variation of this vector, was up to 6.45% lower in the OCPX prices compared to incumbent MDO prices. This suggests OCPX results in a higher degree of price parity in the retail network.

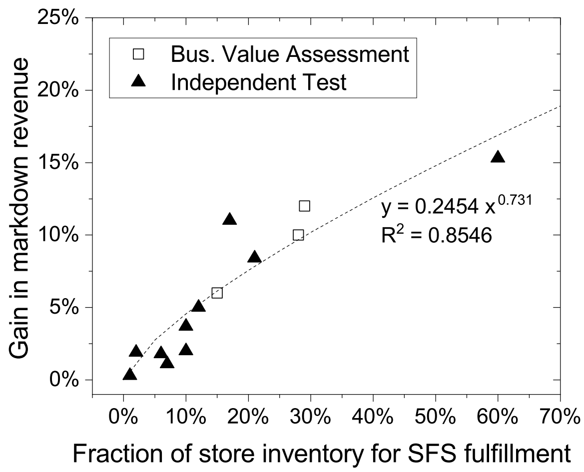

Figure 3 **Gain in markdown revenue plotted against store inventory partitions for SFS fulfillment. Each data point represents a product category in the system integration test.**

As part of the system integration test, projected revenues resulting from the price outputs of OCPX and the incumbent MDO were also computed. The retailer had years of experience with the incumbent demand forecasting engine and its accuracy for various categories. Hence for unbiased revenue projections, it was determined that the projected revenue of both price outputs should be estimated from the incumbent MDO system, with inventory adjusted (using the OCPX partitions) to account for cross-channel fulfillment in OCPX. No adjustments were made to compute the revenue from the incumbent prices since its undiscounted online price results in low online sales, hence, low fulfillment. The difference between the two revenue projections provides a projection of the revenue gain from the OCPX system. We plot the projected gain against the predicted degree of supply-side channel interdependency (Fig. 3) and observe that *the revenue gain is expected to be higher if cross-channel fulfillment is used more.* The horizontal axis of the figure is the OCPX partition for SFS (aggregated over SKUs in a category) divided by the aggregate store inventory, which is the predicted proportion of store inventory used for online fulfillment. The vertical axis is category-wide revenue gain from OCPX, which is computed as the relative difference between the OCPX projected revenue (aggregated over the category) and the incumbent MDO projected revenue (aggregated over the category).

4.3. Commercial release and pilot implementation

After a successful system integration test, a limited commercial release of the proprietary version of the OCPX system went live in March 2016, with the partner retailer as its pilot client. This was 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 the OCPX system. 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 optimality of the recommended prices depend on how accurately the OCPX inventory partitions predict the actual store inventory quantities 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 inventory partitions against the actual fulfillments during the same weeks. For 22 SKUs across 5 consumer electronics categories, 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.

4.3.1. 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.

Instead, to estimate the revenue improvement of OCPX in a real-world implementation, our approach was to analyze the pilot data through causal model analysis via statistical adjustment using appropriate pre-treatment predictors (Gelman and Hill 2006, Chapter 9). This is a variation of the Difference-in-Differences test (see Fisher et al. 2017), wherein we evaluate the change in the markdown season revenue rate between the treatment SKUs and the control SKUs by using the regular season revenue rate as a baseline for both and adjusting for other pre-treatment factors.

We obtained one year 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 were marked for clearance in both channels with a markdown "end date" (last date sold) in Q1 of 2017. We selected SKUs where the full clearance season is in the data. For an 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 with SFS using approximately 20% of store inventory.

To identify the presence of selection bias in the SKUs in control and treatment data, we compute a variety of regular season (pre-treatment) covariates such as the ticket price, average selling

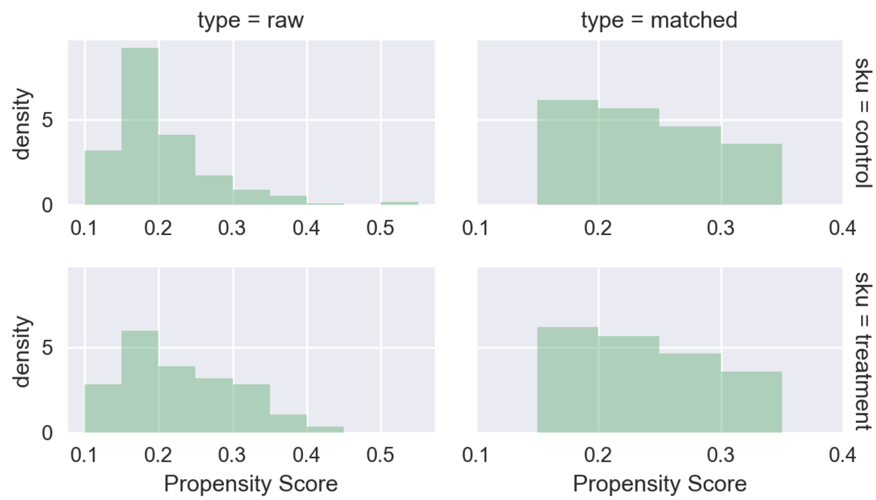

Figure 4 **Histograms of the density of propensity scores for observations before and after matching.**

price, sales rate, revenue rate, and channel shares for all SKUs. We observe that the range of the respective distributions across the control and treatment SKUs has a significant overlap. To formalize we use a non-parametric preprocessing matching method called `MatchIt` (Ho et al. 2007, 2011) that estimates a propensity score – defined as the probability of receiving the treatment given the pre-treatment covariates – and matches observations in treatment with those in control by balancing the propensity score (and thus the covariate distribution) to reduce the selection bias in parametric causal inference. We observe that higher propensity products correspond to those with high regular season promotion depths, strong sales and a large online presence. We expect OCPX to yield relatively higher benefits for higher propensity-score items, and lower benefits for lower propensity-score items characterized by weaker sales rates and a small online presence. This indicates that there is a selection bias, but the left-side panels of Fig. 4 that plot the raw density of the propensity scores for the treated and control SKUs reveal that this bias is not severe. The average propensity score is 0.2 for control and is slightly higher at 0.23 for the treatment SKUs. Next, we use `MatchIt` to select control SKUs in the mid-range scores (between 0.15 to 0.35) with nearest propensity score to the treatment SKUs and discard the remaining unmatched SKUs. The distributions of the propensity scores after matching are plotted in right-side panels of Fig. 4, showing nearly identical distributions with an average propensity score of 0.24 each. For this matched subclass, we conclude that there is no selection bias based on the descriptive covariates of the SKUs. To confirm this, recomputing propensity scores on this matched subclass results in scores of 0.5 ± 0.02 with 92% confidence.

We use the following regression model to estimate the average causal effect of OCPX (across all SKUs and within the matched subclass), where each SKU is an observation:

$$\begin{aligned}\ln(\text{Avg-Weekly-MD-Rev}_i) \sim \alpha_0 + \alpha_1 \ \ln(\text{Avg-Weekly-Reg-Rev}_i) + \alpha_2 \ \text{Avg-Weekly-MD-Inventory}_i \\ + \alpha_3 \ \text{Reg-Online-Share}_i + \alpha_4 \ \text{Treatment}_i \\ + \alpha_5 \ \text{Treatment}_i * \text{Reg-Online-Share}_i + \epsilon \qquad \forall i \in \text{Items}. \qquad (4.1)\end{aligned}$$

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. 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. 3 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). We also add the treatment variable by itself to gather local effects. Moreover, for both treatment and control SKUs, we expect that the higher the online presence, the retailer can clear the inventory better using SFS fulfillment. 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_4 and α_5. Table 3 and Table 4 show the estimation results for the data set with all SKUs and only the matched SKUs, respectively. In each table, we present the results of three regression models: with the main treatment effect only, with the interaction treatment effect only, and with both. For the matched SKUs, the coefficient for Reg-Online-Share (α_3) was not significant in all the three models, hence we report a model without it in Table 4. The tables report results from an F-test for joint significance of α_4 and α_5, where the null hypothesis is that both α_4 and α_5 are zero. Based on the p-values of the F-test, we can reject the null hypothesis on the raw data and the matched data. We observe that the coefficient estimates of the treatment effect (α_4, α_5) has the same sign across models. To compute average improvement due to OCPX on the treatment SKUs, we compute the average of the SKU-level benefits over all treatment SKUs. We note that among the three models, the one with only the interaction term results in the lowest gains, which is what we report. It is 13.7% for the data set with all SKUs and 20% to 24% (the former when we include the Reg-Online-Share control variable) for the model with matched SKUs only. We observe higher gains for the treatment SKUs in the matched subclass since the average sales and average online share (kept same for control SKUs in the matched subclass) is slightly higher compared to the treatment SKUs in the raw data.

Table 3 **Causal model estimated from the pilot data with all the SKUs.**

ln(Avg-Weekly-MD-Rev)	Coefficient	All [1]	All [2]	All [3]
Constant	α_0	1.820***	1.776***	1.819***
		(0.178)	(0.169)	(0.178)
ln(Avg-Weekly-Reg-Rev)	α_1	0.597***	0.598***	0.594***
		(0.021)	(0.021)	(0.021)
Avg-Weekly-MD-Inventory	α_2	0.004***	0.004***	0.004***
		(0.001)	(0.001)	(0.001)
Reg-Online-Share	α_3	0.563**	0.734**	0.636**
		(0.263)	(0.240)	(0.270)
Treatment	α_4		0.195**	0.140
			(0.097)	(0.118)
Treatment ∗ Reg-Online-Share	α_5	0.918*		0.497
		(0.509)		(0.619)
# obs		275	275	275
# treatment		57	57	57
R-sq		0.832	0.832	0.833
F-test for joint significance of α_4 and α_5		*	**	*
Improvement due to OCPX (lowest of above models): 13.7%				

$***p < .001$, $**p < .05$, $*p < .1$. SKUs with very small durations and rate of sales were eliminated. Standard errors are in the parenthesis.

Table 4 **Causal model estimated from the pilot data with matched SKUs only.**

ln(Avg-Weekly-MD-Rev)	Coefficient	Matched [1]	Matched [2]	Matched [3]
Constant	α_0	1.202**	1.078**	1.191**
		(0.374)	(0.364)	(0.373)
ln(Avg-Weekly-Reg-Rev)	α_1	0.664***	0.672***	0.657***
		(0.038)	(0.037)	(0.039)
Avg-Weekly-MD-Inventory	α_2	0.004***	0.004***	0.004***
		(0.001)	(0.001)	(0.001)
Treatment	α_4		0.325**	0.221
			(0.158)	(0.178)
Treatment ∗ Reg-Online-Share	α_5	1.743**		0.119
		(0.843)		(0.950)
# obs		78	78	78
# treatment		39	39	39
R-sq		0.850	0.850	0.853
F-test for joint significance of α_4 and α_5		**	**	*
Improvement due to OCPX (lowest of above models): 24%				

$***p < .001$, $**p < .05$, $*p < .1$. Matched control and treatment SKUs only and the rest are discarded. Standard errors are in the parenthesis.

5. Simulation experiments

The previous section describes a pilot implementation of the omnichannel pricing system at a large U.S. retailer, resulting in an estimated revenue increase compared to the legacy system. In this section, our aim is to have a controlled environment through which we conduct a detailed comparison of the omnichannel policies and the legacy policy in carefully designed simulation experiments.

5.1. Parameter generation

We randomly generate 27 problem instances, where each instance has randomly chosen parameters for the demand models and the zone-to-zone distances. We set $Z=5$ and $T=5$. To model the oftenly observed case that EFC is depleted at the start of a product's clearance period, we set $x_e=0$. To model high-ticket items sold by the partner retailer, we set the minimum and maximum price for the product as \$125 and \$500, respectively. We discretize the set of feasible prices into a finite set of 31 prices, with a stepsize of \$12.50. The salvage value for unsold items is \$50. The zone-to-zone distance matrix is populated with random distances from 0 to 2,500 miles. We then set the shipping cost between zones as $\$4.67+\$0.001\times d$, where d is the inter-zone distance in miles. Note that the shipping costs can take values between \$4.67 to \$7.17.

Given the chosen zone z channel prices at time t, the stochastic demand in channel m is:

$$\tilde{D}^t_{mz}(p_e,p_b)=N^t_z\times\theta_{mz}(p_e,p_b)\times(1+\tilde{w}_{mz}\delta_m),\quad \text{for } m=e,b, \tag{5.1}$$

where θ^t_{mz} is the deterministic MNL function (3.4). The demand stochasticity originates from the random demand factor $\tilde{\xi}^t_{mz}=1+\tilde{w}_{mz}\delta_m$ for all t, where δ_e,δ_b are parameters, and $\tilde{w}_{mz}$ is a random variable with mean of 0. We chose $\delta_b=0.5$ and $\delta_e=0.9$ to reflect our observations from the retailer's data that the forecast error is higher for online demand. Parameter α_{mz} in (3.4) is randomly chosen from the interval [10.8, 13.2]. The online price coefficient β_{ez} in (3.4) is randomly chosen from [0.0243, 0.0297], while the store price coefficient β_{bz} is randomly chosen from [0.0225, 0.0275]. Note that if both channel prices are \$350, then the average own and cross elasticities of the online channel are -6.1 and 5.3, respectively; for the store channel, they are -3.4 and 3.3, respectively. These high elasticity values are typical of products marked for clearance. We set

$$N^t_z=B\left(\frac{t}{T+1};a,b\right)\times\varepsilon_z\times\sigma,\quad t\in[T],z\in[Z], \tag{5.2}$$

where B is the pdf of a beta r.v. with randomly chosen shape parameters $b\in[1,5]$ and $a\in[1,b]$, ε_z is randomly chosen in [0,1], and σ is a scaling factor that ensures $\sum_z\sum_t N^t_z=100$. Thus, in all problem instances, the aggregate market size is always 100. To complete the specification of a problem instance, we set $x_{bz}=\sum_t N^t_z$, i.e., the available inventory at zone z is set to the expected market size at z. Note that the aggregate demand in zone z can be greater than $\sum_t N^t_z$ under some demand realizations.

The zero-mean random variable $\tilde{w}_{mz}$ has the distribution illustrated in Fig. EC.2 and can take values between -1 and 1. We ensure the realizations of w have the property

$$\sum_{m=e,b}\sum_z w_{mz}=0. \tag{5.3}$$

This models our observations that while the zone-level channel demand may deviate significantly from its expected value, the expected aggregate market size is a good predictor of the actual value. This also helps compare across different policies and sample paths that see the same total demand. Given the prices, the set of demand factors $(w_{mz})_{mz}$ determines the actual demand realizations. To conduct the simulation experiments, we randomly generate 10^3 sets of the demand factors.

5.2. Methodology

For any pricing policy, our simulations construct the distribution of realized profits (total revenue minus fulfillment costs) when the policy is applied to a specific problem instance. The distribution is constructed from the 10^3 samples of the demand factors $(w_{mz})_{mz}$.

We next describe how we compute the realized profit under a given sample of demand factors. For each week t starting from week 1, the policy sets channel prices $p_e^t, (p_{bz}^t)_{z\in Z}$ based on the current inventory levels. The actual demand $(D_{ez}^t, D_{bz}^t)_{z\in Z}$ is realized based on (5.1). The channel sales, fulfillment decisions, and profit realization are then determined by a fulfillment engine, which returns the inventory levels for week $t+1$. The fulfillment engine takes in as input the week t realized demand, the week t prices, and the week t store inventory levels. The engine was designed to approximate how fulfilment 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). If week t is divided into K fulfillment 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, the 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. The total profit is the sum of each week's profit and the salvage value from unsold items at the end of week 5.

The expected profit of the optimal pricing policy cannot be computed, and is also difficult to approximate using sample average approximation since the number of variables and constraints of optimization model is proportional to the sample size. Moreover, since the stochastic channel demand does not follow a multinomial distribution, the upper bound in Lemma 1 does not apply. However, as the following lemma states, we can prove an easily computable upper bound which could be used as a benchmark to the profit of any pricing policy.

LEMMA 4. *For any distribution of* $\tilde{w}$, $V^1(x) \le V_{PF} := E_{\tilde{w}}[U_{PF}(\tilde{w})]$, *where*

$$U_{PF}(w) := \max_{p=(p^1,\dots,p^T)} U^1(w,p;x), \tag{5.4}$$

and U^1 *is the optimal value of optimization model* (3.9) *with* $t=1$.

We next describe the pricing policies used in the simulation experiments. We tested our proposed pricing policies D-OCPX and R-OCPX. The parameters of uncertainty set (3.7) we use are $\frac{\Gamma^k}{2Z} = 0.1$, $a^t_{mz} = \delta^t_{mz}$, and $\Delta^t = 0$. This uncertainty set restricts the average absolute percentage error of zone-level demand to less than 10%, and only considers sample paths in which the aggregate weighted percentage error is zero. R-OCPX constructs $\Omega^{t,T}_Z$ for (3.11) by setting lower and upper bounds that are 95% and 105%, respectively, of the optimal price solution to (3.6).

We also test the following pricing policies based on the legacy pricing model:

- Policy Legacy sets the online price and store prices at period t as

$$p^t_e \leftarrow \arg\max_{p\in\Omega}\ p\cdot\min\left(x^t_e,\ \sum_{z\in[Z]} N^t_z\cdot\theta_{ez}(p,p^{nom})\right) \tag{5.5}$$

$$p^t_{bz} \leftarrow \arg\max_{p\in\Omega}\ p\cdot\min\left(x^t_{bz},\ N^t_z\cdot\theta_{bz}(p^{nom},p)\right), \qquad z\in[Z], \tag{5.6}$$

  where $p^{nom} = \$350$, a price level that is approximately the center of the feasible price range. The policy ignores any cross-channel interactions when optimizing the channel prices. The effect of the online price on stores is ignored, and conversely, the effect of store prices on the online demand is ignored. The full store inventory is used to determine store prices. This mimics the output of the legacy system. Since $x^t_e = 0$, p^t_e is set to \$500 (no markdown).

- Policy Legacy+Adjust sets the store prices in period t according to (5.6). Then to set the online price, it solves (5.5) but with the adjustment $x^t_e \leftarrow x^t_e + \nu\sum_{z\in[Z]} x^t_{bz}$ for some parameter $\nu \geq 0$. This policy mimics the manual adjustment process of the merchandise managers to the inputs of the legacy pricing system. Since $x^t_e = 0$, we chose $\nu = 0.27$, which is the average MNL online share if both channel prices are \$350. The store price is unadjusted because, based on our conversations, managers accept the store markdowns from the legacy system as correct, but take action on the undiscounted online price.

- Policy Legacy+OCPX sets the period t online and store prices based on (5.5)–(5.6), but with the adjustments $x^t_e \leftarrow x^t_e + \sum_{z,z'\in[Z]} y^*_{zz'}$ and $x^t_{bz} \leftarrow x^t_{bz} - \sum_{z'\in[Z]} y^*_{zz'}$, where y^* is the optimal inventory partition from (3.6). This policy model ignore cross-channel demand interactions but partitions store inventory to account for cross-channel fulfillment.

We implemented the computational experiments in Python 3.5. The optimization models are solved using CPLEX Optimization Studio 12.7. We set CPLEX to terminate when an integer feasible solution is within 0.01% of optimality. The average run time of D-OCPX and R-OCPX on a sample path are 6.796 sec and 7.895 sec, respectively. The average run time of Legacy, Legacy+Adjust, and Legacy+OCPX are 1.025 sec, 1.712 sec, and 11.046 sec, respectively.

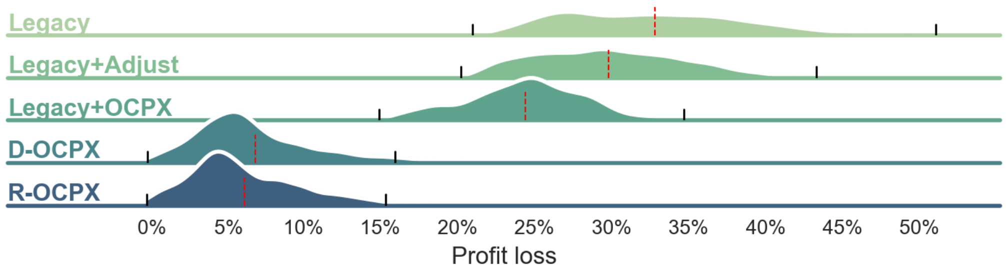

Figure 5 **The profit loss distribution under the pricing policies in one problem instance. The kernel density is estimated from the profit loss in 10^3 sample paths. The black vertical lines mark the range (removing statistical outliers), while the dashed red vertical lines mark the mean of the distribution.**

Table 5 **The average price markdown, average sales, and average unsold units of different pricing policies applied to a specific problem instance. The average is estimated from 1000 sample paths.**

	Avg. Price markdown		Avg. Inventory usage		
	Online	Store	Online sales	Store sales	Unsold
Legacy	0%	41%	0.5%	85%	14%
Legacy+Adjust	26%	41%	10%	79%	11%
Legacy+OCPX	35%	39%	32%	59%	10%
D-OCPX	16%	12%	40%	48%	12%
R-OCPX	13%	11%	36%	50%	14%
Perfect Foresight	11%	10%	42%	54%	4%

5.3. Results

Fig. 5 shows the distributions of profit loss in one specific problem instance. The profit loss on a sample path w is defined as the relative difference between the realized profit of a pricing policy and $U_{PF}(w)$ defined in (5.4). We show only the distributions in one problem instance, since the distributions in all 27 problem instances follow a similar pattern. Particularly, a clear pattern is that the pricing policies based on the legacy model (Legacy, Legacy+Adjust, Legacy+OCPX) have a significantly higher profit loss compared to the omnichannel pricing policies (D-OCPX, R-OCPX). This is further demonstrated in Fig. 6 which uses scatter plots to compare the *average* profit loss of the legacy pricing policies to that of D-OCPX in all 27 problem instances. The average profit loss of D-OCPX is between 5% to 8%. The prices from the unadjusted legacy model (Legacy) results in roughly 20% to 40% average profit loss. Adjusting the pricing policy by injecting online inventory (Legacy+Adjust) reduces the average profit loss slightly. Out of the legacy models, the one that incurs the lowest average profit loss is Legacy+OCPX, which has average profit losses of 15% to 30%.

In order to understand the causes of the above observations, we next compare and contrast the pricing policies with respect to their average markdowns and their average sales. Table 5 presents

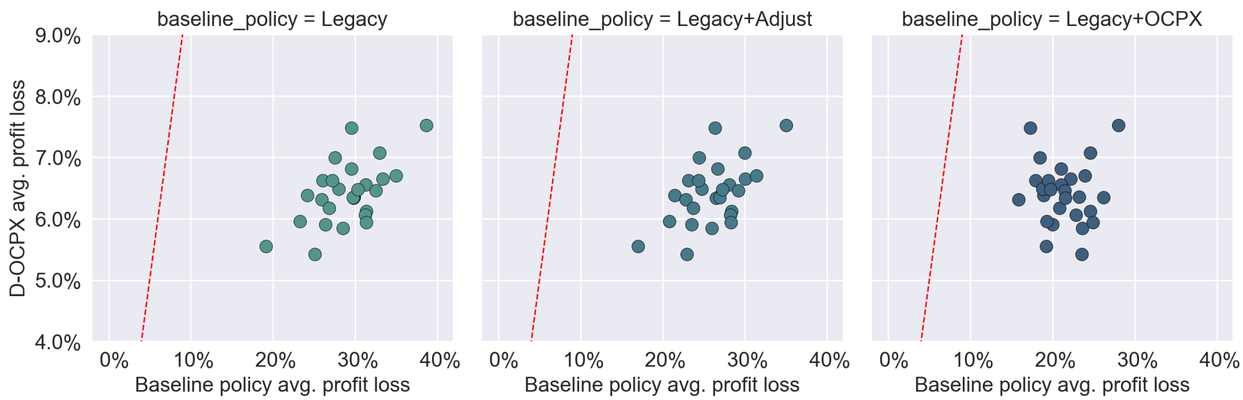

Figure 6 **Scatter plot of D-OCPX average profit loss and a baseline policy (Legacy, Legacy+Adjust, Legacy+OCPX) average profit loss. Points to the right of the red line are problem instances where D-OCPX has a lower average profit loss compared to the baseline.**

these statistics for the problem instance corresponding to Fig. 5. From the table, we observe that policies Legacy and Legacy+Adjust result in only a small percentage of inventory sold through the online channel. This is due to the significantly larger store markdown compared to the online markdown. Since customers consider the online channel and store channel as substitutes, this price difference induces a larger proportion of customers to purchase from stores. Due to the large store markdowns, the resulting effect is a low overall profit margin.

The legacy store markdown optimization model assumes that the store inventory input will be sold only through the store channel. Hence, the big store markdowns of Legacy and Legacy+Adjust are the result of the unadjusted store inventory inputs to the legacy model. But based on Table 5, it is optimal to sell an average of 42% of the store inventory through the online channel. Since the D-OCPX model (3.6) optimizes the inventory partitions which allocate store inventory between the channels, we also tested the Legacy+OCPX pricing policy which adjusts the inventory inputs to the legacy model based on the optimized partitions from (3.6). With this adjustment, the average online and store markdowns are brought closer together (see Table 5), implying that the inventory partitions of (3.6) correctly identify a balanced allocation of inventory between channels. However, the average channel markdowns of 35% and 39% with Legacy+OCPX are larger than the Perfect Foresight markdowns. Thus, even if inventory is balanced across the network, a markdown optimization model that ignores channel substitution and does not jointly optimize omnichannel prices results in significant profit loss.

Fig. 5 shows that for a specific problem instance, R-OCPX incurs a lower average profit loss than D-OCPX. Fig. 7 shows the distribution (over the 10^3 samples) of profit increase due to R-OCPX, with a box plot for each of the 27 problem instances. The profit increase was computed as the D-OCPX profit loss minus the R-OCPX profit loss. The red line in the figure denotes no improvement

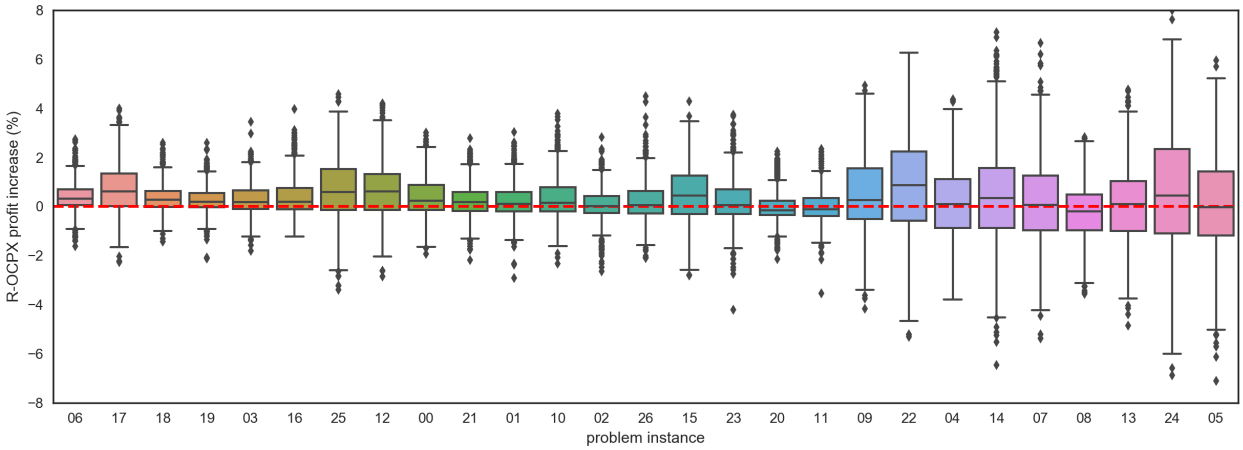

Figure 7 **Boxplots of the R-OCPX profit increase in 10^3 samples. The boxplots have been ordered in decreasing 25% quantile.**

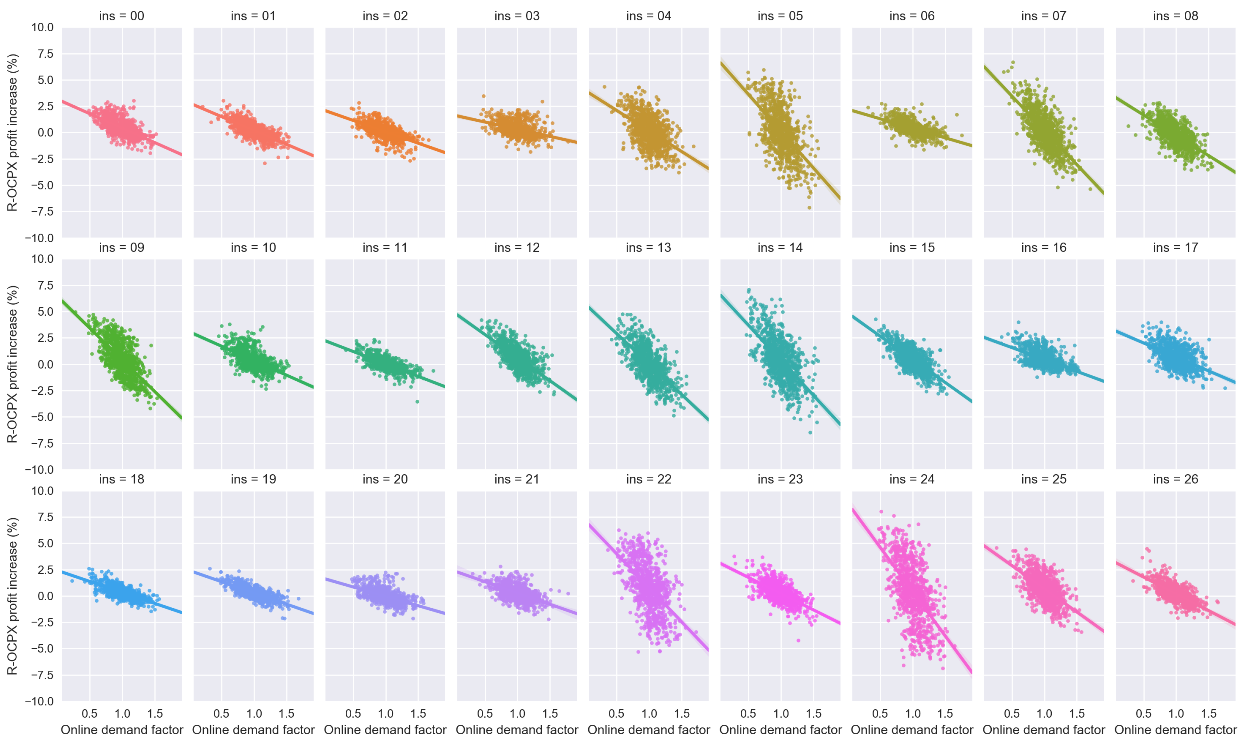

Figure 8 **Scatter plots and linear fit of the profit increase by R-OCPX against the average online demand factor. The plots demonstrate a decreasing relationship between the profit increase and the average online factor.**

in profit. In the first nine problem instances shown in the figure, R-OCPX has a higher profit than D-OCPX in about 75% of the samples. In almost all instances, the median R-OCPX profit increase is positive.

We next investigate which cases result in the largest profit increase by R-OCPX. Define the online demand factor $\bar{\xi}_e := 1 + \delta_e \bar{w}_e$, where $\bar{w}_e = \frac{1}{Z}\sum_{z\in[Z]} w_{ez}$ is the average online percentage deviation. Fig. 8 shows scatter plots of the R-OCPX profit increase plotted against $\bar{\xi}_e$. Note that $\bar{\xi}_e < 1$ if the online demand is lower than its expected value. Due to the procedure for generating the demand factor samples in (5.3), the average store percentage deviation is $-\bar{w}_e$. In other words, if the average online demand is higher than expected (i.e., $\bar{\xi}_e > 1$), then the average store demand is lower than expected. It is clear from these scatter plots that the largest profit increase due to R-OCPX occurs whenever the average online factor is small, i.e., when the actual online demand is lower than the predicted values and actual store demand is higher than predicted.

Observe that the omnichannel demand uncertainty has an asymmetric impact on a retailer's expected profitability. The reason is when store demand is lower than predicted and online demand is higher than forecasted, the omnichannel retailer can gainfully employ SFS fulfilment to satisfy the excess online demand using unused store inventory, resulting in a 'win-win' revenue situation. On the other hand, when store sales are higher than expected (exceeding store inventory), and online demand is correspondingly lower, this results in lost in-store sales, and the SFS option is also of limited use due to weak online demand, resulting in a 'loss-loss' in both channels. Therefore, the robust "adversary" can be expected to choose low online demand and high store demand (particularly, higher store demand for stores with low inventory). As can be seen in Fig. 8, the proposed R-OCPX solution increasingly outperforms D-OCPX when omnichannel demand realizations are highly adversarial ($\bar{\xi}_e$ is very small). This also explains why R-OCPX profitably protects against worst-case demand realization by optimally raising prices relative to D-OCPX (see Table 5).

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. Next we thank Prof. Yan Huang for her valuable suggestions related to propensity score analysis for the causal model.

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* .

Bertsekas, Dimitri P. 1995. *Dynamic programming and optimal control*, vol. 1. Athena scientific Belmont, MA.

Bertsimas, Dimitris, Frans JCT de Ruiter. 2016. Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS Journal on Computing* **28**(3) 500–511.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* **5**(3) 203–229.

Bitran, G., R. Caldentey, S. Mondschein. 1998. Coordinating clearance markdown sales of seasonal products in retail chains. *Operations research* **46**(5) 609–624.

Bitran, G., S.V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Science* **43**(1) 64–79.

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.

Charnes, A., W.W. Cooper. 1962. Programming with linear fractional functionals. *Naval Research logistics quarterly* **9**(3-4) 181–186.

Chen, X., D. Simchi-Levi. 2012. *The Oxford handbook of pricing management*, chap. Pricing and inventory management. Oxford University Press, 784–822.

Cooper, William L. 2002. Asymptotic behavior of an allocation policy for revenue management. *Operations Research* **50**(4) 720–727.

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.

Ettl, Markus, Pavithra Harsha, Shivaram Subramanian. 2017. Distributed optimization method for real-time omnichannel retail operations. Patent. Patent No. US9626646 B1.

Fisher, Marshall, Santiago Gallino, Jun Li. 2017. Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science* .

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. Ettl. 2015. A practical price optimization approach for omni-channel retailing. Submitted.

Ho, Daniel E, Kosuke Imai, Gary King, Elizabeth A Stuart. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis* **15**(3) 199–236.

Ho, Daniel E, Kosuke Imai, Gary King, Elizabeth A Stuart, et al. 2011. Matchit: nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software* **42**(8) 1–28.

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

Jasin, Stefanus, Sunil Kumar. 2013. Analysis of deterministic lp-based booking limit and bid price controls for revenue management. *Operations Research* **61**(6) 1312–1320.

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. 2017. Dynamic joint pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Operations Management, Forthcoming* .

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`.

Sherali, H.D., W.P. Adams. 1999. Reformulation-linearization techniques for discrete optimization problems. *Handbook of combinatorial optimization*. Springer, 479–532.

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.

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.

Zeng, Bo, Long Zhao. 2013. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters* **41**(5) 457–461.

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

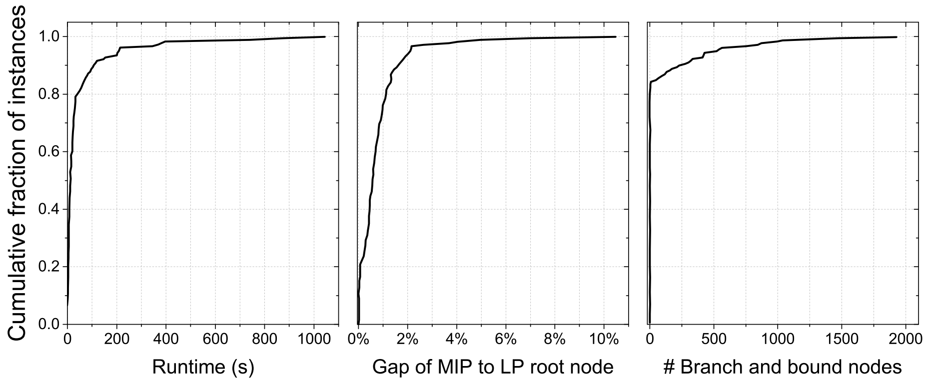

Figure EC.1 **Distribution of the computational runtimes, the percentage gap of the MIP optimal solution with respect to its LP root node, and number of branch and bound nodes.**

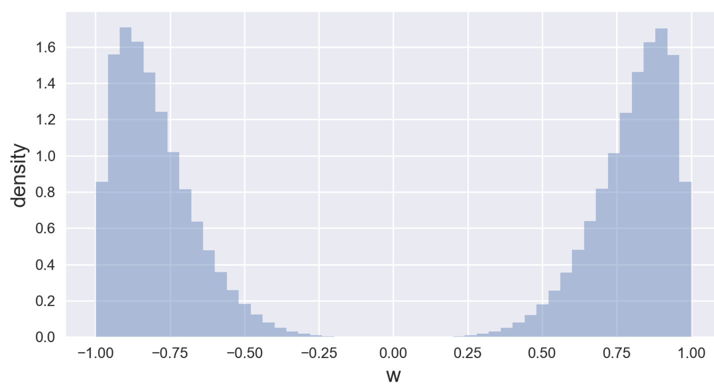

Figure EC.2 **The distribution of the demand uncertainty factors used in the simulations.**

EC.1. Proofs

EC.1.1. Proof of Lemma 1

Note that V^1 defined by the Bellman equations (3.1)–(3.2) gives the optimal expected profit of a pricing policy that sets the same price for customers arriving in each period t. An upper bound for V^1 is the optimal expected profit for a pricing policy that can set a different price and fulfillment for each arriving customer.

Given an arrival to zone z, if the online price and in-store price are p_e and p_b, respectively, then the probability that the customer will choose to buy online is $\theta_{ez}(p_e, p_b)$ and from the store is $\theta_{bz}(p_e, p_b)$. Let us define the functions:

$$\psi^{tn}_{mz}(p_e, p_b) := \frac{N^t_z}{N^t} \cdot \theta_{mz}(p_e, p_b), \quad n \in [N^t],\ t \in [T],\ m = e, b,\ z \in [Z], \tag{A.1}$$

which is the probability of a customer arriving in subperiod n during period t to zone z and choosing to purchase from channel m given prices (p_e, p_b).

Let p_e^{tn} be the online price and p_{bz}^{tn} be the store price observed by an arrival $n \in [N^t]$ to zone $z \in [Z]$ during period $t \in [T]$. Define $\tilde{D}_{mz}^{tn}(p_e^{tn}, p_{bz}^{tn})$ as the stochastic demand for channel $m = e, b$ in zone z during subperiod $n \in [N^t]$ of period t. Note that $E[\tilde{D}_{mz}^{tn}(p_e^{tn}, p_{bz}^{tn})] = \psi_{mz}^{tn}(p_e^{tn}, p_{bz}^{tn})$, and that

$$\sum_{z\in[Z]} \sum_{m=e,b} \psi_{mz}^{tn}(p_e^{tn}, p_{bz}^{tn}) \leq 1, \quad \forall n \in [N^t], \forall t \in [T]. \tag{A.2}$$

Let us assume that there exists a $p_\infty \in \Omega$ such that $\tilde{D}_{bz}^{tn}(p_e, p_\infty) = 0$ for any $p_e \in \Omega$. Define $\tilde{R}_{ez}^{tn}(p_e^{tn}, p_{bz}^{tn}) = p_e^t \tilde{D}_{ez}^{tn}(p_e^{tn}, p_{bz}^{tn})$ and $\tilde{R}_{bz}^{tn}(p_e^{tn}, p_{bz}^{tn}) = p_{bz}^t \tilde{D}_{bz}^{tn}(p_e^{tn}, p_{bz}^{tn})$ as the stochastic revenue from the online and store channel, respectively, in zone z during subperiod $n \in [N^t]$ of period t.

Let $\tilde{\lambda}_{ij}^{tnz}$ be a binary random variable which equals to 1 if and only if a zone z arrival in subperiod n of period t is presented by price p_i online and p_j in the store. Let $\tilde{\mu}_i^{tne}$ be a binary random variable which equals to 1 if and only if a zone z arrival in subperiod n of period t is presented with price p_i online. Let $\tilde{y}_{ez}^{tn}$ and $\tilde{y}_{z'z}^{tn}$ be the random fulfillment from the EFC and the store in zone z', respectively, of an arrival in zone z during subperiod n of period t. Let $\tilde{s}_{bz}^{tn}$ be the random store sales in zone z during subperiod n of period t.

Therefore, $V^1 \leq V^*$, where

$$V^* = \underset{\lambda,\mu,y,s}{\text{maximize}} \quad E\left[\sum_{t\in[T]} \sum_{n\in[N^t]} \sum_{z\in[Z]} \sum_{m=e,b} \sum_{i\in[I]} \sum_{j\in[I]} \tilde{\lambda}_{ij}^{tnz} \tilde{R}_{mz}^{tn}(p_i, p_j)\right] + q\left(x_e + \sum_{z\in[Z]} x_{bz}\right) \tag{A.3a}$$

$$- E\left[\sum_{t\in[T]} \sum_{n\in[N^t]} \sum_{z\in[Z]} \left(q\tilde{s}_{bz}^{tn} + (c_{ez} + q)\tilde{y}_{ez}^{tn} + \sum_{z'\in[Z]} (c_{z'z} + q)\tilde{y}_{z'z}^{tn}\right)\right] \tag{A.3b}$$

subject to

$$\tilde{y}_{ez}^{tn} + \sum_{z'\in[Z]} \tilde{y}_{z'z}^{tn} = \sum_{i\in[I]} \sum_{j\in[I]} \tilde{\lambda}_{ij}^{tnz} \tilde{D}_{ez}^{tn}(p_i, p_j), \quad \forall z \in [Z],\ \forall t \in [T],\ \forall n \in [N^t], \tag{A.3c}$$

$$\tilde{s}_{bz}^{tn} = \sum_{i\in[I]} \sum_{j\in[I]} \tilde{\lambda}_{ij}^{tnz} \tilde{D}_{bz}^{tn}(p_i, p_j), \quad \forall z \in [Z],\ \forall t \in [T],\ \forall n \in [N^t], \tag{A.3d}$$

$$\sum_{t\in[T]} \sum_{n\in[N^t]} \tilde{y}_{ez}^{tn} \leq x_e, \tag{A.3e}$$

$$\sum_{t\in[T]} \sum_{n\in[N^t]} \left(\tilde{s}_{bz}^{tn} + \sum_{z'\in[Z]} \tilde{y}_{z'z}^{tn}\right) \leq x_{bz}, \quad \forall z \in [Z], \tag{A.3f}$$

$$\sum_{j\in[I]} \tilde{\lambda}_{ij}^{tnz} = \tilde{\mu}_i^{tne}, \quad \forall t \in [T],\ n \in [N^t],\ z \in [Z],\ i \in [I], \tag{A.3g}$$

$$\sum_{i\in[I]} \tilde{\mu}_i^{tne} = 1, \quad \forall t \in [T],\ n \in [N^t], \tag{A.3h}$$

$$\tilde{y} \geq 0,\ \tilde{s} \geq 0,\ \tilde{\lambda} \in \{0,1\},\ \tilde{\mu} \in \{0,1\} \tag{A.3i}$$

Let $P_{zt} : \Omega \times \Omega \mapsto [0,1]$ be the marginal probability distribution of zone z prices under the optimal policy π^* to (A.3). Let us define $\bar{\psi}^{tn}_{mz}$ and $\bar{r}^{tn}_{mz}$ be the optimal expected demand rate and revenue rate, respectively, during subperiod n of period t in channel m of zone z. Note that

$$\bar{\psi}^{tn}_{mz} = \sum_{i\in[I]}\sum_{j\in[I]} \psi^{tn}_{mz}(p_i,p_j)P_{zt}(p_i,p_j), \quad m=e,b,\ \forall t\in[T], \forall n\in[N^t], \forall z\in[Z], \tag{A.4}$$

$$\bar{r}^{tn}_{mz} = \sum_{i\in[I]}\sum_{j\in[I]} r^{tn}_{mz}(p_i,p_j)P_{zt}(p_i,p_j), \quad m=e,b,\ \forall t\in[T], \forall n\in[N^t], \forall z\in[Z] \tag{A.5}$$

Also, $\sum_{j\in[I]} P_{zt}(p_i,p_j) = \sum_{j\in[I]} P_{z't}(p_i,p_j) = P_t(p_i)$ for some $P_t(p_i)\in[0,1]$ and any $z,z'\in[Z]$ and $p_i\in\Omega$. Finally, given the optimal demand rate, the probability distribution of fulfillment and store sales under π^* achieves an expected fulfillment cost equal to:

$$
\begin{aligned}
FC^* = \underset{s\geq 0,\ y\geq 0}{\text{minimize}} \quad & \sum_{t\in[T]}\sum_{n\in[N^t]}\sum_{z\in[Z]} \left(q s^{tn}_{bz} + (c_{ez}+q) y^{tn}_{ez} + \sum_{z\in[Z]} y^{tn}_{z'z} \right) && \text{(A.6a)}\\
\text{subject to} \quad & y^{tn}_{ez} + \sum_{z'\in[Z]} y^{tn}_{z'z} = \bar{\psi}^{tn}_{ez}, \qquad \forall z\in[Z], t\in[T], n\in[N^t], && \text{(A.6b)}\\
& s^{tn}_{bz} = \bar{\psi}^{tn}_{bz}, \qquad \forall z\in[Z], t\in[T], n\in[N^t], && \text{(A.6c)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]} y^{tn}_{ez} \leq x_e, && \text{(A.6d)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]} \left(s^{tn}_{bz} + \sum_{z'\in[Z]} y^{tn}_{z'z} \right) \leq x_{bz}, \quad \forall z\in[Z] && \text{(A.6e)}
\end{aligned}
$$

Defining $r^{tn}_{ez}(p_i,p_j) := p_i\psi^{tn}_{ez}(p_i,p_j)$ and $r^{tn}_{bz}(p_i,p_j) := p_j\psi^{tn}_{bz}(p_i,p_j)$, consider the deterministic LP counterpart:

$$
\begin{aligned}
V^{LP} = \underset{\lambda,\mu,y,s}{\text{maximize}} \quad & \sum_{t\in[T]}\sum_{n\in[N^t]}\sum_{z\in[Z]}\sum_{m=e,b}\sum_{i\in[I]}\sum_{j\in[I]} \lambda^{tnz}_{ij} r^{tn}_{mz}(p_i,p_j) + q\left(x_e + \sum_{z\in[Z]} x_{bz} \right) && \text{(A.7a)}\\
& - \sum_{t\in[T]}\sum_{n\in[N^t]}\sum_{z\in[Z]} \left(q s^{tn}_{bz} + (c_{ez}+q) y^{tn}_{ez} + \sum_{z'\in[Z]} (c_{z'z}+q) y^{tn}_{z'z} \right) && \text{(A.7b)}
\end{aligned}
$$

subject to

$$
\begin{aligned}
& y^{tn}_{ez} + \sum_{z'\in[Z]} y^{tn}_{z'z} = \sum_{i\in[I]}\sum_{j\in[I]} \lambda^{tnz}_{ij} \psi^{tn}_{ez}(p_i,p_j), \quad \forall z\in[Z],\ \forall t\in[T],\ \forall n\in[N^t], && \text{(A.7c)}\\
& s^{tn}_{bz} = \sum_{i\in[I]}\sum_{j\in[I]} \lambda^{tnz}_{ij} \psi^{tn}_{bz}(p_i,p_j), \qquad \forall z\in[Z],\ \forall t\in[T],\ \forall n\in[N^t], && \text{(A.7d)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]} y^{tn}_{ez} \leq x_e, && \text{(A.7e)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]} \left(s^{tn}_{bz} + \sum_{z'\in[Z]} y^{tn}_{z'z} \right) \leq x_{bz}, \qquad \forall z\in[Z], && \text{(A.7f)}\\
& \sum_{j\in[I]} \lambda^{tnz}_{ij} = \mu^{tne}_i, \qquad \forall t\in[T],\ n\in[N^t],\ z\in[Z],\ i\in[I], && \text{(A.7g)}\\
& \sum_{i\in[I]} \mu^{tne}_i = 1, \qquad \forall t\in[T],\ n\in[N^t], && \text{(A.7h)}\\
& y\geq 0,\ s\geq 0,\ \lambda\geq 0,\ \mu\geq 0 && \text{(A.7i)}
\end{aligned}
$$

We next construct a feasible solution (λ,μ,s,y) to (A.7) which achieves an objective value equal to V^*. In particular, choose

$$\lambda_{ij}^{tnz} = P_{zt}(p_i,p_j), \quad \forall t\in[T], \forall n\in[N^t], \forall z\in[Z], \forall i\in[I], \forall j\in[I] \tag{A.8}$$

$$\mu_i^{tne} = P_t(p_i), \qquad \forall t\in[T], \forall n\in[N^t], \forall i\in[I]. \tag{A.9}$$

It is easy to verify that (λ,μ) satisfy constraints (A.7g)–(A.7i). We let (s,y) be the solution to a linear program of the following form:

$$\begin{aligned}
FC' = \;& \underset{s\geq 0,\; y\geq 0}{\text{minimize}} \quad \sum_{t\in[T]}\sum_{n\in[N^t]}\sum_{z\in[Z]}\left(q s_{bz}^{tn} + (c_{ez}+q) y_{ez}^{tn} + \sum_{z\in[Z]} y_{z'z}^{tn}\right) && \text{(A.10a)}\\
& \text{subject to} \\
& y_{ez}^{tn} + \sum_{z'\in[Z]} y_{z'z}^{tn} = \sum_{i\in[I]}\sum_{j\in[I]} P_{zt}(p_i,p_j)\psi_{ez}^{tn}(p_i,p_j), \quad \forall z\in[Z], t\in[T], n\in[N^t], && \text{(A.10b)}\\
& s_{bz}^{tn} = \sum_{i\in[I]}\sum_{j\in[I]} P_{zt}(p_i,p_j)\psi_{bz}^{tn}(p_i,p_j), \quad \forall z\in[Z], t\in[T], n\in[N^t], && \text{(A.10c)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]} y_{ez}^{tn} \leq x_e, && \text{(A.10d)}\\
& \sum_{t\in[T]}\sum_{n\in[N^t]}\left(s_{bz}^{tn} + \sum_{z'\in[Z]} y_{z'z}^{tn}\right) \leq x_{bz}, \quad \forall z\in[Z] && \text{(A.10e)}
\end{aligned}$$

Note that due to (A.10b)–(A.10e), we have that (λ,μ,s,y) also satisfy (A.7c)–(A.7f). Additionally, from definition (A.4), the right-hand side of (A.10b) and (A.10c) are equal to $\bar{\psi}_{ez}^{tn}$ and $\bar{\psi}_{bz}^{tn}$, respectively. Hence, $FC^* = FC'$. Therefore,

$$V^{LP} \geq \sum_{t\in[T]}\sum_{n\in[N^t]}\sum_{z\in[Z]}\sum_{m=e,b}\underbrace{\sum_{i\in[I]}\sum_{j\in[I]} P_{zt}(p_i,p_j) r_{mz}^{tn}(p_i,p_j)}_{=\bar{r}_{mz}^{tn}} + q\left(x_e + \sum_{z\in[Z]} x_{bz}\right) - FC' = V^* \geq V^1.$$

What is left to prove Lemma 1 is to show that in each period t, a stationary solution is optimal for (A.7). Note that for any optimal solution (λ,μ,y,s), a stationary solution $(\bar{\lambda},\bar{\mu},\bar{y},\bar{s})$ is also optimal, where for any $t\in[T]$ and $n\in[N^t]$, we set

$$\bar{\lambda}_{ij}^{tnz} = \frac{1}{N^t}\sum_{n'\in[N^t]}\lambda_{ij}^{tn'z}, \quad \bar{y}_{ez}^{tn} = \frac{1}{N^t}\sum_{n'\in[N^t]} y_{ez}^{tn'}, \quad \bar{y}_{z'z}^{tn} = \frac{1}{N^t}\sum_{n'\in[N^t]} y_{z'z}^{tn'}, \quad \bar{s}_{bz}^{tn} = \frac{1}{N^t}\sum_{n'\in[N^t]} s_{bz}^{tn'}. \quad \square$$

EC.1.2. Proof of Lemma 2

To prove the first part of the lemma, it suffices to prove that $V_D^t(x^t) = \bar{V}^*$, where

$$\bar{V}^* := \underset{\lambda,\mu,s,y,u}{\text{maximize}} \quad \sum_{k=t}^{T}\sum_{z\in[Z]}\sum_{m=e,b}\sum_{i\in[I]} p_i s_{mzi}^k - \sum_{z\in[Z]} c_{ez}y_{ez} - \sum_{z\in[Z]}\sum_{z'\in[Z]} c_{zz'}y_{zz'} + q\left(u_e + \sum_{z\in[Z]} u_{bz}\right) \quad \text{(A.11a)}$$

subject to Constraints (3.6c) – (3.6e),

$$\sum_{i\in[I]} s_{mzi}^k \le \sum_{i\in[I]}\sum_{j\in[I]} \lambda_{ij}^{kz} d_{mz}^k(p_i,p_j), \quad k=t,\dots,T,\ m=e,b,\ z\in[Z], \quad \text{(A.11b)}$$

$$s_{ezi}^k \le \sum_{j\in[I]} \lambda_{ij}^{kz} d_{ez}^k(p_i,p_j), \quad k=t,\dots,T,\ z\in[Z],\ i\in[I], \quad \text{(A.11c)}$$

$$s_{bzj}^k \le \sum_{i\in[I]} \lambda_{ij}^{kz} d_{bz}^k(p_i,p_j), \quad k=t,\dots,T,\ z\in[Z],\ i\in[I], \quad \text{(A.11d)}$$

$$\sum_{i\in[I]} \mu_i^{ke} = 1, \quad \sum_{j\in[I]} \mu_j^{kz} = 1, \quad \forall z\in[Z], \quad \text{(A.11e)}$$

$$\lambda_{ij}^{kz} \le \mu_i^{ke}, \quad \lambda_{ij}^{kz} \le \mu_j^{kz} \quad \forall z\in[Z],\ i\in[I],\ j\in[I], \quad \text{(A.11f)}$$

$$s\ge 0, y\ge 0, u\ge 0,\ \lambda,\mu\in\{0,1\} \quad \text{(A.11g)}$$

Since the feasible price set is discrete, then (3.6) is equivalent to the optimization model with the objective (A.11a), and with constraints (3.6c)–(3.6e) and (A.11b)–(A.11g). Let us denote by $(\hat{\lambda},\hat{\mu},\hat{s},\hat{y},\hat{u})$ the optimal solution to this model, which has an objective value of $V_D^t(x^t)$. Defining

$$s_{ezi}^k = \hat{\mu}_i^{ke}\hat{s}_{ez}^k, \qquad k=t,\dots,T,\ \forall z\in[Z],\ \forall i\in[I],$$
$$s_{bzj}^k = \hat{\mu}_j^{kz}\hat{s}_{bz}^k, \qquad k=t,\dots,T,\ \forall z\in[Z],\ \forall j\in[I],$$

we can easily check that $(\hat{\lambda},\hat{\mu},s,\hat{y},\hat{u})$ is a feasible solution to (A.11) and achieves the objective value $V_D^t(x^t)$. Hence, $V_D^t(x^t) \le \bar{V}^*$.

We next show that $V_D^t(x^t) \ge \bar{V}^*$. Let us denote the maximizer of (A.11) as $(\bar{\lambda},\bar{\mu},\bar{s},\bar{y},\bar{u})$, which achieves the optimal value $\bar{V}^*$. From the binary constraints of the MIP, and from constraints (A.11e), it follows that for time k, there exist price indices $(i_{ke}, j_{kz_1},\dots,j_{kz_n})$ such that:

$$\bar{\mu}_{i_{ke}}^{ke} = 1, \qquad \bar{\mu}_i^{ke} = 0, \quad \forall i\ne i_{ke}, \quad \text{(A.12)}$$

$$\bar{\mu}_{j_{kz}}^{kz} = 1, \quad \forall z\in[Z], \qquad \bar{\mu}_j^{kz} = 0, \quad \forall j\ne j_{kz},\ \forall z\in[Z]. \quad \text{(A.13)}$$

These then imply from constraints (A.11f) that $\bar{\lambda}_{ij}^{kz} = 0$ for all i,j where $i\ne i_{ke}$ or $j\ne j_{kz}$, for all $z\in[Z]$. Therefore, from constraint (A.11c), it follows that $\bar{s}_{ezi}^k = 0$ for all $i\ne i_{ke}$, $z\in[Z]$. Similarly, from constraint (A.11d), it follows that $\bar{s}_{bzj}^k = 0$ for all $j\ne j_{kz}$, $z\in[Z]$. Define

$$s_{ez}^k = \bar{s}_{ezi_{ke}}^k, \qquad k=t,\dots,T,\ \forall z\in[Z],$$
$$s_{bz}^k = \bar{s}_{bzj_{kz}}^k, \qquad k=t,\dots,T,\ \forall z\in[Z]$$

Based on the definitions (A.12)–(A.13), we have that $\bar{s}_{ezi}^k = \bar{\mu}_i^{ke}s_{ez}^k$ and $\bar{s}_{bzj}^k = \bar{\mu}_j^{kz}s_{bz}^k$ for any $k\in[T]$, $z\in[Z]$ and $i\in[I]$. Therefore $(\bar{\lambda},\bar{\mu},s,\bar{y},\bar{u})$ is a feasible solution to (3.6) with an objective value $\bar{V}^*$. Hence, $\bar{V}^* \le V_D^t(x^t)$, which proves the lemma. □

To prove the second part of the lemma, assume the MNL model (3.4) of choice for arrivals in zone $z \in [Z]$. Under this demand model, assuming that N_z^t is the expected number of zone z arrivals in period t, $d_{mz}^t(p_i, p_j) = N_z^t \cdot \theta_{mz}(p_i, p_j)$ for all $i, j \in [I]$. Let us introduce the following constants:

$$\gamma_{ezi}^k := \exp\left(\alpha_{ez} - \beta_{ez} p_i\right), \quad k = t, \ldots, T, \, \forall z \in [Z], \forall i \in [I], \tag{A.14}$$

$$\gamma_{bzj}^k := \exp\left(\alpha_{bz} - \beta_{bz} p_j\right), \quad k = t, \ldots, T, \, \forall z \in [Z], \forall j \in [I]. \tag{A.15}$$

Therefore, by introducing binary decision variables for the price decisions, we can reformulate optimization model (3.6) as:

$$\begin{aligned}
V_D^t(x^t) = \underset{\mu, s, y, u}{\text{maximize}} \quad & \sum_{k=t}^{T} \sum_{z \in [Z]} \sum_{i \in [I]} p_i \left(s_{ez}^k \mu_i^{ke} + s_{bz}^k \mu_i^{kz}\right) - \sum_{z \in [Z]} c_{ez} y_{ez} - \sum_{z, z' \in [Z]} c_{zz'} y_{zz'} + q \left(u_e + \sum_{z \in [Z]} u_{bz}\right) && \text{(A.16a)} \\
\text{subject to} \quad & \text{Constraints } (3.6\text{c}) - (3.6\text{f}), (\text{A.11e}), \\
& s_{ez}^k \leq \frac{N_z^t \sum_{i \in [I]} \gamma_{ezi}^k \mu_i^{ke}}{1 + \sum_{i \in [I]} \gamma_{ezi}^k \mu_i^{ke} + \sum_{j \in [I]} \gamma_{bzj}^k \mu_j^{kz}}, \quad k = t, \ldots, T, \, \forall z \in [Z], && \text{(A.16b)} \\
& s_{bz}^k \leq \frac{N_z^t \sum_{j \in I_{bz}} \gamma_{bzj}^k \mu_j^{kz}}{1 + \sum_{i \in I_e} \gamma_{ezi}^k \mu_i^{ke} + \sum_{j \in I_{bz}} \gamma_{bzj}^k \mu_j^{kz}}, \quad k = t, \ldots, T, \, \forall z \in [Z], && \text{(A.16c)} \\
& \mu_i^{ke} \in \{0, 1\}, \mu_i^{kz} \qquad k = t, \ldots, T, \, \forall z \in [Z], \forall i \in [I] && \text{(A.16d)}
\end{aligned}$$

Hence, it suffices to prove the following lemma.

Lemma EC.1. *Under a multinomial logit demand model, $V_D^t(x^t) = \bar{V}_2^*$, where*

$$\begin{aligned}
\bar{V}_2^* = \underset{\mu, s, y, u, g, h}{\text{maximize}} \quad & \sum_{k=t}^{T} \sum_{z \in [Z]} \sum_{m=e,b} \sum_{i \in [I]} p_i s_{mzi}^k - \sum_{z \in [Z]} c_{ez} y_{ez} - \sum_{z, z' \in [Z]} c_{zz'} y_{zz'} + q \left(u_e + \sum_{z \in [Z]} u_{bz}\right) && \text{(A.17a)} \\
\text{subject to} \quad & \text{Constraints } (3.6\text{c}) - (3.6\text{f}), (\text{A.11e}), \\
& \sum_{i \in [I]} s_{mzi}^k \leq N_z^t \sum_{i \in [I]} \gamma_{mzi}^k h_{mzi}^k, \qquad k = t, \ldots, T, \, m = e, b, \, \forall z \in [Z], && \text{(A.17b)} \\
& s_{mzi}^k \leq N_z^k \gamma_{mzi}^k h_{mzi}^k, \qquad k = t, \ldots, T, \, m = e, b, \, \forall z \in [Z], \forall i \in [I], && \text{(A.17c)} \\
& \sum_{i \in [I]} h_{mzi}^k = g_z^k, \qquad k = t, \ldots, T, \, m = e, b, \, \forall z \in [Z], && \text{(A.17d)} \\
& h_{ezi}^k \leq \mu_i^{ke}, \qquad k = t, \ldots, T, \, \forall z \in [Z], \forall i \in [I], && \text{(A.17e)} \\
& h_{bzj}^k \leq \mu_j^{kz}, \qquad k = t, \ldots, T, \, \forall z \in [Z], \forall j \in [I], && \text{(A.17f)} \\
& g_z^k + \sum_{i \in I_e} \gamma_{ezi}^k h_{ezi}^k + \sum_{j \in I_{bz}} \gamma_{bzj}^k h_{bzj}^k = 1 \qquad k = t, \ldots, T, \, \forall z \in Z, && \text{(A.17g)} \\
& g \geq 0, \, h \geq 0, \, \mu \in \{0, 1\} && \text{(A.17h)}
\end{aligned}$$

Proof. Let us denote by $(\hat{\mu}, \hat{s}, \hat{y}, \hat{u})$ as the optimal solution to (A.16) with objective value $V_D^t(x^t)$. Defining

$$
\begin{aligned}
&s_{ezi}^k = \hat{\mu}_i^{ke}\hat{s}_{ez}^k, && k=t,\ldots,T,\ \forall z\in[Z],\ \forall i\in[I],\\
&s_{bzj}^k = \hat{\mu}_j^{kz}\hat{s}_{bz}^k, && k=t,\ldots,T,\ \forall z\in[Z],\ \forall j\in[I],\\
&g_z^k = \frac{1}{1+\sum_{i\in[I]}\gamma_{ezi}^k\hat{\mu}_i^{ke}+\sum_{j\in[I]}\gamma_{bzj}^k\hat{\mu}_j^{kz}}, && k=t,\ldots,T,\ \forall z\in[Z],\\
&h_{ezi}^k = \hat{\mu}_i^{ke}g_z^k, && k=t,\ldots,T,\ \forall z\in[Z],\ \forall i\in[I],\\
&h_{bzj}^k = \hat{\mu}_j^{kz}g_z^k, && k=t,\ldots,T,\ \forall z\in[Z],\ \forall j\in[I],
\end{aligned}
$$

we can easily check that $(\hat{\mu}, s, \hat{y}, \hat{u}, g, h)$ is a feasible solution to the mixed integer program (A.17) and achieves the objective value $V_D^t(x^t)$. Hence, $V_D^t(x^t) \le \bar{V}_2^*$.

We next show that $V_D^t(x^t) \ge \bar{V}_2^*$. Let us denote the maximizer of model (A.17) as $(\bar{\mu}, \bar{s}, \bar{y}, \bar{u}, \bar{g}, \bar{h})$, which achieves the optimal value $\bar{V}_2^*$. From the binary constraints of (A.17), and from constraints (A.11e), it follows that for time k, there exist price indices $(i_{ke}, j_{kz_1}, \ldots, j_{kz_n})$ such that:

$$
\bar{\mu}_{i_{ke}}^{ke} = 1, \qquad \bar{\mu}_i^{ke} = 0, \quad \forall i \ne i_{ke}, \tag{A.18}
$$

$$
\bar{\mu}_{j_{kz}}^{kz} = 1, \quad \forall z\in[Z], \qquad \bar{\mu}_j^{kz} = 0, \quad \forall j \ne j_{kz}, \forall z\in[Z]. \tag{A.19}
$$

From constraints (A.18) and (A.17c), it follows that $\bar{s}_{ezi}^k = 0$ for all $i \ne i_{ke}$, $z\in[Z]$. Similarly, from constraints (A.19) and (A.17c), it follows that $\bar{s}_{bzj}^k = 0$ for all $j \ne j_{kz}$, $z\in[Z]$. We set $s_{ez}^k = \bar{s}_{ezi_{ke}}^k$ and $s_{bz}^k = \bar{s}_{bzj_{kz}}^k$ for all $z\in[Z]$. Thus, it is easy to check based on the definitions (A.18)–(A.19) that:

$$
\bar{s}_{ezi}^k = \bar{\mu}_i^{ke}s_{ez}^k, \qquad k=t,\ldots,T,\ \forall z\in[Z],\ \forall i\in[I], \tag{A.20}
$$

$$
\bar{s}_{bzj}^k = \bar{\mu}_j^{kz}s_{bz}^k, \qquad k=t,\ldots,T,\ \forall z\in[Z],\ \forall j\in[I], \tag{A.21}
$$

From constraints (A.18) and (A.17e), it follows that $\bar{h}_{ezi}^k = 0$ for all $i \ne i_{ke}$, $z\in[Z]$. Similarly, from constraints (A.19) and (A.17f), it follows that $\bar{h}_{bzj}^k = 0$ for all $j \ne j_{kz}$, $z\in[Z]$. Thus, these imply from constraints (A.17d) that $\bar{h}_{ezi_{ke}}^k = \bar{g}_z^k$ for all $z\in[Z]$, and that $\bar{h}_{bzj_{kz}}^k = \bar{g}_z^k$ for all $z\in[Z]$. Thus, it is easy to check based on the definitions (A.18)–(A.19) that we have the following relationships:

$$
\bar{h}_{ezi}^k = \bar{\mu}_i^{ke}\bar{g}_z^k, \qquad k=t,\ldots,T,\ \forall z\in[Z],\ \forall i\in[I], \tag{A.22}
$$

$$
\bar{h}_{bzj}^k = \bar{\mu}_j^{kz}\bar{g}_z^k, \qquad k=t,\ldots,T,\ \forall z\in[Z],\ \forall j\in[I], \tag{A.23}
$$

From (A.22)–(A.23), and from (A.17g), it follows that $1 = \bar{g}_z^k + \gamma_{ezi_{ke}}^k\bar{g}_z^k + \gamma_{bzj_{kz}}^k\bar{g}_z^k$. Thus,

$$
\bar{g}_z^k = \frac{1}{1+\gamma_{ezi_{ke}}^k+\gamma_{bzj_{kz}}^k} = \frac{1}{1+\sum_{i\in[I]}\gamma_{ezi}^k\bar{\mu}_i^{ke}+\sum_{j\in[I]}\gamma_{bzj}^k\bar{\mu}_j^{kz}}. \tag{A.24}
$$

Substituting (A.22)–(A.23) into (A.17b), and using the relationship (A.24), we have

$$
s_{ez}^t \le N_z^k\gamma_{ezi_{ke}}^k\bar{g}_z^k = \frac{N_z^k\sum_{i\in[I]}\gamma_{ezi}^k\bar{\mu}_i^{ke}}{1+\sum_{i\in[I]}\gamma_{ezi}^k\bar{\mu}_i^{ke}+\sum_{j\in[I]}\gamma_{bzj}^k\bar{\mu}_j^{kz}}, \tag{A.25}
$$

$$
s_{bz}^t \le N_z^k\gamma_{bzj_{kz}}^k\bar{g}_z^k = \frac{N_z^k\sum_{j\in[I]}\gamma_{bzj}^k\bar{\mu}_j^{kz}}{1+\sum_{i\in[I]}\gamma_{ezi}^k\bar{\mu}_i^{ke}+\sum_{j\in[I]}\gamma_{bzj}^k\bar{\mu}_j^{kz}}. \tag{A.26}
$$

Thus, (A.25)–(A.26) proves that $(s, \bar{y}, \bar{u}, \bar{\mu})$ is feasible for model (A.16). Moreover, since we have (A.20)–(A.21), then this solution achieves the objective value $\bar{V}_2^*$. Thus, $\bar{V}_2^* \le V_D^t(x^t)$, proving the lemma. $\square$

EC.1.3. Proof of Proposition 1

If w and p are fixed, then (3.9) is a linear program. Hence, by strong LP duality, $U^t(w,p;x^t)$ is equivalent to a minimization LP. Thus, the worst-case retailer profit is equivalent to:

$$
\begin{aligned}
U_R^t(p;x^t) = \quad & \underset{\alpha,\beta,w}{\text{minimize}} \quad \sum_{k=t}^{T}\sum_{z\in[Z]}\sum_{m=e,b} d_{mz}^k(p_e^k,p_{bz}^k)\left(1+\delta_{mz}^k w_{mz}^k\right)\alpha_{mz}^k + \beta_e x_e^t + \sum_{z\in[Z]}\beta_{bz}x_{bz}^t && \text{(A.27a)}\\
& \text{subject to} \quad \alpha_{bz}^k + \beta_{bz}^k \geq p_{bz}^k, \qquad k=1,\ldots,T,\ \forall z\in[Z], && \text{(A.27b)}\\
& \alpha_{ez}^k + \beta_e^k \geq p_e^k - c_{ez}, \qquad k=1,\ldots,T,\ \forall z\in[Z], && \text{(A.27c)}\\
& \alpha_{ez}^k + \beta_{bz'}^k \geq p_e^k - c_{z'z}, \qquad k=1,\ldots,T,\ \forall z,z'\in[Z], && \text{(A.27d)}\\
& \alpha \geq 0,\ \beta\geq q,\ w\in W_{\Gamma,\Delta}^t && \text{(A.27e)}
\end{aligned}
$$

where α,β are the variables in the dual of $U^t(w,p;x^t)$. To prove Proposition 1, we need the following result:

Lemma EC.2. *If $(\bar\alpha,\bar\beta,\bar w)$ is the optimal solution for (A.27), then $\bar\alpha_{bz}^k \leq p_{bz}^k - q$ and $\bar{\alpha_{ez}^k} \leq p_e^k - c_z^{\min} - q$ for all $z\in[Z]$, where $c_z^{\min} := \min\left\{c_{ez}, \min_{z'\in[Z]} c_{z'z}\right\}$.*

Proof. Since the demands are nonnegative for all realizations, then given $\bar\beta$, $\bar\alpha$ must take the smallest feasible value allowed by constraints (A.27b)–(A.27d). Thus, $\bar\alpha_{bz}^k = \max\left(0,\ p_{bz}^k - \bar\beta_{bz}\right)$, and $\bar\alpha_{ez}^k = \max\left(0,\ p_e^k - c_{ez} - \bar\beta_e, \max_{z'\in[Z]}\left(p_e^k - c_{z'z} - \bar\beta_{bz'}\right)\right)$. Since the $\bar\beta$ variables are bounded below by q, we have the following upper bounds for the $\bar\alpha$ variables: $\bar\alpha_{bz}^k \leq p_{bz}^k - q$ and $\bar\alpha_{ez}^k \leq p_e^k - c_z^{\min} - q$ for all $z\in[Z]$. □

Note that α variables are the shadow prices for the demand constraints of $U^t(w,p;x^t)$, while the β variables are the shadow prices for its inventory constraints. Thus the upper bounds in Lemma EC.2 are natural 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). Moreover, due to Assumption 1, these upper bounds on α are nonnegative.

Optimization problem (A.27) is nonconvex due to the bilinear term $\alpha^\top w$ in the objective. Note however that the w variables are bounded between -1 and 1. Moreover, if we define the parameters $p^{\max} := \max\{p : p\in\Omega\}$, $A_{bz}^k := p^{\max} - q$, and $A_{ez}^k := p^{\max} - c_z^{\min} - q$, then from Lemma EC.2, we can add the constraints $\alpha_{mz}^k \leq A_{mz}^k$ to (A.27) without changing its optimal value. Hence the α variables are also bounded. Therefore, using these bounds on w and α, we can "lift" the optimization problem (A.27) onto a higher dimensional space by introducing variables $\eta_{mz}^k = \alpha_{mz}^k w_{mz}^k$, which linearizes the objective. This results in a linear program whose optimal value $U_l^t(p;x^t)$ is a lower bound to $U_R^t(p;x^t)$, where

$$
\begin{aligned}
U_l^t(p;x^t) := \quad & \underset{\alpha,\beta,w,\nu,\eta}{\text{minimize}} \quad \sum_{k=t}^{T}\sum_{z\in[Z]}\sum_{m=e,b} d_{mz}^k(p_e^k,p_{bz}^k)\left(\alpha_{mz}^k+\delta_{mz}^k \eta_{mz}^k\right) + \beta_e x_e^t + \sum_{z\in[Z]}\beta_{bz}x_{bz}^t && \text{(A.28a)}\\
& \text{subject to} \quad \alpha_{bz}^k + \beta_{bz}^k \geq p_{bz}^k, \qquad k=1,\ldots,T,\ \forall z\in[Z], && \text{(A.28b)}\\
& \alpha_{ez}^k + \beta_e^k \geq p_e^k - c_{ez}, \qquad k=1,\ldots,T,\ \forall z\in[Z], && \text{(A.28c)}\\
& \alpha_{ez}^k + \beta_{bz'}^k \geq p_e^k - c_{z'z}, \qquad k=1,\ldots,T,\ \forall z,z'\in[Z], && \text{(A.28d)}\\
& \alpha \geq 0,\ \beta\geq q,\ w\in W_{\Gamma,\Delta}^t, && \text{(A.28e)}\\
& \left|\eta_{mz}^k\right| \leq \alpha_{mz}^k, \qquad m=e,b,\ k=1,\ldots,T,\ \forall z\in[Z], && \text{(A.28f)}\\
& \left|A_{mz}^k w_{mz}^k - \eta_{mz}^k\right| \leq A_{mz}^k - \alpha_{mz}^k, \qquad m=e,b,\ k=1,\ldots,T,\ \forall z\in[Z] && \text{(A.28g)}
\end{aligned}
$$

Note that constraints (A.28f), (A.28g) and (3.7) can be linearized. The new constraints (A.28f) and (A.28g) are valid inequalities that are satisfied by any feasible solution to the linearized problem. We next derive the form (3.10) in Proposition 1. Note that due to strong duality, we can reformulate (A.28) as:

$$U_l^t(p;x^t) = \underset{\substack{s,y,u,r,R,g,G,\\ h,H,\phi,\Phi,l,L,f}}{\text{maximize}} \quad \sum_{k=t}^{T}\sum_{z\in[Z]}\left(p_{bz}^k s_{bz}^k + \left(p_e^k - c_{ez}\right) y_{ez}^k + \sum_{z'\in[Z]}\left(p_e^k - c_{z'z}\right) y_{z'z}^k\right) \tag{A.29a}$$

$$- \sum_{k=t}^{T}\left(\Gamma^k f^k + \sum_{z\in[Z]}\sum_{m=e,b}\left(\phi_{mz}^k + \Phi_{mz}^k\right) + \sum_{z\in[Z]}\sum_{m=e,b} A_{mz}^k\left(g_{mz}^k + G_{mz}^k\right)\right) \tag{A.29b}$$

$$- \sum_{k=t}^{T}\Delta^k\left(l^k + L^k\right) + q\left(u_e + \sum_{z\in[Z]} u_{bz}\right) \tag{A.29c}$$

subject to

$$s_{bz}^k + r_{bz}^k + R_{bz}^k - g_{bz}^k - G_{bz}^k \le d_{bz}^k(p_e^k, p_{bz}^k), \qquad k=t,\dots,T,\ \forall z\in[Z], \tag{A.29d}$$

$$y_{ez}^k + \sum_{z'\in[Z]} y_{z'z}^k + r_{ez}^k + R_{ez}^k - g_{ez}^k - G_{ez}^k \le d_{ez}^k(p_e^k, p_{bz}^k), \qquad k=t,\dots,T,\ \forall z\in[Z], \tag{A.29e}$$

$$\sum_{k=t}^{T}\sum_{z\in[Z]} y_{ez}^k + u_e = x_e^t, \tag{A.29f}$$

$$\sum_{k=t}^{T} s_{bz}^k + \sum_{k=t}^{T}\sum_{z'\in[Z]} y_{zz'}^k + u_{bz} = x_{bz}^t, \qquad \forall z\in[Z], \tag{A.29g}$$

$$\phi_{mz}^k - \Phi_{mz}^k + h_{mz}^k - H_{mz}^k - A_{mz}^k\left(g_{mz}^k - G_{mz}^k\right) + a_{mz}^k\left(l^k - L^k\right) = 0, \qquad m=e,b,\ k=t,\dots,T,\ \forall z\in[Z], \tag{A.29h}$$

$$- f^k + h_{mz}^k + H_{mz}^k = 0, \qquad m=e,b,\ k=t,\dots,T,\ \forall z\in[Z], \tag{A.29i}$$

$$r_{mz}^k - R_{mz}^k + g_{mz}^k - G_{mz}^k = \delta_{mz}^k d_{mz}^k(p_e^k, p_{bz}^k),, \tag{A.29j}$$

$$s, y, u, r, R, g, G, h, H, \phi, \Phi, l, L \ge 0. \tag{A.29k}$$

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,

$$U_l^t(p;x^t) = \underset{\substack{s,y,u,r,R,g,G,\\ h,H,\phi,\Phi,l,L,f}}{\text{maximize}} \quad \sum_{k=t}^{T}\sum_{z\in[Z]}\left(p_e^k s_{ez}^k + p_{bz}^k s_{bz}^k\right) - \sum_{z\in[Z]}\left(c_{ez} y_{ez} + \sum_{z'\in[Z]} c_{z'z} y_{z'z}\right) \tag{A.30a}$$

$$- \sum_{k=t}^{T}\left(\Gamma^k f^k + \sum_{z\in[Z]}\sum_{m=e,b}\left(\phi_{mz}^k + \Phi_{mz}^k\right) + \sum_{z\in[Z]}\sum_{m=e,b} A_{mz}^k\left(g_{mz}^k + G_{mz}^k\right)\right) \tag{A.30b}$$

$$- \sum_{k=t}^{T}\Delta^k\left(l^k + L^k\right) + q\left(u_e + \sum_{z\in[Z]} u_{bz}\right) \tag{A.30c}$$

subject to

Constraints (3.6c)–(3.6f), (A.29h)–(A.29k),

$$s_{mz}^k + r_{mz}^k + R_{mz}^k - g_{mz}^k - G_{mz}^k \le d_{mz}^k(p_e^k, p_{bz}^k), \quad m=e,b,\ k=t,\dots,T,\ z\in[Z] \tag{A.30d}$$

Let us introduce the following variable transformations:

$$
\begin{aligned}
\chi^k_{mz} &= \phi^k_{mz} - \Phi^k_{mz}, & X^k_{mz} &= \phi^k_{mz} + \Phi^k_{mz},\\
\upsilon^k_{mz} &= g^k_{mz} - G^k_{mz}, & \Lambda^k_{mz} &= g^k_{mz} + G^k_{mz},\\
\psi^k_{mz} &= r^k_{mz} - R^k_{mz}, & \Psi^k_{mz} &= r^k_{mz} + R^k_{mz},\\
\upsilon^k_{mz} &= h^k_{mz} - H^k_{mz}, & \Upsilon^k_{mz} &= h^k_{mz} + H^k_{mz},\\
\vartheta^k &= l^k - L^k, & \Theta^k &= l^k + L^k.
\end{aligned}
$$

Therefore, by replacing these new variables into model (A.30), we have its equivalent formulation:

$$
\begin{aligned}
U^t_l(p;x^t) = \underset{\substack{s,y,u,\chi,X,\upsilon,\Lambda,\\ \psi,\Psi,\upsilon,\Upsilon,\vartheta,\Theta,f}}{\text{maximize}} \quad & \sum_{k=t}^{T}\sum_{z\in[Z]}\left(p^k_e s^k_{ez} + p^k_{bz}s^k_{bz}\right) - \sum_{z\in[Z]}\left(c_{ez}y_{ez} + \sum_{z'\in[Z]} c_{z'z}y_{z'z}\right) && \text{(A.31a)}\\
& -\sum_{k=t}^{T}\left(\Gamma^k f^k + \Delta^k\Theta^k + \sum_{z\in[Z]}\sum_{m=e,b}\left(X^k_{mz} + A^k_{mz}\Lambda^k_{mz}\right)\right) + q\left(u_e + \sum_{z\in[Z]} u_{bz}\right) && \text{(A.31b)}\\
\text{subject to} \quad & \text{Constraints (3.6c)–(3.6f)}, && \text{(A.31c)}\\
& s^k_{mz} \le d^k_{mz}(p^k_e, p^k_{bz}) + \Lambda^k_{mz} - \Psi^k_{mz}, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31d)}\\
& \chi^k_{mz} + \upsilon^k_{mz} - A^k_{mz}\upsilon^k_{mz} + a^t_{mz}\vartheta^k = 0, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31e)}\\
& \Upsilon^k_{mz} = f^k, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31f)}\\
& \upsilon^k_{mz} + \psi^k_{mz} = \delta^k_{mz} d^k_{mz}(p^k_e, p^k_{bz}), \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31g)}\\
& X^k_{mz} \ge |\chi^k_{mz}|, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31h)}\\
& \Lambda^k_{mz} \ge |\upsilon^k_{mz}|, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31i)}\\
& \Psi^k_{mz} \ge |\psi^k_{mz}|, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31j)}\\
& \Upsilon^k_{mz} \ge |\upsilon^k_{mz}|, \quad m=e,b,\ k=t,\dots,T,\ z\in[Z], && \text{(A.31k)}\\
& \Theta^k \ge |\vartheta^k|, \quad k=t,\dots,T && \text{(A.31l)}
\end{aligned}
$$

Note that since $\Upsilon^k_{mz} = f^k$, then we can eliminate the Υ^k_{mz} variables by replacing the constraint (A.31k) by $f^k \ge |\upsilon^k_{mz}|$. By equation (A.31e), we know $\upsilon^k_{mz} = A^k_{mz}\upsilon^k_{mz} - \chi^k_{mz} - a^t_{mz}\vartheta^k$. Due to the maximizing objective, in the optimal solution, we have $X^k_{mz} = |\chi^k_{mz}|$ and $\Theta^k = |\vartheta^k|$. Next, because the coefficient of s^k_{mz} in the objective is positive and Ψ^k_{mz} reduces the upper bound for s^k_{mz} without impacting the objective, an alternative optimal solution is obtained when Ψ^k_{mz} is reduced and set equal to $|\psi^k_{mz}|$. Lastly, note that the coefficient of Λ^k_{mz} is negative, and it increases the bound on s^k_{mz} through equation (A.31d). However, since A^k_{mz} 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^k_{mz} to the smallest feasible value and incur the least penalty A^k_{mz}. Hence, in the optimal solution $\Lambda^k_{mz} = |\upsilon^k_{mz}|$. Hence, we have that model (A.31) is equivalent to (3.10). □

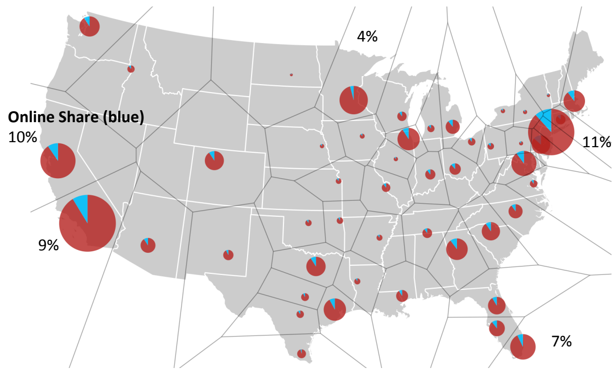

Figure EC.3 **Distribution of sales over 50 zones for a product category. Sales volume is proportional to the pie size. The pie in each zone shows the relative frequency of brick-and-mortar sales and e-commerce sales.**

EC.1.4. Proof of Lemma 4

Consider a clairvoyant who knows the future demand factors $w = (w_{mz})_{mz}$ prior to making decisions, and thus is able to earn the highest profit with his or her price, sales, and fulfillment decisions by solving (5.4). On each sample path w, the "perfect foresight" realized profit $U_{PF}(w)$ is an upper bound on the realized profit of any pricing policy. Hence, the expected perfect foresight profit V_{PF} is an upper bound on the optimal expected profit. $\square$

EC.2. 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 the MNL function Eq. (3.4) 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:

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

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

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

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.2.1. Estimation procedure.

The γ and β coefficients in Eqs. (B.1–B.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

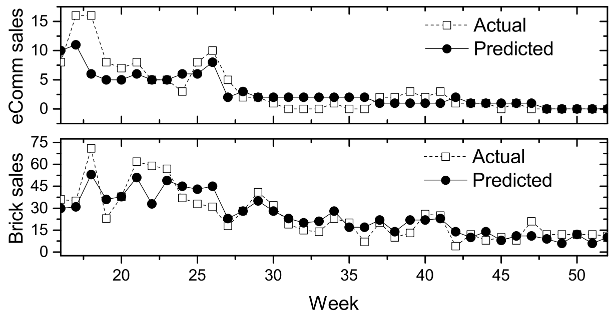

Figure EC.4 Predicted and actual e-commerce sales, and brick-and-mortar sales for a specific SKU-zone.

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.2.2. 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 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%

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

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.

EC.2.3. 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).