

Dynamic Procurement of New Products with Covariate Information: The Residual Tree Method

Gah-Yi Ban

London Business School, Regent's Park, London, NW1 4SA, United Kingdom.
gban@london.edu

Jérémie Gallien

London Business School, Regent's Park, London, NW1 4SA, United Kingdom.
jgallien@london.edu

Adam Mersereau

Kenan-Flagler Business School, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 USA.
ajm@unc.edu

Problem definition: We study the practice-motivated problem of dynamically procuring a new, short life-cycle product under demand uncertainty. The firm does not know the demand for the new product but has data on similar products sold in the past, including demand histories and covariate information such as product characteristics.

Academic/practical relevance: The dynamic procurement problem has long attracted academic and practitioner interest, and we solve it in an innovative data-driven way with proven theoretical guarantees. This work is also the first to leverage the power of covariate data in solving this problem.

Methodology: We propose a new, combined forecasting and optimization algorithm called the Residual Tree method, and analyze its performance via epi-convergence theory and computations. Our method generalizes the classical Scenario Tree method by using covariates to link historical data on similar products to construct demand forecasts for the new product.

Results: We prove, under fairly mild conditions, that the Residual Tree method is asymptotically optimal as the size of the data set grows. We also numerically validate the method for problem instances derived using data from the global fashion retailer Zara. We find that ignoring covariate information leads to systematic bias in the optimal solution, translating to a 6–15% increase in the total cost for the problem instances under study. We also find that solutions based on trees using just 2–3 branches per node, which is common in the existing literature, are inadequate, resulting in 30–66% higher total costs compared with our best solution.

Managerial implications: The Residual Tree is a new and generalizable approach that uses past data on similar products to manage new product inventories. We also quantify the value of covariate information and of granular demand modeling.

Key words: new product, inventory management, data-driven operations, Scenario Tree method, Residual Tree method, demand uncertainty

History: Original: February 24, 2017. This version: March 28, 2018

1. Introduction

Matching supply with demand is a critical challenge for retailers selling short life cycle products in markets with fast-changing demand trends, such as fashion apparel and consumer electronics. In response, retailers and academics have considered specific operational strategies, such as quick response and fast fashion, for facilitating this matching. Such strategies make use of many levers at the firm's disposal, including product design (e.g., component commonality), forecasting and market tests (e.g., early order incentives), manufacturing locations (e.g., on-shoring and near-shoring), manufacturing sequencing (e.g., delayed differentiation), pricing (e.g., clearance or in-season mark-downs), and logistics (Fisher et al. 2001).

Another recognized strategy involves the dynamic allocation of demand with different levels of uncertainty to manufacturing facilities characterized by different lead times and costs. The key idea is to selectively use fast and expensive supply sources after having acquired as much information as possible about products with high initial demand uncertainty (risk-based supply portfolio strategy; see Fisher and Raman 1996) or to supply the portion of demand that is associated with more uncertainty (base/surge demand strategy; Allon and Van Mieghem 2010).

This paper is motivated by our interaction with the Spain-based apparel retailer Zara, which has implemented such a procurement segmentation approach in its supply chain. Specifically, Zara typically purchases some initial quantity of a new product from China several months before the season, but depending on sales during the season it can subsequently place one of several in-season replenishment orders for the same product from more expensive but closer vendors located in places like Turkey or Portugal (Patel 2012). A key related challenge reported to us by Zara managers is the related decision complexity. For example, it was often unclear to them exactly how much of a given article should initially be ordered from China as a function of the different costs and lead times associated with all capable suppliers worldwide and of the anticipated decrease in forecast uncertainty once the selling season starts.

The second and third authors of this paper have been helping Zara adapt existing stochastic programming approaches to address this dynamic procurement challenge and to implement this work as part of the industrial decision support system shown in Figure 1. This system is designed for use by a fashion buyer charged with making procurement decisions for a new product. It includes tools for the buyer to visualize and edit the demand forecasts (line chart in Figure 1) and to see, edit, and optimize the procurement plan (bar chart and the right-hand-side table in Figure 1). Field testing of this system with Zara buyers has revealed additional critical challenges, however. Firstly, the system currently relies on Zara's legacy process for forecasting demand of new products, which

requires a selection by the user of several past articles from previous seasons deemed “comparable” to the new one under consideration. As buyers realized that this choice of prior comparable articles was a subjective one often having a substantial impact on purchasing recommendations, they started asking for automated suggestions of which comparable articles to use. Secondly, the system currently relies on an *ad hoc* model of demand learning dynamics which requires time and expertise to estimate and maintain. For this reason, the learning model is shared by all articles within the same product subfamily (e.g., women’s shirts), which for Zara can comprise a relatively large set of products. As a result, a relatively well-understood article similar to ones sold in previous seasons, and thus having little demand uncertainty, could get the same recommended initial purchase quantity from China as a fashion-forward and untested article for which much demand uncertainty is resolved after the first weeks of sales are observed. This results in suboptimal decisions; for example, the benefits of a second purchase from Turkey or Portugal are lower for the well-understood garment than for the fashion-forward one so that, all else being equal, the well-understood garment should have a larger initial purchase quantity from China.

The present study lays the foundations for addressing these challenges as part of a more data-driven, accurate and automated version of the existing decision support system shown in Figure 1. Although we hope to help Zara implement such an enhanced system in the future as part of our ongoing relationship with the firm, we believe that the relevance of this foundation work extends much beyond Zara. Specifically, this paper presents, analyzes and tests a practice-oriented, generalizable and tractable data-driven computational method for optimizing the procurement of a new, short life cycle product from multiple sources over time. As there is no historical data for a new product, the key to our solution is the use of covariate (also known as feature/attribute/predictor) information to link the new product to similar products that were sold in the past. Relevant product covariates for apparel articles may include the procurement cost, retail price, colour, item type (e.g., sweater, t-shirt, etc.), fabric, design style (e.g., sporty, classic), expert predictions of the product’s popularity, etc. We refer to covariates tied to product attributes that do not change over time as “static.” Our approach can also account for “dynamic” covariates—e.g., lagged demand observations included to naturally account for demand auto-correlations. Importantly, covariates are essential for accounting for product heterogeneity by providing a means for aligning demand observations across different products. This allows us to construct a large number of estimated demand trajectories relevant to the new product, which we can then use in a scenario-based model of forecast updating dynamics.

Our method employs a new approach which we call the Residual Tree method, which can be used to solve general multi-stage stochastic programs where there is no information about the

underlying uncertainty, but there are data on historical data trajectories of other stochastic processes related through covariate information. The Residual Tree method extends the classical Scenario Tree method of stochastic programming (see Shapiro et al. 2009), which relies on scenario samples from the underlying stochastic process of interest (in our case, demand for the new product). In contrast, our approach replaces these scenario paths with “residual paths”, composed of residuals from a regression model relating the uncertainty in question and the available covariate information.

We show that our Residual Tree method can solve, to an arbitrary accuracy, a generic version of the dynamic procurement problem by following three steps: (i) a least-squares or lasso regression providing demand forecasts capturing the covariate information; (ii) binning of output residuals for improved computational efficiency; and (iii) solving a large linear program in which the new product demand is approximated by the binned residual paths. We prove, under natural assumptions, the method just described is asymptotically optimal as the number of similar products increases.

Lastly, we numerically validate the Residual Tree method on stylized problems based on data provided by Zara. The numerical results are promising in that the Residual Tree method finds solutions quickly that have near-optimal finite-sample performance. Our numerical results further show the importance of covariate information. For example, we find that ignoring static covariates results in out-of-sample cost increases on the order of 6-15%, with statistical significance at the 1% level in a four-period problem. We also find that considering just 2-3 branches per node for the dynamic procurement problem, as is the case in much of the literature to date, is too simplistic and may result in solutions that underperform substantially, with additional costs of 30 – 66% of the optimal total cost compared with our solution with 10 branches per node, also with statistical significance at the 1% level.

Our approach integrates regression with multi-stage stochastic programming in a way that is intuitive while retaining desirable theoretical properties such as asymptotic optimality. The approach is practical in the sense that the main input is data a retailer is likely to have. While we assume a linear relationship between a product’s demand and relevant covariates, we do not impose any distributional assumptions on the residuals. We believe this work is the first to leverage the power of covariate information in multi-stage procurement optimization.

We discuss the related literature and our contributions in Section 2, the optimization model in Section 3, our proposed solution approach in Section 4, asymptotic optimality results in Section 5, and numerical experiments in Section 6. We conclude in Section 7. The proofs of all mathematical results are placed in the online Appendix. Some of the data presented in this paper have been disguised to protect its confidentiality, and we emphasize that the views presented in this paper do not necessarily represent those of the Inditex Group.

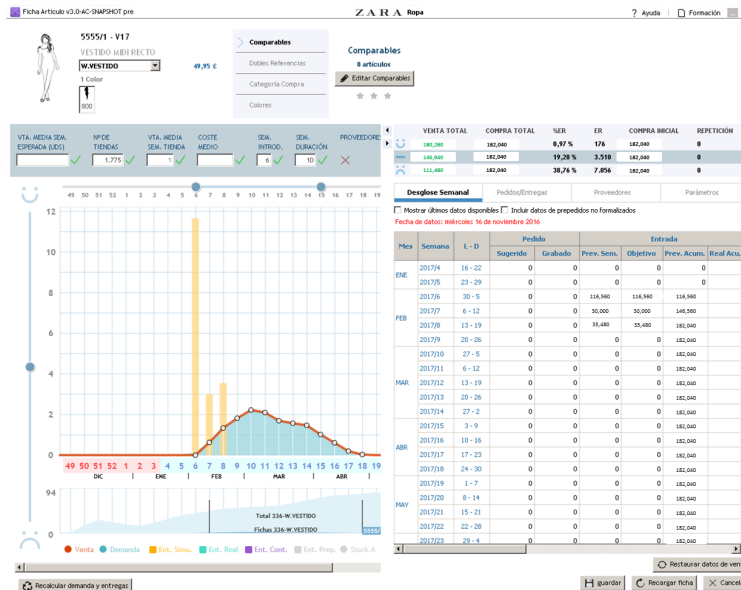


Figure 1: Graphical user interface component used by Zara as part of a decision support system adapted to the dynamic procurement problem in its environment.

2. Literature Review

This paper relates to the literature on production/inventory systems investigating how one or multiple supply sources should be used when demand forecast updates become available over a finite time horizon. As discussed in the recent survey by Serel (2016), these feature a variety of modeling approaches, including multi-period extensions of the newsvendor model (Lau and Lau 1996), Bayesian updating (Eppen and Iyer 1997, Burnetas and Gilbert 2001), the martingale model of forecast evolution (Iida and Zipkin 2006, Lu et al. 2006, Wang et al. 2012) and the forecast band refinement model (Kaminsky and Swaminathan 2001, 2004). Within this body of work, our paper is closest to the practice-oriented studies motivated by collaboration with industry, such as Fisher and Raman (1996), Fisher et al. (2001), Jones et al. (2001) and Peng et al. (2012). We also highlight Escudero et al. (1993) and Higle and Kempf (2011), which propose computational approaches based on stochastic programming to address production/inventory problems with demand uncertainty. Finally, our work is closely related to Agrawal et al. (2002), which is a practice-based study reporting the use of stochastic programming to address a dynamic procurement problem arising in a retail environment. In contrast to all papers above, our approach is data-driven, captures a potentially large amount of covariate information, and is justified by an asymptotic optimality result.

This paper also contributes to the growing literature on data-driven approaches to operations management challenges, which seeks to inform operational decisions using available historical data

and minimal distributional assumptions. There are now many studies on data-driven inventory management (e.g. Burnetas and Smith 2000, Huh and Rusmevichientong 2009, Kunnumkal and Topaloglu 2008, Godfrey and Powell 2001, Levi et al. 2007, Levi et al. 2015, Ban and Rudin (2018), Ban 2018 and references therein). A number of recent studies are further motivated by the availability of large data sets that contain covariate data (also referred to as features, characteristics, attributes or explanatory variables), as is this paper. These include Ban and Rudin (2018), Ferreira et al. (2015), Chen et al. (2015), Cohen et al. (2016), Qiang and Bayati (2016), Hu et al. (2016), Ban and Keskin (2017) and Baardman et al. (2017). Our work is distinct from the data-driven operations management literature just discussed, which considers covariate information for single-period optimization problems only. In contrast, we focus on the dynamic procurement problem, which requires analyzing a more complex multi-stage stochastic program.

From a methodological standpoint, our Residual Tree method adds to the literature on Scenario Tree methods for solving multi-stage stochastic programs. We refer the reader to Dupačová et al. (2000) for a review up to 2000, and Pennanen and Koivu (2002), Dupačová et al. (2003), Heitsch and Römisch (2003), Casey and Sen (2005), Pflug and Hochreiter (2003) and Rios et al. (2015) for more recent developments. In the classical Scenario Tree method, knowledge of the distribution of the underlying stochastic process is assumed, enabling the decision-maker to simulate arbitrary quantities of data from it. In contrast, our Residual Tree method is intended to be used when the decision-maker does not know the distribution of the underlying stochastic process of interest, but instead has finite observations of other stochastic processes that are related to the one in question through covariate information. In this sense the Residual Tree method is a covariate data-driven generalization of the classical Scenario Tree approach. Our proof of asymptotical optimality relies on the epi-convergence theory of Pennanen (2005) and Pennanen (2009) developed for classical Scenario Tree methods. We note that our proof must show that both discretization error (which also arises in the classical Scenario Tree method) and estimation error (which arises from our regression modeling of the data) vanish, and is therefore a nontrivial extension of past work.

Methodologically, the most closely related paper to ours is Rios et al. (2015), who use functional regression based on historical data to estimate a stochastic model of electricity demand, then use their model to generate scenarios for a Scenario Tree formulation of an electricity generation problem. There are a few key differences with our paper. First, Rios et al. (2015) assume limited data on the process of interest along with some temporal covariates (that are revealed over time) compared with our assumptions of no data on the process of interest and the existence of both static and temporal covariates. Second, our approach directly constructs demand scenarios from regression

residuals, whereas theirs first fits distributions to residuals and then samples from the distributions. Thirdly, Rios et al. (2015) does not include any theoretical guarantees, whereas we prove asymptotic optimality of our method.

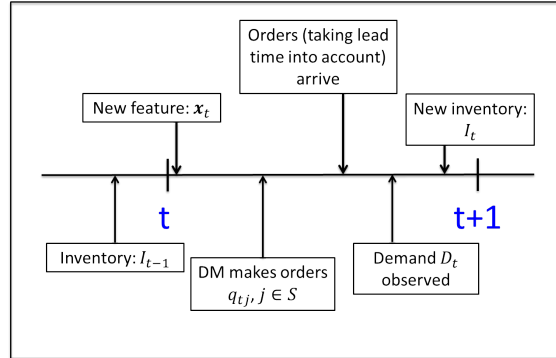


Figure 2: Timeline of events.

3. Model

We formulate our dynamic procurement problem in two steps. We first describe our assumed demand process and connect it to data that is likely to be available to a fashion retailer. We then formulate the stochastic optimization problem.

3.1. Demand Model and Description of Existing Data

Stochastic demand for a new product arrives in every period, described by a random vector $\{D_{0t}\}_{t=1}^T$ defined on the probability space $(\Xi_0, \mathcal{F}_0, \mathbf{P}_0)$, where each $D_{0t} : \Xi_{0t} \rightarrow \mathbb{R}_+$, $1 \leq t \leq T$ is a random variable on P_{0t} , and $\mathbf{P}_0 = P_{01} \times \dots \times P_{0T}$ is the product measure. (We use the index 0 to refer to the new, focal product; data from n other, historical products will be indexed by $k = 1, \dots, n$.) Note the demands up to time t induce a natural filtration $\mathcal{F}_{0t} = \sigma(\{D_{0s}^{-1}(A) : 1 \leq s \leq t, A \in \Omega\})$ on \mathcal{F}_0 .

We assume that the demand D_{0t} linearly depends on a m_t -dimensional random covariate vector \mathbf{X}_{0t} , where each element may consist of product features such as price, color and style code, as well as time-series factors such as seasonality and past demands. Note the subscript t in m_t is used to allow for the number of covariates to differ across periods. In other words, we assume the random demand D_{0t} for our focal product in period t is given by:

$$D_{0t} | (\mathbf{X}_{0t} = \mathbf{x}_{0t}) = \alpha_t + \beta_t^\top \mathbf{x}_{0t} + \varepsilon_t, \quad (1)$$

where $\alpha_t \in \mathbb{R}$ is a scalar intercept coefficient, $\beta_t \in \mathbb{R}^{m_t}$ is a vector of model coefficients, $\mathbf{x}_{0t} \in \mathbb{R}^{m_t}$ is a vector of m_t realized covariates at time t , and ε_t is the random error described shortly.

Since the product is new, the coefficients α_t and β_t and some components of \mathbf{x}_{0t} (e.g. past demands $d_{0(t-1)}$ if there are time-series dependencies) are unknown to the firm. However, in our problem setting the firm has access to a large amount of data on similar products sold in the past, using which it can forecast the demand for the new product. We now describe the demand model which makes this *cross-learning* possible.

Assume the following demand model for the family of $n + 1$ products (the focal product 0 plus n other similar products sold in the past):

$$D_{kt} | \mathbf{X}_{kt} = \alpha_t + \beta_t^\top \mathbf{X}_{kt} + \varepsilon_t, \quad k = 0, 1, \dots, n, \quad (2)$$

where α_t , β_t and ε_t are as in (1), and for each $k = 0, \dots, n$, \mathbf{X}_{kt} is an independent m_t -dimensional random covariate vector with a common distribution $P_{\mathbf{X}_t}$, where we use the notation \mathbf{X}_t to denote the generic random covariate vector. The random errors ε_t , $t = 1, \dots, T$ are assumed to be independent across time and of the feature vectors \mathbf{X}_{kt} , $k = 0, \dots, n$ and $t = 1, \dots, T$. They are also assumed to be drawn independently from a (common) continuous distribution with strictly increasing cumulative distribution function F_{ε_t} (i.e., an inverse cdf exists), with mean zero and variance σ_t^2 , for each $t = 1, \dots, T$. Finally, we assume that $\mathbb{E}[\mathbf{X}_t \mathbf{X}_t^\top]$, where $\mathbf{X}_t := [\mathbf{X}_{1t}, \dots, \mathbf{X}_{nt}]$, exists and is a symmetric positive definite matrix.

As for the data, we assume the firm has a database of past demands and covariates of n similar products sold in the past, $\mathcal{D} = \{(\mathbf{x}_{k1}, d_{k1})_{k=1}^n, \dots, (\mathbf{x}_{kT}, d_{kT})_{k=1}^n\}$, drawn from the demand model (2). We assume that $n^{-1} \sum_{k=1}^n \mathbf{X}_t \mathbf{X}_t^\top$ are finite symmetric positive definite matrices. At the beginning of the selling season, the firm also has the covariates for the new product, $[\hat{\mathbf{x}}_{01}, \dots, \hat{\mathbf{x}}_{0T}]$, where we use the $\hat{\cdot}$ notation to highlight the fact that some non-static elements of the covariate vector (e.g. past demands, weather or economic indicators) may be estimates of values to be revealed later (as the season unfolds, the firm will observe $\mathbf{X}_{0t} = \mathbf{x}_{0t}$, $t = 2, \dots, T$). (Note that if the demand only depends on static, product-characteristic covariates, then $\hat{\mathbf{x}}_{0t} = \mathbf{x}_{0t}$ for all $t = 1, \dots, T$.)

As an example, if the new product 0 were of a particular t-shirt design, the database \mathcal{D} would include the demand histories for other t-shirt products offered last season. For any product $k = 0, 1, \dots, n$ and any time period t , the feature vector \mathbf{x}_{kt} could include numerical representations of t-shirt k 's color, fabric, and design; the historical demand strength of t-shirts during the time of year represented by period t ; and the observed demand for t-shirt k in period $t - 1$. We remark that in describing the database \mathcal{D} we are not distinguishing between observed sales and demand.

Comments on the demand model (2) and the available data

1. *Linearity.* The linear demand model (2) is a widely-used and useful model that can arbitrarily approximate nonlinear models and is general enough to subsume many time-series models (e.g. autoregressive and martingale models such as the Martingale Model of Forecast Evolution (MMFE) of Heath and Jackson 1994 and Graves et al. 1986). We refer the readers to Ban and Rudin (2018) for details on how smooth nonlinear models can be approximated by linear ones through Taylor expansions and enlargement of the covariate dimension, and on how time-series models and MMFE models can be captured by the linear demand model.
2. *Assumptions on the residual error ε_t .* As the residual error is not product-dependent, the underlying assumption is that the demands for products within the same product family have the same distribution once adjusted for their means. For example, the demands for two different t-shirt designs both being normally distributed with the same variance but different means would satisfy the assumptions of our model. Thus the assumption of a common error distribution at each time $t = 1, \dots, T$ allows for cross-learning of the product demands. We believe this is a reasonable assumption, since all effects due to other factors are meant to be captured by the linear demand model, leaving a noise term with a common distribution across products.
3. *Assumptions on the design matrix \mathbf{X}_t .* The assumptions on $\mathbb{E}[\mathbf{X}_t \mathbf{X}_t^\top]$ and its empirical counterpart $n^{-1} \sum_{k=1}^n \mathbf{X}_t \mathbf{X}_t^\top$ are standard requirements for least-squares regression estimation to be possible. Violation of the assumptions (i.e. the matrices are not positive definite) implies the existence of redundant covariate dimension(s), and this can be dealt with by removing the problematic dimension(s) from consideration.
4. *Data-driven learning.* The inclusion of an arbitrary selection of covariates realized over time, including past demand observations, enables the learning component of our model. This learning would typically be reflected by an estimated positive correlation structure between the initial observations of weekly sales and sales subsequently predicted, and/or by an estimated decreasing variance pattern for the random error term ε_t over time. While a variety of alternative demand learning models have been used in past studies of related dynamic production/inventory problems (e.g., the Bayesian model in Burnetas and Gilbert 2001, the MMFE in Wang et al. 2012, the band refinement model in Kaminsky and Swaminathan 2001, 2004), we emphasize that our approach is data-driven. That is, the starting point assumed for the demand model (2), the combined estimation and optimization method to be presented in Section 4 and the related asymptotic optimality theory to be developed in Section 5 is a data set of historical observations, as opposed to exogenously given distributional knowledge.

3.2. Optimization Problem Formulation

We model the management of a single new product over a finite, T -period time horizon indexed in the forward direction by $t = 1, \dots, T$. In each period t , the firm may order some quantity of the product from a vendor chosen from a set \mathcal{S} of potential suppliers, letting $S = |\mathcal{S}|$. Orders from supplier j placed in period t then arrive after a deterministic lead time ℓ_j at a cost of c_{tj} .

The firm has initial inventory I_0 . Any unmet demand is assumed lost, incurring opportunity costs of b_t per unit. The firm also incurs a holding cost of h_t per unit for inventory held at the end of periods $t = 1, \dots, T - 1$. At the end of the horizon, any remaining units of inventory are salvaged, generating revenue v per unit. We assume $h_t > 0$ and $b_t \geq b_{t+1}$ for all $t = 1, \dots, T$ and $b_T > v$. The assumption $b_t \geq b_{t+1}$ is justified because retail prices, and hence stockout opportunity costs, typically decrease in time. The condition $b_T > v$ precludes giving the retailer an incentive to withhold stock from customers in period T in order to make more money in salvage.

Figure 2 illustrates the sequence of events within a period. Orders are chosen and placed with the covariate vector $\hat{\mathbf{x}}_{0t}$, but before past orders arrive and the new demand D_{0t} is observed; note this timeline allows for orders from suppliers with zero lead time to arrive in the same period the order is placed. Let q_{tj} denote the quantity ordered in period t from supplier j . We have $q_{tj} \in Q_{tj}$, where Q_{tj} is a bounded subset of \mathbb{R}_+ . We require $q_{tj} \in \mathcal{F}_{0t}$, i.e. adapted to the natural filtration of the demand process up to time t .

The multi-stage stochastic procurement problem we wish to solve is thus:

$$\min_{\substack{q_{tj}, I_t, l_t \\ 1 \leq t \leq T, 1 \leq j \leq S}} \mathbb{E} \left[\sum_{t=1}^{T-1} h_t I_t + \sum_{t=1}^T b_t l_t + \sum_{t=1}^T \sum_{j=1}^S c_{tj} q_{tj} - v I_T \right] \quad (\text{PP})$$

s.t.

$$q_{tj} \in \mathcal{F}_{0t}, \quad \forall 1 \leq t \leq T, \quad \forall 1 \leq j \leq S, \quad (\text{PPa})$$

$$I_t = \left(I_{t-1} + \sum_{j \in \mathcal{S}} \sum_{\substack{1 \leq \tau \leq t-1: \\ \tau = t - \ell_j}} q_{\tau j} - D_{0t} \right)^+, \quad \forall 1 \leq t \leq T \quad (\text{PPb})$$

$$l_t = \left(D_{0,t} - I_{t-1} - \sum_{j \in \mathcal{S}} \sum_{\substack{1 \leq \tau \leq t: \\ \tau = t - \ell_j}} q_{\tau j} \right)^+, \quad \forall 1 \leq t \leq T \quad (\text{PPc})$$

$$q_{tj} \in Q_{tj}, \quad \forall 1 \leq t \leq T, \quad \forall 1 \leq j \leq S \quad (\text{PPd})$$

$$I_t \geq 0, \quad l_t \geq 0, \quad \forall 1 \leq t \leq T, \quad (\text{PPe})$$

where PP is short for “procurement problem” and the expectation is taken over the measure induced by $\{D_{01}, \dots, D_{0T}\}$. Note $\sum_{\tau=t-\ell_j}^{1 \leq \tau \leq t} q_{\tau j}$ denotes the total order quantity arriving from supplier j during period t , l_t the lost sales incurred during period t and I_t the inventory on-hand at the end of time t . The constraint (PPa) is a requirement that decisions are adapted to the history, (PPb) captures the inventory dynamics, (PPc) captures the lost sales dynamics and (PPd)–(PPE) define the domains of the decision variables. Note constraints (PPb) and (PPc) apply almost surely. We assume the parameters of problem (PP) are such that it is feasible and thus at least one optimal solution exists.

Denote the solution to (PP) by $(\mathbf{q}^*, \mathbf{I}^*, \mathbf{l}^*)$, where $\mathbf{q}^* = [q_{t1}^*, \dots, q_{tS}^*]_{t=1}^T$, $\mathbf{I}^* = [I_1^*, \dots, I_T^*]$ and $\mathbf{l}^* = [l_1^*, \dots, l_T^*]$. Denoting the available information at the beginning of period t by $\mathbf{H}_t := \{I_0, \dots, I_{(t-1)}, l_0, \dots, l_{(t-1)}, \mathbf{x}_{01}, \dots, \mathbf{x}_{0t}, [q_{\tau 1}^*, \dots, q_{\tau S}^*]_{\tau=1}^{(t-1)}\}$, we note that q_{tj}^* , $j = 1, \dots, S$, I_t^* and l_t^* are policies (functions) that map \mathbf{H}_t to the reals. Finally, we note I_0 , the initial inventory on-hand, is a constant assumed given.

Some features of optimization formulation (PP) are motivated by considerations in Zara’s environment. We allow for an arbitrary number of supply options with differing lead times, in keeping with the number and diversity of the vendors and shipping options available to Zara. While lead-times are assumed to be deterministic for tractability reasons, we understand the typical variability of vendor delivery times to be lower than, or at worst comparable to, our intended planning period of one week. We also emphasize that the assumption of boundedness imposed on sets Q_{tj} allows us to capture realistic supply constraints and cost structures such as limited vendor capacity, minimum order quantities, and quantity discount schemes. Finally, the assumed lost sales dynamics align with Zara’s observations of typical customer behaviors when confronted with stock-outs of fashion products. While these lost sales dynamics can present specific tractability challenges (Zipkin 2008), the following proposition establishes that the inventory and lost sales constraints (PPb)–(PPc) in (PP) have equivalent linear expressions in this case.

PROPOSITION 1. *The multi-stage stochastic program (PP) is equivalent to the following problem, in that the optimal solution(s) and the optimal objective values coincide:*

$$\min_{\substack{q_{tj}, I_t, l_t \\ 1 \leq t \leq T, 1 \leq j \leq S}} \mathbb{E} \left[\sum_{t=1}^{T-1} h_t I_t + \sum_{t=1}^T b_t l_t + \sum_{t=1}^T \sum_{j=1}^S c_{tj} q_{tj} - v I_T \right] \quad (\text{PP2})$$

s.t.

$$q_{tj} \in \mathcal{F}_{0t}, \quad \forall 1 \leq t \leq T, \quad \forall 1 \leq j \leq S, \quad (\text{PP2a})$$

$$I_t = I_{t-1} + l_t + \sum_{j \in \mathcal{S}} \sum_{\substack{1 \leq \tau \leq t-1: \\ \tau = t - \ell_j}} q_{\tau j} - D_{0t}, \quad \forall 1 \leq t \leq T. \quad (\text{PP2b})$$

$$q_{tj} \in Q_{tj}, \quad \forall 1 \leq t \leq T, \quad \forall 1 \leq j \leq S \quad (\text{PP2c})$$

$$I_t \geq 0, \quad l_t \geq 0, \quad \forall 1 \leq t \leq T. \quad (\text{PP2d})$$

4. Solution Approach: the Residual Tree Method

A standard numerical approach to solving multi-stage stochastic programs such as (PP2) is to model the uncertainty using Scenario Trees that directly represent the evolution of the key uncertainty drivers such as the demand. Such an approach would be justified for (PP2) if the decision-maker (DM) had sufficiently many iid observations of the demand process, or, if the DM knew the distributions of the underlying demand process (from which s/he could simulate an arbitrarily many observations of iid demand trajectories). Clearly, neither situation is applicable to our problem of procuring a product that has never been sold before, thus the Scenario Tree method (which is a well-known and studied technique in stochastic programming; see the references in the literature review for the latest on this method) is not applicable in our setting.

To solve Zara's procurement problem (PP2), we thus develop a new method, called the Residual Tree method, which circumvents the lack of direct demand observations for the new product by using historical data on similar items sold in the past. While the DM could use the demand data on similar products to construct a Scenario Tree directly, this would yield decisions that are systematically biased, due to differences in the products. Instead, the Residual Tree method employs a two-step approach whereby in the first step, least-squares (or lasso) regression is carried out to not only learn the parameters of the demand model (2), but also to construct artificial demand data for the never-before sold new product using the residuals from the regression. The second step solves an approximated version of (PP2) using the constructed data set.

We present two versions of the Residual Tree method to estimate the solution of (PP2), $(\mathbf{q}^*, \mathbf{I}^*, \mathbf{l}^*)$. The first version (presented in Section 4.2) is to be used when the number of features, m_t , is small relative to n (perhaps identified through a prior feature-selection step). The second version (presented in Appendix B) is designed for situations when m_t is large, possibly much larger than n , and only a subset of the features have a non-negligible role in the demand model (2). In Section 5 we prove the asymptotic optimality of both versions of our approach. For comparison, we first present the standard Scenario Tree approach to solving (PP2) under the hypothetical assumption that demand observations for the new product were available.




Figure 3: A schematic of the Residual Tree, when there is no serial correlation in the demand (otherwise, the second-stage residuals depend on the first-stage residuals) (left plot). A schematic of histogram binning to reduce the size of the Residual Tree, where twelve residual scenarios are reduced to six. Note that the distribution of the residuals need not be normal, and the bell-curve in the schematic is for illustrative purposes only (right plot).

4.1. Scenario Tree Algorithm

In the hypothetical situation where the DM had K observations of the demand path for the new product, $\{d_{01}^k, \dots, d_{0T}^k\}_{k=1}^K$, a textbook Scenario Tree approach to the problem (PP2) is to solve the following approximation:

$$\begin{aligned} \min_{\substack{q_{tj}^k, I_t^k, l_t^k \\ 1 \leq t \leq T, 1 \leq j \leq S, 1 \leq k \leq K}} \quad & \frac{1}{K} \sum_{k=1}^K \left[\sum_{t=1}^{T-1} h_t I_t^k + \sum_{t=1}^T b_t l_t^k + \sum_{t=1}^T \sum_{j=1}^S c_{tj} q_{tj}^k - v I_T^k \right] \quad (\text{SC}) \\ \text{s.t.} \quad & \forall 1 \leq k \leq K \ \& \ \forall 1 \leq t \leq T: \\ & q_{tj}^k = q_{tj}^\ell \quad \forall k, \ell \text{ with same demand path up to } t-1, \forall 1 \leq j \leq S, \quad (\text{SCa}) \\ & I_t^k = I_{t-1}^k + l_t^k + \sum_{\substack{j \in S \\ 1 \leq \tau \leq t-1: \\ \tau = t - \ell_j}} q_{\tau j} - d_{0t}^k, \quad (\text{SCb}) \\ & q_{tj}^k \in Q_{tj} \quad \forall 1 \leq j \leq S \quad (\text{SCc}) \\ & I_t^k \geq 0, \ l_t^k \geq 0. \quad (\text{SCd}) \end{aligned}$$

The key changes to (PP2) by the approximation (SC) are: (i) the decision processes become K -dimensional vectors, one for each demand path scenario; (ii) the objective is now an empirical average of the costs over the K scenario paths, rather than an expectation; (iii) the non-anticipativity constraint (PP2a) simplifies to (SCa); and (iv) the almost-sure random constraint (PP2b) is replaced by K deterministic constraints (SCb).

The Scenario Tree approach is one of the standard approaches to solving a multi-stage stochastic program as (PP2). In the special case where the demand process is a discrete random vector with K equally-weighted scenario paths, (SC) is equivalent to (PP2); otherwise (SC) is a computable

approximation of (PP2). For a textbook treatment of the topic we refer the readers to Shapiro et al. (2009), and for recent developments, the references mentioned in the literature review.

Although the Scenario Tree approach is popular and well-studied, the assumptions that underpin it renders its application to our problem impossible, because the DM does not have demand observations for a brand-new product (either directly or through the knowledge of the demand distribution), which appears in (SCb). Our Residual Tree method remedies this shortcoming by constructing trees based on residuals from regressions on similar products sold in the past.

4.2. Residual Tree Algorithm (Least-squares Regression version)

Step 1. *Residual Tree construction.*

- (a) *Least-squares regression.* For each $t \in [1, \dots, T]$, perform least-squares regression on available data on the n historical demands of similar products sold in the past:

$$\min_{\alpha_t \in \mathbb{R}, \beta_t \in \mathbb{R}^{m_t}} \sum_{k=1}^n (d_{kt} - \alpha_t - \beta_t^\top \mathbf{x}_{kt})^2. \quad (3)$$

Let $(\hat{\alpha}_t, \hat{\beta}_t)$ denote the optimal solution, and $\{\varepsilon_{kt}\}_{k=1}^n$ the residuals. We can construct n “samples” of the demand for the new product by combining the results of (3):

$$\{\hat{d}_{0t}^k = \hat{\alpha}_t + \hat{\beta}_t^\top \hat{\mathbf{x}}_{0t} + \varepsilon_{kt}\}_{k=1}^n, \quad (4)$$

where we use the $\hat{\cdot}$ notation in \hat{d}_{0t}^k to highlight the fact that d_{0t}^k is not a draw from a known distribution of the new product demand (as is the case in the Scenario Tree method), but an estimated sample from cross-learning from similar products, and in $\hat{\mathbf{x}}_{0t}$ to highlight the fact that some components of $\hat{\mathbf{x}}_{0t}$ may themselves be estimated from data (e.g. time-series components), as discussed in Sec. 3.1. We now have n independent demand estimates at every point in time for the new product, equalling n^T scenario paths in total.

- (b) *Histogram binning (for efficient computation if n is too large).* We could use the demand samples constructed in Step 1 (a) directly, but if n is too large, we can reduce the computation time by binning the residuals into a histogram (see Fig. 3, right plot).

Choose $1 \leq B_t \leq n$, the number of histogram bins at time t , for $t = 1, \dots, T$, and $\bar{\varepsilon}_{bt}$, $b = 1, \dots, B_t$ the center of the b -th bin at time t . B_t is also equal to the number of branches emanating from each node in the tree at period t . Let p_{bt} denote the empirical probability of a residual ε_{kt} being in the b -th bin at time t .

We now have B_t sample predictions for the demand at time $t = 1, \dots, T$:

$$\{\bar{d}_{0t}(b) := \hat{\alpha}_t + \hat{\beta}_t^\top \hat{\mathbf{x}}_{0t} + \bar{\varepsilon}_{bt}\}_{b=1}^{B_t}, \quad (5)$$

each occurring with probability p_{bt} . The total number of scenarios is thus reduced to $B := B_1 \times \dots \times B_T$. Let $\mathcal{B} = \{\{b_1, \dots, b_T\} : b_t \in \{1, \dots, B_t\}, t = 1, \dots, T\}$ denote the set of labels for all possible scenario paths, and denote a particular scenario path by $\mathbf{b} = [b_1, \dots, b_T]$, which occurs with probability $p_{\mathbf{b}} = p_{b_1} \times \dots \times p_{b_T}$. Denote the discretized state space by $\hat{\Xi}_n = \hat{\Xi}_1 \times \dots \times \hat{\Xi}_T$ and the corresponding discretized probability measure by $\hat{\mathbf{P}}_n = \hat{P}_1 \times \dots \times \hat{P}_T$.

Step 2. *Solve an estimated problem using the Residual Tree from Step 1.*

Compute replenishment decisions by solving the following optimization problem (denoted EPP, for “estimated procurement problem”):

$$\min_{\substack{q_{tj}(\cdot), I_t(\cdot), l_t(\cdot) \\ 1 \leq t \leq T, 1 \leq j \leq S}} \sum_{\mathbf{b} \in \mathcal{B}_n} p_{\mathbf{b}} \left[\sum_{t=1}^{T-1} h_t I_t(\mathbf{b}) + \sum_{t=1}^T b_t l_t(\mathbf{b}) + \sum_{t=1}^T \sum_{j=1}^S c_{tj} q_{tj}(\mathbf{b}) - v I_T(\mathbf{b}) \right] \quad (\text{EPP})$$

s.t. $\forall 1 \leq t \leq T :$

$$q_{tj}(\mathbf{b}) = q_{tj}(\mathbf{b}'), \quad \forall \mathbf{b}, \mathbf{b}' \in \mathcal{B}, \quad \text{s.t. } \bar{d}_{0\tau}(b_\tau) = \bar{d}_{0\tau}(b'_\tau), \quad \forall 1 \leq \tau \leq t-1, \quad \forall 1 \leq j \leq S, \quad (\text{EPPa})$$

$$I_t(\mathbf{b}) = I_{t-1}(\mathbf{b}) + l_t(\mathbf{b}) + \sum_{j \in S} \sum_{\substack{1 \leq \tau \leq t-1: \\ \tau = t - \ell_j}} q_{\tau j}(\mathbf{b}) - \bar{d}_{0,t}(b_t), \quad \forall \mathbf{b} \in \mathcal{B}, \quad (\text{EPPb})$$

$$q_{tj}(\mathbf{b}) \in Q_{tj}, \quad \forall 1 \leq j \leq S, \quad \forall \mathbf{b} \in \mathcal{B} \quad (\text{EPPc})$$

$$I_t(\mathbf{b}) \geq 0, \quad l_t(\mathbf{b}) \geq 0, \quad \forall \mathbf{b} \in \mathcal{B}, \quad (\text{EPPd})$$

where we indicate the dependence of the demand data and decisions on the particular scenario path \mathbf{b} explicitly. Denote the solution to (EPP) by $(\hat{\mathbf{q}}(n), \hat{\mathbf{I}}(n), \hat{\mathbf{l}}(n)) \in (\mathbb{R}^{TS} \times \mathbb{R}^T \times \mathbb{R}^T)^B$, where $\hat{\mathbf{q}}(n) = [\hat{q}_{1j}(\cdot), \dots, \hat{q}_{Tj}(\cdot)]_{j=1}^S \in \mathbb{R}^{TS}$, $\hat{\mathbf{I}}(n) = [\hat{I}_1(\cdot), \dots, \hat{I}_T(\cdot)] \in \mathbb{R}^T$ and $\hat{\mathbf{l}}(n) = [\hat{l}_1(\cdot), \dots, \hat{l}_T(\cdot)] \in \mathbb{R}^T$.

Remarks.

1. *Decisions after $t = 1$:* The estimated optimal ordering policy $\hat{\mathbf{q}}(n)$ from (EPP) is a matrix of values that maps the initial state (starting inventory, I_0) and the n constructed demand paths to ordering decisions. At time $t = 1$, the decision-maker orders \hat{q}_{1j} from each supplier $j = 1, \dots, S$, the actual demand for the new product d_{01} is observed, with which inventory I_1 and lost sales l_1 can be computed. For subsequent decisions in periods $t = 2, \dots, T$, the decision-maker resolves (EPP) with ordering quantities, inventories and lost sales up to period $t - 1$ fixed by their

realized values. In other words, the decision-maker solves the procurement problem via a forward rolling-horizon approach.

2. *Histogram binning*: There are a myriad of ways to construct the histogram in Step 1 (b), with the total number of bins, B_t , the centres $\bar{\varepsilon}_{bt}$, $b = 1, \dots, B_t$ and the widths δ_{bt} of the bins at time t being free parameters. (Clearly, a smaller number of bins is more computationally efficient but has greater discretization error; we illustrate this tradeoff in Section 6.) Note that, however, for the asymptotic optimality results in Section 5, the decision-maker needs to increase the total number of bins with n such that the resulting histogram measure is asymptotically equivalent to the sample-average measure. We state this requirement precisely in the statement of Theorem 1 in the next section.
3. *Unknown covariates for the new product*: Whenever constructing a residual tree, we construct sample demand predictions $\bar{d}_{0t}(b)$ using equation (5) for a bin b in a period $t > 0$ using mid-points from bin b 's parent bins as values for the relevant lagged demand covariates. When (EPP) is solved in a forward rolling-horizon manner (see Remark 1), as the season progresses the firm will observe actual dynamic covariate values for the new product and so can progressively replace any estimated components of the covariate vector \hat{x}_{0t} , $t = 1, \dots, T$, before each re-solve of (EPP).
4. We leave the sets Q_{tj} unspecified in (EPP) for the sake of expositional simplicity, reminding the reader that the theory to be developed in Section 5 only requires these sets to be bounded. Various possible choices of set types with that property lead to different specifications of (EPP) for computational purposes. For example, for closed intervals $[0, K_j]$ where K_j is a capacity limit for supply source j (EPP) is a linear program. Sets of the type $\{0\} \cup [O_j, K_j]$ would also capture a minimum order quantity O_j and would lead to (EPP) being a mixed integer program.
5. *Comparison with Scenario Tree algorithm of Sec 4.1*: Apart from fundamental differences in the starting assumptions, the key mathematical difference between the two approaches is in the inventory balancing equations, whereby in (EPPb), the estimated demand sample is used as opposed to the actual demand sample in (SCb). Thus, the Residual Tree method is subject to estimation errors as well as discretization errors, unlike the Scenario Tree method which is subject to only the latter. Our theory in Section 5 extends results for the Scenario Tree method by reconciling both errors.

In the online Appendix B, we extend the algorithm just stated to the high-dimensional setting where the DM may have a number of redundant covariates in the data set, by replacing the least-squares regression of Step 1 (a) with lasso regression. That is, the decision-maker has access to a superset of $r_t > m_t$ covariates (among which the identities of the true m_t covariates are unknown).

Applying least-squares regression (3) using all r_t covariates becomes problematic in this case because of ill-conditioning of the design matrix, hence the need for a modification to the algorithm.

In the following section we state the asymptotic optimality of the procurement solutions obtained using the Residual Tree Algorithm (Least-squares Regression version). The lasso version of the algorithm is also asymptotically optimal; we leave this statement to Appendix B.

5. Theory: Asymptotic Optimality of the Residual Tree Method

The main result of this section is that if nature generates product family demand data from the random models (1)–(2), then both the optimal value and the optimal solution of (EPP) converge to the optimal objective value and the optimal solution of (PP) as the number of products n tends to infinity, which is known as epi-convergence in the stochastic programming literature. (Note that we use the terms ‘epi-convergence’ and the more self-evident ‘asymptotic optimality’ interchangeably.) We state this formally below. The proofs of the results are deferred to Appendix C, and below we provide some qualitative explanations of the key ingredients.

THEOREM 1. *Assume that, for all $y \in \mathbb{R}$ and $1 \leq t \leq T$,*

$$\left| \sum_{1 \leq b \leq B_t} p_{bt} \mathbb{I}(\bar{d}_{0t}(b) \leq y) - \frac{1}{n} \sum_{k=1}^n \mathbb{I}(\hat{d}_{0t}^k \leq y) \right| \xrightarrow{P} 0 \quad (6)$$

as $n \rightarrow \infty$, where \xrightarrow{P} denotes convergence in probability. Then the optimal values of (EPP) obtained using Residual Tree Algorithm (Least-squares Regression version) converge to that of (PP) as $n \rightarrow \infty$ and all cluster points of $\{\hat{q}_{11}(n), \dots, \hat{q}_{1S}(n), \hat{I}_1(n), \hat{l}_1(n)\}_{n \in \mathbb{N}}$ are optimal first-stage solutions of (PP).

The assumption (6) requires the binned histograms to converge to the true density as the sample-average distribution. A necessary requirement for this is for B_t to grow linearly with n , i.e., $B_t = Cn$ for some constant $C \in (0, 1)$. One intuitive way to bin the residuals so that (6) is satisfied is to bin them into finer intervals as n grows (e.g. by binning into quartiles at first, then by deciles, then in ever smaller percentiles as n grows).

An immediate consequence of Theorem 1 is that decisions made in periods $t = 2, 3, \dots, T$ by the forward recursion process described in Remark 1. of Sec 4.2 also converge to the optimal t -stage solutions of (PP), conditioned on the demand path $d_{01}, \dots, d_{0(t-1)}$ and orders $[\hat{q}_{1j}, \dots, \hat{q}_{(t-1)j}]_{j=1}^S$ made up to period $t - 1$. The result is immediate because conditioned on the demand path and orders up to period $t - 1$, the T -stage problem (PP) reduces to a $T - t + 1$ -stage problem, whose first-stage solution is the decision for period t .

To prove Theorem 1, we make use of recent results of Pennanen (2005) and Pennanen (2009), who establish sufficient conditions for discretized multi-stage stochastic programs to epi-converge (i.e. both the objective and the first-stage solutions converge). We extend these results to the Residual Tree method, which are subject to estimation errors as well as discretization errors, as discussed in Remark 5 in Sec. 4.2. The full proof is rather intricate and we simplify the presentation by first establishing three separate lemmas on the nature of the true problem (PP), one on the continuity property of the random objective function, one on the expectation of the objective function, and one on the property of the demand model. Apart from the well-behaved nature of (PP), asymptotic optimality requires the consistency of the estimated demand processes, with or without the residual binning, and convergence of the probabilities of the scenario paths as the number of similar products grows. We show that our Residual Trees, with and without binning, satisfy both requirements.

6. Numerical Study

In this section we present results of a numerical study designed to illuminate the properties and behavior of the Residual Tree algorithm. Our primary goals in this section are to understand the finite-sample performance of the algorithm—namely, the dependence of the performance on the amount of available data (both covariate information and the number of observations), the resolution of the Residual Tree, and the running time of the algorithm—thereby providing insight into its potential for wide-spread applicability in practice. We also seek to verify that the policy output of our computational algorithm has qualitative properties that are consistent with the known optimal solutions of related stylized analytical models described in the literature. All results we report on were calculated using MATLAB 8.6, calling Gurobi 6.00 to solve linear programs, running on a single Macintosh desktop computer with a 3.2 GHz Intel i5 processor. We view this as a conservative choice relative to industrial computing clusters available at many firms.

Our tests focus on two sets of procurement problem instances differing primarily in how we simulate the demand data inputs. The first set of problem instances are two-period problems in which the demand does not depend on static covariates but is serially correlated across periods. The second includes four-period problems in which the generative model is based on both static and time-series covariates. In both cases, we first specify the true demand model (either by construction or from raw Zara data), then simulate data sets of various sizes from them to evaluate the convergence behavior of the Residual Tree algorithm.

For both sets of examples we use the Residual Tree Algorithm (Least-squares Regression version) as described in Section 4, and we reduce the size of our Residual Trees using the binning technique

described in Step 2, which controls the number of bins (i.e., the number of tree branches) at each node of the Residual Tree. The bin breakpoints are set at equal percentiles of the residual data used to fit the tree, and we represent each bin in the tree by its median residual value. It follows that all the bins $b = 1, \dots, B_t$ at a node are assigned equal probabilities $p_{bt} = 1/B_t$. Note that this binning method satisfies the required condition specified in (6) if B_t is made to grow linearly with n .

We note that for the results reported here we solve the problem (EPP) once at the beginning of the horizon and apply the single policy obtained over the whole horizon. Thus the performance results are for the largest problem instance the DM has to solve, and we omit reporting results for the entire forward rolling horizon approach described in Remark 1 of Section 4.2 for the sake of brevity and simplicity. To implement the Residual Tree policy without re-solving, we interpret each bin as capturing a range of possible residual values, such that the set of bins $\{1, \dots, B_t\}$ represents a mutually exclusive and exhaustive partitioning of the possible residual realizations in period t . This enables us to trace out-of-sample residual paths through the tree. It also means that a single solution to (EPP) yields an ordering policy that can be applied to new products with previously unseen error sequences, even without re-solving in a rolling-horizon manner. (In practice, the DM may as well re-solve the problem as the selling season unfolds and new demands are observed.) When we do not re-solve (EPP) in subsequent periods, we obviously forego the opportunity to replace dynamic covariate estimates with realized values, as described in Remark 3 in Section 4.2.

6.1. Two-Period Instances without Static Covariates

For the instances considered in this subsection, there are two periods with ordering opportunities and demand realizations in each period. There are three supply options: the firm can place an order to be delivered prior to period 1 at cost 0.5 (“pre-season” supplier), the firm can place an order prior to period 1 to be delivered between the two periods at cost 0.5 (“slow” supplier), and the firm can deliver after period 1 for delivery prior to period 2 at cost 1 (“fast” supplier). We assume a penalty cost of $b = b_1 = b_2 = 11$ per unit of unmet demand, a holding cost of $h_1 = 0.25$ per unit carried over between the two periods, and no salvage value for units left over at the end of the problem horizon.

We simulate demand trajectories such that demand realizations do not depend on any static covariates but are linearly dependent across periods. For each path, demand in period 1 is sampled from a normal distribution with mean 1000 and standard deviation 100, and demand in period 2 is sampled from a normal distribution with standard deviation σ_2 and mean equal to the period 1 demand realization. That is, we sample from the following population demand model for $k = 1, \dots, n$,

where n is the number of similar products in the product family for which the firm has demand data.

$$\begin{aligned} D_{k1} &= 1000 + \varepsilon_1, & \varepsilon_1 &\sim N(0, 100^2), \\ D_{k2} &= 0 + 1D_{k1} + \varepsilon_2, & \varepsilon_2 &\sim N(0, \sigma_2^2), \end{aligned} \quad (7)$$

Note that in demand model (7) the standard deviation σ_2 of the second error term ε_2 is a parameter capturing the amount of learning occurring between the first and the second period. This is because smaller values of σ_2 induce larger correlations between D_{k1} and D_{k2} in our problem, so that observations of the first-period demand are more predictive of the second-period demand when σ_2 is small. As a result, the problem instance defined in this subsection is qualitatively similar to the two period models with learning considered for example in Fisher et al. (2001), Milner and Kouvelis (2002, 2005) and Li et al. (2009). Later, we will present some results examining how order quantities produced by our approach vary with σ_2 . Until then, we set $\sigma_2 = 100$.

For each of the results reported here, we run the algorithm on 200 independently generated training sets, then we evaluate the resulting policies on a single, common, and independently generated “test set.” We vary the number n of demand trajectories in the training sets, but we use a common test set of size 100,000. (Since the instances we consider here do not depend on static covariates, the same policy can be applied to all paths in the test set. We take a different approach in the next subsection, where demand depends on static covariates that differ across test set paths.) Times, costs, and orders reported in this subsection are averaged across 200 runs. Unless indicated otherwise, we keep the number of bins per node constant (i.e., $B = B_t$ for all t) for each instance.

Table 1 reports the performances of policies generated using training sets ($n = 50, 200, 1000$) and trees of different sizes ($B = B_1 = B_2 = 1, 2, 3, 5, 10, 25, 50$). We express the average percentage costs in Table 1 as percentages relative to the average cost 1813.46 of the best-performing policy, calculated using $n = 1000$ and $B = 50$ and assuming that the coefficients in the generative demand model (7) are known. We use this as a proxy for the cost of the true optimal policy, noting that this average cost is within 1% of that produced using the next-largest tree. We give results for versions of the model in which the coefficients in the generative demand model are known, unknown, and estimated from the training data, and estimated using a misspecified intercept-only set of models that ignore the dependence between the first and second period demands. Note the case with known coefficients eliminates the estimation error, leaving only the discretization error (defined in Section 4.2 Remark 5) and is thus equivalent to a situation in which the firm has n sample paths of past demand for the focal product and applies the Scenario Tree method as outlined in Section 4.1.

B	Known coefficients			Estimated coefficients			Intercept-only		
	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$
1	67.1	65.6	65.1	67.0	65.6	65.1	66.8	65.8	65.1
2	27.7	26.1	25.5	27.8	26.2	25.5	23.1	22.0	21.8
3	10.0	8.9	8.7	10.1	8.9	8.7	11.8	10.5	10.4
5	3.7	3.0	2.8	3.8	3.0	2.8	5.4	4.4	4.3
10	1.5	1.0	0.9	1.6	1.0	0.9	2.6	1.9	1.7
25	0.7	0.2	0.1	0.9	0.2	0.1	3.8	3.0	2.9
50	0.7	0.1	0.0	0.8	0.1	0.0	3.5	2.6	2.5

Table 1: Out-of-sample performance of three versions of the proposed policy for a two-period problem without static covariates. Results are presented as percentage cost increase compared with the best-performing policy, marked in bold. We provide standard errors for these estimates in online Appendix D.

Interestingly, the “estimated coefficients” policy based on a relatively small data set of $n = 50$ training paths and a relatively small Residual Tree with $B = 10$ bins per node achieves an out-of-sample cost performance just 1.6% larger than our best policy. This is a promising result from an implementation standpoint, as $n = 50$ is a realistically small number of similar items in a product family. $B = 10$ bins per node achieves an attractive balance in this case between oversimplifying the stochasticity in the problem and computational efficiency.

Directionally, the results conform to expectations. We see that the out-of-sample performance of the policy generally improves as the number of training paths (n) increases and as the number of bins per node (B) in the Residual Tree increases for each of the three versions of the policy. This is intuitive: a larger training set reduces estimation error by yielding more accurate regression coefficient estimates and by allowing for more accurate calibration of the discrete model of errors, while more branching reduces the discretization error by imposing a finer model of demand. The best cost performances are achieved for instances with largest n and B , which is consistent with performance convergence both as a function of the amount of training data and with the size of the tree. In addition, for fixed training set size and tree size we see that the policy based on known coefficients outperforms the policy based on estimated coefficients, which in turn outperforms the policy based on the misspecified intercept-only demand model for all instances except for those using the smallest trees.

A closer examination of Table 1 shows a strong dependence of the results on the size of the tree used in the algorithm but relatively small differences between policies computed with 50, 200, and 1000 training paths. By comparing the results of the policy with estimated coefficients with those for known coefficients, we conclude that the value of knowing the true demand model coefficients (which corresponds to using the hypothetical Scenario Tree algorithm of Sec. 4.1) is surprisingly

marginal for the instances considered, even for the case with $n = 50$. This shows that the additional estimation error introduced by the Residual Tree method, as explained in Remark 2 of Sec. 4.2, is dominated by the discretization error that both it and the Scenario Tree method are subject to.

Using an over-simplified demand model (the “intercept-only” model), however, has a larger cost penalty relative to using a correctly specified model, leading to an out-of-sample cost overage of 2.5% for the case even when $n = 1000$ and $B = 50$. This indicates the presence of a systematic bias due to model misspecification.

Figure 4 (left plot) illustrates convergence of the policy itself by showing the average initial orders placed as a function of the tree size B . Results shown are for $n = 50$ and the “estimated coefficients” version of the policy. Plots for the cases $n = 200$ and $n = 1000$, omitted for brevity, show average initial orders to be nearly identical to those for $n = 50$. Overall, we conclude that the numerical results are consistent with asymptotic optimality as a function of both n and B .

The computation time required to compute the policy depends primarily on the size of the Residual Tree. The largest tree we considered ($B = 50$) required an average computation time of approximately 24 seconds to compute a policy based on a single training set of 1000 demand paths.

In all the computations discussed above, we have built Residual Trees with the same number of bins for each node. However, the architecture of the Residual Tree can be optimized for a particular instance. Figure 4 (right plot) shows the out-of-sample performance of the policies generated from Residual Trees with more or fewer bins in the first period relative to the second period. Letting B_1 and B_2 indicate the number of bins per node in the first and second periods, respectively, we consider 9 policies with $B_1 = 1, 2, 3, 4, 6, 9, 12, 18, 36$ and $B_2 = 36/B_1$. For the particular problem considered, the best policies result from slightly more branching in the first period relative to the second ($B_1 = 9, B_2 = 4$). In practice, the best tree architecture will depend on the problem instance, the underlying uncertainty, and its evolution over time.

We have also investigated the sensitivity of the policies generated by our approach to various problem parameters. We have found the sensitivities to the cost parameters to be quite intuitive, and we omit the details for brevity. For example, we have seen that the average total order quantity increases with the shortage penalty cost b and decreases with the supply costs. We have also found that as we increase the costs of the pre-season and slow suppliers relative to the fast supplier, the retailer (intuitively) tends to substitute the fast supplier for the less responsive suppliers.

Figure 5 shows how the average orders placed with the various supply options depend on the learning parameter σ_2 defined earlier. Total orders tend to increase with σ_2 , which can be understood from basic newsvendor arguments; that is, the retailer stocks more stock to buffer added demand




Figure 4: Initial orders under the policy as a function of the number of bins per node (B) in the Residual Tree (left plot). Out-of-sample performance of policies based on Residual Trees with more or fewer bins in the first period relative to the second period (right plot).

uncertainty. This increase comes from larger orders placed with the slow supplier. We observe that orders with the pre-season supplier are unaffected by σ_2 in these instances. (The demand uncertainty in period 1 of our problem is relatively small compared with the expected demand in period 2, implying that the problem effectively decomposes across periods.) The average orders placed with the fast supplier actually decrease with σ_2 , which is more clearly evident in the right-hand plot of Figure 5. Because smaller values of σ_2 reflect a greater amount of learning occurring between the first and second periods, our results suggest that the attractiveness of the fast supplier increases with the ability to better predict second period demand based on the observation of demand in the first period. We have confirmed this insight by measuring the increase in cost when we make the fast supplier unavailable. Indeed, this cost increase exceeds 7% for $\sigma_2 = 25$ but is only 2% for $\sigma_2 = 200$. We note that these results are qualitatively consistent with the numerical findings for the value of flexibility reported in analytical studies of two period ordering problems with learning, such as Milner and Kouvelis (2002, 2005). In the context of Zara, they are also consistent with the intuition that orders from China should decrease with the speed of demand learning at the beginning of the selling season (see Section 1).

Figure 6 illustrates the policy we compute for ordering from the fast supply option as a function of realized first-period demand. We see that fast supplier orders given $\sigma_2 = 200$ are no larger than orders given $\sigma_2 = 25$, in keeping with Figure 5. Both computed policies therefore involve a minimum threshold for the second period order based on the first period demand realization, and a linear increase for the second period order beyond that threshold. This exactly matches the (s, S) policy structure that is proven to be optimal or near-optimal in analytical studies of two period ordering problems with specific distributional assumptions, such as Milner and Kouvelis (2005) and Li et al.




Figure 5: Sensitivity of policy to standard deviation σ_2 of ε_2 (left plot). Average fast supplier orders shown with smaller graph scaling (right plot). (Policies are averaged over all runs and all test set demand paths for training sets of size $n = 200$ and Residual Trees with $B_1 = B_2 = 25$ bins per node.)




Figure 6: Policy for ordering from fast supplier as a function of the first-period demand realization. (Each policy is computed based on a single training set path.)

(2009). Furthermore, we note that the ordering threshold for the second order decreases with the amount of learning, and that the slope of the increasing portion of both plots is approximately two. This comes from two sources: (1) each additional unit of first-period sales depletes stock by one unit that must be replenished, and (2) each additional unit of first-period demand raises the expectation for second-period demand by one unit, arising from the slope of “1” in the equation for D_2 in (7). We conclude that the solutions produced by our computational approach exhibit the same key qualitative properties that were previously established by available analytical and parametric studies of similar two period problem instances.

6.2. Instances with Static Covariates

As discussed in Section 3.1, a key innovation of our approach is its accounting for covariates in its assumed demand process. While the demand model considered in the previous subsection includes lagged demand as a “dynamic” covariate in the second period, we would also like to consider demand models that include “static” covariates that reflect information known about the product prior to

the selling season. As mentioned earlier, in practice these static covariates may capture features such as price, design, or expected popularity.

Let us consider four-period problems in which the population demand model includes both static and dynamic covariates. Specifically, the demand models for the $k = 1, \dots, n$ products are:

$$\begin{aligned} D_{k1} &= \{738 + 1072X_{k1} - 403X_{k2} + 2.8X_{k3} - 55X_{k4} + \varepsilon_1\}^+, & \varepsilon_1 &\sim N(0, 1044^2), \\ D_{k2} &= \{-399 - 638X_{k1} + 53X_{k4} + 0.854D_{k1} + \varepsilon_2\}^+, & \varepsilon_2 &\sim N(0, 781^2), \\ D_{k3} &= \{-5 + 0.955D_{k2} + \varepsilon_3\}^+, & \varepsilon_3 &\sim N(0, 601^2), \\ D_{k4} &= \{874 - 46X_{k4} + 0.516D_{k2} + 0.318D_{k3} + \varepsilon_4\}^+, & \varepsilon_4 &\sim N(0, 820^2), \end{aligned}$$

where the static covariates are distributed as follows:

$$\begin{aligned} X_{k1} &\stackrel{iid}{\sim} \text{Bernoulli}(0.25), \\ X_{k2} &\stackrel{iid}{\sim} \text{Bernoulli}(0.5), \\ X_{k3} &\stackrel{iid}{\sim} \text{Normal}(900, 200^2), \\ X_{k4} &\stackrel{iid}{\sim} \text{Discrete Uniform on } \{7.95, 8.95, \dots, 22.95\}. \end{aligned}$$

This demand model is based on regression models fit using stepwise forward variable selection method to real sales data for a single category of 70 garments introduced by Zara during a single season. Here, the covariates X_{k1} and X_{k2} correspond to binary features present (or absent) in each garment, X_{k3} to the number of stores in which the garment was initially introduced, and X_{k4} to the retail price of the garment. The distributions of these covariates are loosely based on observed data for the category.

We assume two supply options available each period: a fast supplier able to supply with zero lead time $\ell_1 = 0$ at cost $c_1 = c_{11} = c_{21} = c_{31} = c_{41} = 1$, a slow supplier able to supply with lead time $\ell_2 = 1$ and cost $c_2 = c_{21} = c_{22} = c_{32} = c_{42} = 0.5$, a holding cost of $h = h_1 = h_2 = h_3 = h_4 = 0.25$ per unit per period, and a penalty cost of $b = b_1 = b_2 = b_3 = b_4 = 11$ per unit of unmet demand. We assume salvage value of zero for inventory left over at the end of the horizon.

As previously, we simulate data sets to serve as training and test data for fitting and evaluating the policies, respectively. Here, for each demand path we also generate a set of static covariates, randomly and independently drawn according to the distributions specified above. Solving the optimization problem (EPP) yields an ordering policy conditional on a single set of static covariates. Therefore, to generate the results in this section we must re-solve the policy for each test demand path. For this reason, we make use of a smaller test set than in Section 6.1. Specifically, the test set comprises

200 independently generated combinations of covariates and demand paths. As before, we vary the size of the training sets and the size of the Residual Trees.

In Table 2, we present the out-of-sample performance of policies generated using training sets of sizes $n = 50, 200, 1000$ and Residual Trees of increasing sizes ($B = B_1 = B_2 = B_3 = B_4 = 1, 2, 3, 5, 7, 10$). We compute four versions of the policy in order to understand the value of covariate information: one in which we assume that the true regression coefficients used to generate the data are known, one in which the form of the true regression model is known but the coefficients must be estimated from the training data, one in which the static covariates are unavailable, and one that ignores both static and dynamic (i.e., lagged demand) covariates. As previously, we present the results as percentages relative to the best-performing policy, which was computed assuming known coefficients based on 1000 training paths and a Residual Tree with 10 bins per node.

As we saw in the previous 2-period example, in Table 2 we see that the policy performances improve with the amount of training data, with the size of the Residual Tree, and with better-specified regression models. Strong performance can be achieved even with a relatively limited amount of training data; for example, for trees with 10 bins per node, the performance difference between the $n = 50$ and $n = 1000$ policies is just 2.0% for the policy based on known coefficients and 4.0% for the policy based on estimated coefficients. Convergence of the policy performances is evident as the underlying Residual Trees increase in size.

For $B = 10$ and for a realistically small training set of $n = 50$, the performance of the estimated problem is within 4.9% of the best-performance, again supporting the practical validity of the algorithm. The value of the covariate information is evident; for the same data set size $n = 50$ and coarseness $B = 10$, the loss in the total cost by ignoring static covariates only and by ignoring all covariates are $10.8\% - 4.9\% = 5.9\%$ and $20.3\% - 4.9\% = 15.4\%$ of the total cost of the best-performing policy, respectively. A two-sample t-test reveals that these differences are statistically significant at the 1% level.

Interestingly, we also observe that the performance of the algorithm heavily depends on the number of bins used. For instance, for $n = 50$ with all covariates, using $B = 2$ yields solutions that are worse than then $B = 10$ by $71.2\% - 4.9\% = 66.3\%$ of the total cost of the best-performing policy. For $B = 3$, the difference is $34.6\% - 4.9\% = 29.7\%$ of the total cost of the best-performing policy. Both differences are statistically significant at the 1% level. The literature on the dynamic procurement problem to date often considers only 2–3 demand branches per node, and our results indicate that such simplified models may not yield good solutions in practice.

B	Known coefficients			Estimated coefficients			Ignoring static covariates			Ignoring all covariates		
	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$
1	181.2	178.4	178.9	184.3	183.0	179.8	199.8	190.7	192.8	195.2	189.7	191.4
2	55.3	52.1	53.0	71.2	58.9	53.8	69.1	66.5	69.8	75.6	75.1	74.3
3	26.3	21.3	22.6	34.6	25.3	26.7	34.6	34.2	34.0	46.2	46.2	45.3
5	6.3	6.2	4.9	15.7	8.0	6.2	19.6	17.5	15.6	27.6	26.5	24.5
7	2.8	3.0	2.6	7.6	5.4	2.1	14.3	14.2	11.6	22.0	22.4	21.5
10	2.0	0.4	0.0	4.9	1.5	0.9	10.8	9.1	9.0	20.3	18.8	19.0

Table 2: Out-of-sample performance of four versions of the proposed policy for a four-period problem with static covariates. Results are presented as percentage cost increase compared with best-performing policy, marked in bold. We provide standard errors for these estimates in online Appendix D.

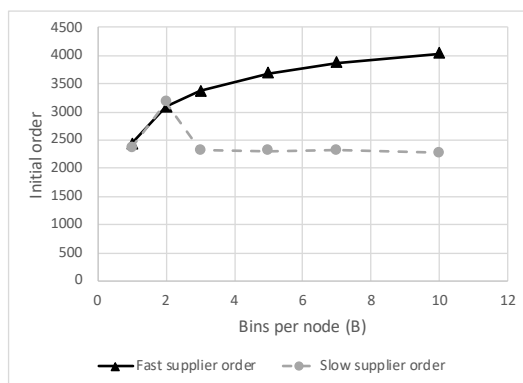


Figure 7: Initial orders under the policy as a function of the number of bins per node (B) in the Residual Tree.

Figure 7 plots average initial orders to the fast and slow suppliers as a function of the size of the underlying Residual Tree. Results reflect the policy based on the correctly specified regression model with estimated coefficients and training sets of $n = 50$ demand paths, although we note that using larger training sets makes no discernible difference in the appearance of this plot. As predicted by our theory, we see the policy decisions converging as the tree increases in size and granularity.

6.3. Computation Time

One limitation of the Residual Tree method, as with the Scenario Tree method, is that the computational complexity of a tree-based approach will grow exponentially with the horizon T .

For the purpose of understanding computation times, we solve instances of the problem (EPP) for horizons $T = 4, 6, 8$, for simulated training sets of size $n = 1000$, and for various choices for the binning configuration $[B_1, \dots, B_T]$. We assume two suppliers, and we sample training data from a discrete-time martingale population demand model with increments given by normal random variables and no static covariates.




Figure 8: Average computation time required to compute the “estimated coefficients” policy based on a single set of 1000 training paths for various horizons and binning configurations.

The times plotted in Figure 8 represent times to solve (EPP) averaged over 20 instances on a single desktop computer processor. (As mentioned previously, this is a conservative choice relative to industrial computing clusters available to some retailers.) We find that the times in Figure 8 are well-described ($R^2 = 0.9994$) by a quadratic function of the count of nodes in the tree $\sum_{t=1}^T \prod_{\tau=1}^t B_{\tau}$:

$$\text{Computation time (seconds)} \approx 0.4832 + 0.0067(\text{Node count}) + (1.178 \times 10^{-6})(\text{Node count})^2.$$

We can use this quadratic relationship to predict the time required to solve similar instances of various horizons with various binning configurations. We present some example predictions in Table 3, where we estimate both the time to solve a single T -period problem (“solve time”) and the total time required to re-solve in each period t using bin configuration $[B_1, \dots, B_{T-t+1}]$ (“re-solve time”). We see that (EPP) remains computationally feasible for 8- and 10-period problems with 3 bins per period, but becomes prohibitively time-consuming for larger trees. One opportunity for improvement would be to consider trees with differing numbers of bins per period; for example, the last line in Table 3 shows that we expect that a 10-period configuration with 10 bins in the first period and 2 bins per period for periods 5-10 will be solved faster than the configuration with 3 bins per period. Some preliminary experiments have suggested that the best choice of binning configuration depends on the evolution of uncertainty over time.

It is also possible to use solutions to short-horizon versions of our residual tree method as heuristic policies for longer-horizon problems. For example, even if replenishment decisions can be made weekly in practice, we might consider using a coarser time granularity in the model, as there may be limited incremental benefit to modeling more ordering opportunities within a selling season. Other promising ideas include smart scenario reduction methods, parallelization, limited lookahead heuristics, re-solving on rolling horizons, and hybrid approaches that combine some of these ideas.

To characterize the relative performance of these heuristics over a representative set of problem instances is a significant task that we leave for future work. We point out that, while in practice

T	B	Node count	Est. solve time (s)	Tot. re-solve time (s)	T	B	Node count	Est. solve time (s)	Tot. re-solve time (s)
4	[5 5 5 5]	780	6	9	8	[2 2 2 2 2 2 2 2]	510	4	11
4	[7 7 7 7]	2,800	29	33	8	[3 3 3 3 3 3 3 3]	9,840	181	231
4	[10 10 10 10]	11,110	221	232	8	[5 5 5 5 5 5 5 5]	488,280	2.8e+05	3.0e+05
6	[3 3 3 3 3 3]	1,092	9	15	10	[2 2 2 2 2 2 2 2 2 2]	2,046	19	39
6	[5 5 5 5 5 5]	19,530	581	635	10	[3 3 3 3 3 3 3 3 3 3]	88,572	9,839	11,297
6	[7 7 7 7 7 7]	137,256	23,119	23,738	10	[5 5 5 5 5 5 5 5 5 5]	12,207,030	1.8e+08	1.8e+08
6	[10 10 ... 10]	1,111,110	1.5e+06	1.5e+06	10	[10 5 4 3 2 2 2 2 2 2]	76,460	7,403	10,160

Table 3: Predicted computation times to solve (EPP) for selected horizons and binning configurations.

a retailer would require a policy determined on a single demand path, for us to rigorously study policy performance with either static covariates or policy re-solving requires us to solve (EPP) over each demand path. This implies that the computational limits bind more tightly when studying performance than when implementing in practice.

7. Conclusion

As mentioned in the beginning, this work arose out of an engagement with Zara in which we initially adapted standard stochastic programming approaches for the dynamic procurement problem they regularly face when introducing new apparel items. The data-driven approach described in this paper has not yet been implemented at Zara, due to various exogenous time constraints and resource allocation choices. As part of our ongoing relationship with that firm however, we hope to leverage this work and learn more about related implementation aspects by helping them develop an enhanced data-driven version of their current decision support system for dynamic procurement. To that end, we note that for implementation purposes Zara and other retailers may need to further specify the relatively generic formulation (EPP) in order to capture important features of their business environment, such as supplier-specific quantity discount schemes and minimum order quantity constraints. A wide range of such features may be captured using integer programming formulation techniques (Bertsimas and Weismantel 2005), and we emphasize again that the results in Section 5 establish that the asymptotic optimality of the Residual Tree method continues to hold for any resulting mixed integer programming formulation as long as it corresponds to bounded sets Q_{t_j} .

Finally, the Residual Tree method is a general data-driven method for solving multi-stage stochastic programs, and as such may be valuable for other applications that employ stochastic programming. Examples include financial options pricing, where historical prices and features such as economic indicators and market factors abound; and energy utility management, where historical demands and demographic and weather feature data are available (see Rios et al. 2015).

Acknowledgments

The authors thank Yiangos Papanastasiou for input in early stages of this project and Miguel Díaz, Francisco Babio Fernandez, Juan José Lema, Julio López and Olaia Vázquez from Zara/Inditex for sharing real-world context and data and for many helpful discussions. The third author thanks the Sarah Graham Kenan Foundation for support.

References

- Agrawal, Narendra, Stephen A Smith, Andy A Tsay. 2002. Multi-vendor sourcing in a retail supply chain. *Production and Operations Management* **11**(2) 157–182.
- Allon, Gad, Jan A Van Mieghem. 2010. Global dual sourcing: Tailored base-surge allocation to near-and offshore production. *Management Science* **56**(1) 110–124.
- Baardman, Lennart, Georgia Perakis, Divya Singhvi. 2017. Leveraging comparables for new product sales forecasting. *Proceedings of MSOM Conference* .
- Ban, Gah-Yi. 2018. Confidence intervals for data-driven inventory policies with demand censoring. *SSRN Working Paper* .
- Ban, Gah-Yi, N. Bora Keskin. 2017. Personalized dynamic pricing. Available at SSRN: <https://ssrn.com/abstract=2972985> .
- Ban, Gah-Yi, Cynthia Rudin. 2018. The big data newsvendor: Practical insights from machine learning. *forthcoming in Operations Research* .
- Bertsimas, Dimitris, Robert Weismantel. 2005. *Optimization Over Integers*. Dynamic Ideas, Belmont, Massachusetts.
- Burnetas, Apostolos, Stephen Gilbert. 2001. Future capacity procurements under unknown demand and increasing costs. *Management Science* **47**(7) 979–992.
- Burnetas, Apostolos N, Craig E Smith. 2000. Adaptive ordering and pricing for perishable products. *Operations Research* **48**(3) 436–443.
- Casey, Michael S, Suvrajeet Sen. 2005. The scenario generation algorithm for multistage stochastic linear programming. *Mathematics of Operations Research* **30**(3) 615–631.
- Chen, Xi, Zachary Owen, Clark Pixton, David Simchi-Levi. 2015. A statistical learning approach to personalization in revenue management .
- Cohen, Maxime C, Ilan Lobel, Renato Paes Leme. 2016. Feature-based dynamic pricing. *SSRN Working Paper* .
- Dupačová, Jitka, Giorgio Consigli, Stein W Wallace. 2000. Scenarios for multistage stochastic programs. *Annals of operations research* **100**(1-4) 25–53.

- Dupačová, Jitka, Nicole Gröwe-Kuska, Werner Römisch. 2003. Scenario reduction in stochastic programming. *Mathematical programming* **95**(3, Ser. A) 493–511.
- Eppen, Gary D, Ananth V Iyer. 1997. Improved fashion buying with Bayesian updates. *Operations research* **45**(6) 805–819.
- Escudero, Laureano F., Pasumarti V. Kamesam, Alan J. King, Roger J-B. Wets. 1993. Production planning via scenario modelling. *Annals of Operations Research* **43** 311–335.
- Ferreira, Kris Johnson, Bin Hong Alex Lee, David Simchi-Levi. 2015. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* **18**(1) 69–88.
- Fisher, Marshall, Kumar Rajaram, Ananth Raman. 2001. Optimizing inventory replenishment of retail fashion products. *Manufacturing & Service Operations Management* **3**(3) 230–241.
- Fisher, Marshall, Ananth Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations research* **44**(1) 87–99.
- Friedman, Jerome, Trevor Hastie, Robert Tibshirani. 2009. *The elements of statistical learning, 2nd ed.*. Springer series in statistics Springer, Berlin.
- Godfrey, Gregory A, Warren B Powell. 2001. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science* **47**(8) 1101–1112.
- Graves, Stephen C, Harlan C Meal, Sriram Dasu, Yuping Qui. 1986. Two-stage production planning in a dynamic environment. *Multi-stage production planning and inventory control*. Springer, 9–43.
- Greene, William H. 2011. *Econometric analysis*. 7th ed. Pearson International Edition.
- Heath, David C, Peter L Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE transactions* **26**(3) 17–30.
- Heitsch, Holger, Werner Römisch. 2003. Scenario reduction algorithms in stochastic programming. *Computational optimization and applications* **24**(2-3) 187–206.
- Higle, Julia H., Karl G. Kempf. 2011. Production planning under supply and demand uncertainty: A stochastic programming approach. G. Infanger, ed., *Stochastic Programming*, chap. 14. International Series in Operations Research and Management Science, Springer Science+Business Media, 297–315.
- Hu, Kejia, Jason Acimovic, Francisco Erize, Douglas J Thomas, Jan A Van Mieghem. 2016. Forecasting product life cycle curves: Practical approach and empirical analysis .
- Huh, Woonghee Tim, Paat Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* **34**(1) 103–123.
- Iida, Tetsuo, Paul H Zipkin. 2006. Approximate solutions of a dynamic forecast-inventory model. *Manufacturing & Service Operations Management* **8**(4) 407–425.

- Ioffe, Alexander D. 1977. On lower semicontinuity of integral functionals. i. *SIAM Journal on Control and Optimization* **15**(4) 521–538.
- Jones, Philip C., Timothy J. Lowe, Rodney D. Traub, Greg Kegler. 2001. Matching supply and demand: The value of a second chance in producing hybrid seed corn. *Manufacturing & Service Operations Management* **3**(2) 122–137.
- Kaminsky, Philip, Jayashankar M Swaminathan. 2001. Utilizing forecast band refinement for capacitated production planning. *Manufacturing & Service Operations Management* **3**(1) 68–81.
- Kaminsky, Philip, Jayashankar M Swaminathan. 2004. Effective heuristics for capacitated production planning with multiperiod production and demand with forecast band refinement. *Manufacturing & Service Operations Management* **6**(2) 184–194.
- Knight, Keith, Wenjiang Fu. 2000. Asymptotics for lasso-type estimators. *Annals of statistics* 1356–1378.
- Kunnunkal, Sumit, Huseyin Topaloglu. 2008. Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research* **56**(3) 646–664.
- Lau, Hon-Shiang, Amy Hing-Ling Lau. 1996. The newsstand problem: A capacitated multiple-product single-period inventory problem. *European Journal of Operational Research* **94**(1) 29–42.
- Levi, Retsef, Georgia Perakis, Joline Uichanco. 2015. The data-driven newsvendor problem: new bounds and insights. *Operations Research* **63**(6) 1294–1306.
- Levi, Retsef, Robin O Roundy, David B Shmoys. 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–839.
- Li, Jian, Suresh Chand, Maqbool Dada, Shailendra Mehta. 2009. Managing inventory over a short season: models with two procurement opportunities. *Manufacturing & Service Operations Management* **11**(1) 174–184.
- Lu, Xiangwen, Jing-Sheng Song, Amelia Regan. 2006. Inventory planning with forecast updates: Approximate solutions and cost error bounds. *Operations research* **54**(6) 1079–1097.
- Milner, Joseph M, Panos Kouvelis. 2002. On the complementary value of accurate demand information and production and supplier flexibility. *Manufacturing & Service Operations Management* **4**(2) 99–113.
- Milner, Joseph M, Panos Kouvelis. 2005. Order quantity and timing flexibility in supply chains: The role of demand characteristics. *Management Science* **51**(6) 970–985.
- Patel, Jalpa. 2012. Optimization-based decision support system for retail sourcing. Master’s thesis, Massachusetts Institute of Technology.
- Peng, Chen, Feryal Erhun, Erik F. Hertzler, Karl G. Kempf. 2012. Capacity planning in the semiconductor industry: Dual-mode procurement with options. *Manufacturing & Service Operations Management* **14**(2) 170–185.

- Pennanen, Teemu. 2005. Epi-convergent discretizations of multistage stochastic programs. *Mathematics of Operations Research* **30**(1) 245–256.
- Pennanen, Teemu. 2009. Epi-convergent discretizations of multistage stochastic programs via integration quadratures. *Mathematical Programming* **116**(1-2) 461–479.
- Pennanen, Teemu, Matti Koivu. 2002. Integration quadratures in discretization of stochastic programs. *Stochastic Programming E-Print Series* .
- Pflug, Georg, Ronald Hochreiter. 2003. Scenario generation for stochastic multi-stage decision processes as facility location problems. Tech. rep., Department of Statistics and Decision Support Systems, University of Vienna.
- Qiang, Sheng, Mohsen Bayati. 2016. Dynamic pricing with demand covariates. *SSRN Working Paper* .
- Rios, Ignacio, Roger JB Wets, David L Woodruff. 2015. Multi-period forecasting and scenario generation with limited data. *Computational Management Science* **12**(2) 267–295.
- Serel, D. A. 2016. Intelligent procurement systems to support fast fashion supply chains in the apparel industry. Tsan-Ming Choi, ed., *Information Systems for the Fashion and Apparel Industry*, chap. 7. Elsevier, Amsterdam, The Netherlands, 121–144.
- Shapiro, Alexander, Darinka Dentcheva, Andrzej Ruszczyński. 2009. *Lectures on stochastic programming: modeling and theory*. SIAM.
- Wang, Tong, Atalay Atasü, Mümin Kurtuluş. 2012. A multiordering newsvendor model with dynamic forecast evolution. *Manufacturing & Service Operations Management* **14**(3) 472–484.
- Zipkin, Paul. 2008. On the structure of lost-sales inventory models. *Operations research* **56**(4) 937–944.

Electronic Companion to “Dynamic Procurement of New Products with Covariate Information: the Residual Tree Method”

Appendix A: Proof of Proposition 1

We proceed by first proving (A) that a feasible solution to problem (PP) is feasible in problem (PP2). Then we prove (B) that an optimal solution to problem (PP2) is feasible in (PP). Because the objective functions of the two problems are identical, it follows that the optimal solutions to the two problems and their objective values must coincide. We show both reductions pathwise, for a fixed demand path $\mathbf{d}_0 = \{d_{01}, \dots, d_{0T}\}$ and fixed order decisions $\mathbf{q} = \{q_{1j}, \dots, q_{Tj}\}$. As a shorthand, we let $Q_t = \sum_{j \in \mathcal{S}} \sum_{\substack{1 \leq \tau \leq t-1 \\ \tau = t - \ell_j}} q_{\tau j}$ indicate the total shipment received in period t given the demand path.

To prove claim (A), let the pair $(\mathbf{I}^*, \mathbf{l}^*)$, where $\mathbf{I}^* = \{I_1^*, \dots, I_T^*\}$ and $\mathbf{l}^* = \{l_1^*, \dots, l_T^*\}$, denote a feasible solution to (PP). Fix t and consider the sign of $I_{t-1}^* + Q_t - d_{0t}$. If $I_{t-1}^* + Q_t - d_{0t} \geq 0$ then $I_t^* = I_{t-1}^* + Q_t - d_{0t}$ by constraint (PPb) (which holds almost surely) and $l_t^* = 0$ by constraint (PPc). Therefore constraint (PP2b) holds on the given demand path for t . Furthermore, $I_t^* \geq 0$ so $I_t^* \in \mathbb{R}_+$ and $l_t^* \geq 0$ so $l_t^* \in \mathbb{R}_+$. Similarly, if $I_{t-1}^* + Q_t - d_{0t} < 0$ then $I_t^* = 0$ by constraint (PPb) and $l_t^* = d_{0t} - I_{t-1}^* - Q_t$ by constraint (PPc) so constraint (PP2b) holds and we have $I_t^* \geq 0$ and $l_t^* \geq 0$. Since this holds for all \mathbf{d}_0 and \mathbf{q} , we have (A).

To prove claim (B), now let the pair $(\mathbf{I}^*, \mathbf{l}^*)$, where $\mathbf{I}^* = \{I_1^*, \dots, I_T^*\}$ and $\mathbf{l}^* = \{l_1^*, \dots, l_T^*\}$, denote an optimal solution to (PP2). We prove (B) via the following lemma.

LEMMA EC.1. *Let $(\mathbf{I}^*, \mathbf{l}^*)$ be an optimal solution to problem (PP2) given the assumed demand path. Then $l_t^* > 0$ implies $I_t^* = 0$ for all $t \in \{1, \dots, T\}$.*

(Proof of Lemma EC.1.) Suppose $l_\tau^* \geq \epsilon$ and $I_\tau^* \geq \epsilon$ for some $\tau \in \{1, \dots, T\}$ and some $\epsilon > 0$. We contradict the optimality of this solution by constructing an alternative solution $(\hat{\mathbf{I}}, \hat{\mathbf{l}})$ with strictly lower cost. Our construction of this alternative solution depends on τ .

If $\tau = T$, then let

$$\begin{aligned} \hat{l}_t &= l_t^*, \quad \forall t \in \{1, \dots, T-1\} \\ \hat{l}_T &= l_T^* - \epsilon \\ \hat{I}_t &= I_t^*, \quad \forall t \in \{1, \dots, T-1\} \\ \hat{I}_T &= I_T^* - \epsilon. \end{aligned}$$

It is straight-forward to show that $(\hat{\mathbf{I}}, \hat{\mathbf{l}})$ is feasible in (PP2). Furthermore, the cost of solution $(\mathbf{I}^*, \mathbf{l}^*)$ less the cost of $(\hat{\mathbf{I}}, \hat{\mathbf{l}})$ is strictly positive:

$$b_T (l_T^* - \hat{l}_T) - v (I_T^* - \hat{I}_T) = (b_T - v)\epsilon > 0.$$

If $\tau < T$, then let

$$\begin{aligned}\hat{l}_t &= l_t^*, \quad \forall t \in \{1, \dots, T\} \setminus \{\tau, \tau + 1\} \\ \hat{l}_\tau &= l_\tau^* - \epsilon \\ \hat{l}_{\tau+1} &= l_{\tau+1}^* + \epsilon \\ \hat{I}_t &= I_t^*, \quad \forall t \in \{1, \dots, T\} \setminus \{\tau\} \\ \hat{I}_\tau &= I_\tau^* - \epsilon.\end{aligned}$$

Again, $(\hat{\mathbf{I}}, \hat{\mathbf{l}})$ is feasible in (PP2). The cost of solution $(\mathbf{I}^*, \mathbf{l}^*)$ less the cost of $(\hat{\mathbf{I}}, \hat{\mathbf{l}})$ is strictly positive:

$$b_\tau (l_\tau^* - \hat{l}_\tau) + b_{\tau+1} (l_{\tau+1}^* - \hat{l}_{\tau+1}) + h_\tau (I_\tau^* - \hat{I}_\tau) = (b_\tau - b_{\tau+1})\epsilon + h_\tau \epsilon > 0,$$

by assumptions on the parameters stated in Sec. 3.2. Since this holds for all \mathbf{d}_0 and \mathbf{q} , we have (B). \square

Returning to claim (B), we fix t and consider two cases. First, suppose $l_t^* = 0$. Then constraint (PP2b) implies $I_t^* = I_{t-1}^* - d_{0t} + Q_t$ and the non-negativity of I_t^* implies that $I_t^* = I_{t-1}^* - d_{0t} + Q_t \geq 0$. Therefore (PPb) holds. Furthermore, $I_{t-1}^* - d_{0t} + Q_t \geq 0$ implies that $(d_{0t} - I_{t-1}^* - Q_t)^+ = 0$ so (PPc) holds.

Second, suppose $l_t^* > 0$. Lemma EC.1 then implies $I_t^* = 0$, and constraint (PP2b) implies $l_t^* = d_{0t} - I_{t-1}^* - Q_t > 0$ so (PPb) holds. Furthermore, $d_{0t} - I_{t-1}^* - Q_t > 0$ implies that $(I_{t-1}^* + Q_t - d_{0t})^+ = 0$ so (PPc) holds.

Appendix B: Residual Tree Algorithm (Lasso Regression version): High Dimensional Data

In our demand model as described in Section 3.1, the demand in period t depends on an m_t -dimensional vector of covariates. Supposing that the decision-maker has access to a superset of $r_t > m_t$ covariates (among which the true m_t covariates are unknown), then $r_t - m_t$ of the covariates are irrelevant or redundant. Applying least-squares regression (3) using all r_t covariates is then problematic because of model misspecification and the increased likelihood of ill-conditioned covariance of the design matrix. Thus, for this case, we propose modifying Step 1 (a) of the Residual Tree Algorithm (Least-squares Regression version) by replacing the least squares regression with Lasso regression as follows:

Step 1. (a)' *Lasso Regression*. For each $t \in [1, \dots, T]$, perform Lasso regularized regression on available data on the n existing products:

$$\min_{\alpha_t \in \mathbb{R}, \beta_t \in \mathbb{R}^{m_t}} \sum_{k=1}^n (d_{kt} - \alpha_t - \beta_t^\top \mathbf{x}_{kt}^{Lasso}) + \lambda_n \|\beta_t\|_1, \quad (\text{EC.1})$$

where $\lambda_n > 0$ is the regularization parameter, and $\|\cdot\|_1$ denotes the L_1 norm. Denote the solution vector by $(\hat{\alpha}_t^{Lasso}, \hat{\beta}_t^{Lasso})$. We can thus construct n sample estimates for the new product by using the estimated parameters and the residuals from the Lasso regression:

$$\{\hat{d}_{0t}^{Lasso} = \hat{\alpha}_t^{Lasso} + \hat{\beta}_t^{Lasso \top} \hat{\mathbf{x}}_{0t} + \hat{\varepsilon}_{kt}^{Lasso}\}_{k=1}^n. \quad (\text{EC.2})$$

If histogram binning Step 1 (b) is employed after Lasso Step 1 (a)', we have the following B_t samples of constructed data:

$$\{\bar{d}_{0t}^{Lasso}(b) := \hat{\alpha}_t^{Lasso} + \hat{\beta}_t^{Lasso \top} \hat{\mathbf{x}}_{0t} + \bar{\varepsilon}_{bt}^{Lasso}\}_{b=1}^{B_t}, \quad (\text{EC.3})$$

where $\bar{\varepsilon}_{bt}^{Lasso}$, $b = 1, \dots, B_t$ are the centres of the histogram bins for the residuals from Lasso regression (EC.1), $\{\hat{\varepsilon}_{kt}^{Lasso}\}_{k=1}^n$.

The Residual Tree Algorithm (Lasso Regression version) is also asymptotically optimal. We defer the proof to Appendix C.

THEOREM EC.1. *Assume (6) and $\lambda_n/n = o(1)$. Then the optimal values of (EPP) using Residual Tree Algorithm (Lasso Regression version) converge to that of (PP) as $n \rightarrow \infty$ and all cluster points of $\{\hat{q}_{11}(n), \dots, \hat{q}_{1S}(n), \hat{I}_1(n), \hat{l}_1(n)\}_{n \in \mathbb{N}}$ are optimal first-stage solutions of (PP).*

Remark on the choice of the regularization parameter λ_n : If λ_n is too large, the resulting parameter estimates $(\hat{\alpha}_t^{Lasso}, \hat{\beta}_t^{Lasso})$ will have large biases; whereas if λ_n is too small, the parameter estimates will retain dimensions that should not be in the model as the penalization effect will be small. A sufficient condition for asymptotic consistency of $(\hat{\alpha}_t^{Lasso}, \hat{\beta}_t^{Lasso})$ is $\lambda_n/n = o(1)$, i.e. λ_n should decrease faster than the rate $O(n)$ [Knight and Fu (2000)], which is one of the assumptions we make for Theorem EC.1. In practice, for fixed n , λ_n can be chosen via cross-validation; that is, we can tune λ_n over a grid of values based on some data put aside for validation purposes (see Friedman et al. 2009 for further details).

Appendix C: Proofs of results in Section 5.

The proofs of Theorem 1 and Theorem EC.1 rely on the following three lemmas.

LEMMA EC.2. *The function $f: \mathbb{R}^{T(S+2)} \times \Xi_0 \rightarrow (-\infty, \infty]$ given by*

$$f(\mathbf{q}, \mathbf{I}, \mathbf{l}; \mathbf{D}_{0,[1,T]}) = \sum_{t=1}^{T-1} h_t I_t + \sum_{t=1}^T b_t l_t + \sum_{t=1}^T \sum_{j=1}^S c_{tj} q_{tj} - v I_T, \quad (\text{EC.4})$$

where $\mathbf{q} = [q_{1j}, \dots, q_{Tj}]_{1 \leq j \leq S}$, $\mathbf{I} = [I_1, \dots, I_T]$, and $\mathbf{l} = [l_1, \dots, l_T]$ may depend on the random vector $\mathbf{D}_{0,[1,T]} = [D_{01}, \dots, D_{0T}]$, is lower semi-continuous in the first argument and has the lower compactness property of Ioffe (1977).

LEMMA EC.3. *The following conditions hold for the stochastic optimization problem (PP): for every feasible $\mathbf{z} \in \mathbb{R}^{T(S+2)}$, there is a uniformly bounded sequence $\mathbf{y}_\mu \rightarrow \mathbf{z}$ of non-anticipative, \mathbf{P}_0 -a.s. continuous functions $\mathbf{y}_\mu: \Xi_0 \rightarrow \mathbb{R}^{T(S+2)}$ such that*

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{\mathbf{P}_n} f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \leq \mathbb{E}_{\mathbf{P}_0} f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}), \quad \forall \mu \in \mathbb{Z}_+ \text{ and} \quad (\text{a})$$

$$\limsup_{\mu \rightarrow \infty} \mathbb{E}_{\mathbf{P}_0} f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \leq \mathbb{E}_{\mathbf{P}_0} f(\mathbf{z}(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}). \quad (\text{b})$$

LEMMA EC.4. *The stochastic demand model (1) is a time series model with uniform innovations, i.e. for each $t = 1, \dots, T$, the demand is of the form*

$$D_{0t} | \mathbf{X}_{0t} = g_t(D_{00}, D_{01}, \dots, D_{0,t-1}, \omega_t), \quad (\text{EC.5})$$

where D_{00} is a (non-random) constant, $\omega_1, \dots, \omega_T$ are mutually independent random variables, with ω_t uniformly distributed in the unit interval $[0, 1]$, and $g_t: \Xi_{01} \times \dots \times \Xi_{0,t-1} \times [0, 1] \rightarrow \Xi_{0t}$.

(*Proof of Lemma EC.2.*) Firstly, $f(\cdot; \mathbf{D}_{0,[1,T]})$ is continuous for every $\mathbf{D}_{0,[1,T]} \in \Xi_0$ so it is lower semi-continuous. Secondly, the function $f: \mathbb{R}^{T(S+2)} \times \Xi_0 \rightarrow (-\infty, \infty]$ satisfies the lower compactness property of [Ioffe \(1977\)](#) if there exists a non-decreasing real-valued function g on $[0, +\infty)$ and a real number b such that

$$f(\mathbf{z}; \mathbf{D}_{0,[1,T]}) \geq -g(\|\mathbf{z}\|_\infty) - b, \quad \forall \mathbf{z} \in \mathbb{R}^{T(S+2)}, \quad \forall \mathbf{D}_{0,[1,T]} \in \Xi_0.$$

Now

$$f(\mathbf{q}, \mathbf{I}, \mathbf{l}; \mathbf{D}_{0,[1,T]}) \geq -T(c_\infty S \|\mathbf{q}\|_\infty + (h_\infty \vee v) \|\mathbf{I}\|_\infty + b_\infty \|\mathbf{l}\|_\infty), \quad \forall (\mathbf{q}, \mathbf{I}, \mathbf{l}) \in \mathbb{R}^{T(S+2)}, \quad \forall \mathbf{D}_{0,[1,T]} \in \Xi_0,$$

where $c_\infty = \max\{c_{1j}, \dots, c_{Tj}\}_{1 \leq j \leq S}$, $h_\infty = \max\{h_1, \dots, h_{T-1}\}$ and $b_\infty = \max\{b_1, \dots, b_T\}$, thus $f(\cdot; \cdot)$ satisfies the lower compactness property of [Ioffe \(1977\)](#). \square

(*Proof of Lemma EC.3.*) The condition (a) holds because by (6), $\hat{\mathbf{P}}_n$ converges to the sample average measure as n goes to infinity, then by the Weak Law of Large Numbers (WLLN) on the sample average measure.

For condition (b), let $\mathbf{z} = (\mathbf{q}, \mathbf{I}, \mathbf{l}) \in \mathbb{R}^{T(S+2)}$ be an admissible policy to (PP). Consider, for $\mu \in \mathbb{Z}_+$, the function

$$\mathbf{y}_\mu = \mathbf{z} - \frac{\mathbf{1}_0}{\mu},$$

where $\mathbf{1}_0$ is a vector with $T(S+2)$ ones and a zero at the end (corresponding to the location of I_T in the \mathbf{z} vector). Then

$$\|\mathbf{y}_\mu\|_\infty \leq \|\mathbf{z}\|_\infty + 1,$$

hence \mathbf{y}_μ is uniformly bounded for every \mathbf{z} .

We also have, by elementary algebra,

$$f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) = f(\mathbf{z}(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) - \frac{C}{\mu},$$

where C is the positive constant

$$C = \sum_{t=1}^{T-1} h_t + \sum_{t=1}^T b_t + \sum_{t=1}^T \sum_{j=1}^S c_{tj},$$

which is clearly dominated by $f(\mathbf{z}; \mathbf{D}_{0,[1,T]})$ for all $\mu \in \mathbb{Z}_+$. Thus we have

$$\mathbb{E}_P f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \leq \mathbb{E}_P f(\mathbf{z}(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]})$$

for all $\mu \in \mathbb{Z}_+$, and (b) follows. \square

(*Proof of Lemma EC.4.*) Given $\mathbf{X}_{0t} = x_{0t}$, define the function

$$g_t(D_{00}, \dots, D_{0,t-1}, \omega_t) := \alpha_t + \beta_t^\top \mathbf{x}_{0t} + F_{\varepsilon_t}^{-1}(\omega_t), \quad (\text{EC.6})$$

where $F_{\varepsilon_t}^{-1}$ is the inverse cdf of the random error term ε_t in the demand model (2). We immediately have

$$D_{0t} | (\mathbf{X}_{0t} = \mathbf{x}_{0t}) = g_t(D_{00}, \dots, D_{0T}, \omega_t), \quad (\text{EC.7})$$

which satisfies the definition of a time series model with uniform innovations. \square

Before proceeding with the proofs, let us first introduce some definitions.

Define the Lebesgue space $Z(\mathbf{P}_0) := L^\infty(\Xi_0, \mathcal{F}_0, \mathbf{P}_0; \mathbb{R}^{T(S+2)})$, and the set of $(\mathcal{F}_{0t})_{t=1}^T$ -adapted elements of $Z(\mathbf{P}_0)$

$$\mathcal{N}(\mathbf{P}_0) := \{\mathbf{z} \in Z(\mathbf{P}_0) \mid \mathbf{z} \text{ contains an } (\mathcal{F}_{0t})_{t=1}^T \text{-adapted function}\}.$$

In other words, $\mathcal{N}(\mathbf{P}_0)$ is the set of feasible solutions (each of which are, technically speaking, equivalence classes, as solutions that differ on sets of measure zero have equivalent objective values) of (PP).

Now consider the discretized demand space $\hat{\Xi}_n$. Let $\hat{\Xi}^{bt}$ denote the b -th bin out of $B_t(n)$ bins at time t . Define, for each $n \in \mathbb{Z}_+$ and $t = 1, \dots, T$, the functions $s_n : \Xi_0 \rightarrow \Xi_0$ and $\psi_{t,n} : \Xi_t \rightarrow \mathbb{R}$ given by

$$s_n(\mathbf{d}_{0,[1,T]}) := \bar{\mathbf{d}}_{0,[1,T]}(\mathbf{b}), \text{ and}$$

$$\psi_{t,n}(\mathbf{d}_{0,t}) := \frac{p_{bt}}{\mathbb{P}(\hat{\Xi}^{bt})},$$

where $\mathbf{b} = [b_1, \dots, b_T]$ is such that d_{0t} belongs to the b_t -th histogram bin $\hat{\Xi}^{b_t t}$ at time t , and $\bar{\mathbf{d}}_{0,[1,T]}(\mathbf{b}) = [\bar{d}_{01}(b_1), \dots, \bar{d}_{0T}(b_T)]$, where $\bar{d}_{0t}(\cdot)$, $t = 1, \dots, T$ are as defined in (5). In words, $s_n(\cdot)$ maps a vector of demand path for the new product to one of the demand paths defined by the Residual Tree by determining which histogram bin each demand belongs to; and $\psi_{t,n}(\cdot)$ denotes the ratio of the empirical histogram measure of a given demand path to the true measure of the histogram bins $\hat{\Xi}^{b_t t}$, $t = 1, \dots, T$. Thus, the closer $s_n(\cdot)$ is to the identity map, and the closer $\psi_{t,n}(\cdot)$ is to one, the better the Residual Tree approximation.

Finally, let $\mathcal{N}_n(\mathbf{P})$ be the set of $\{(s_n)^{-1}(\mathcal{F}_{0t})\}$ -adapted elements of $X(\mathbf{P}_0)$.

(*Proof of Theorem 1*). By Corollary 3.1. of Pennanen (2005), epi-convergence follows if our problem satisfies the following two conditions:

- *Condition 1*: the sequence $\{\hat{\mathbf{P}}_n\}_{n=1}^\infty$ of discretized measures is such that $\mathcal{N}_n(\mathbf{P}) \subset \mathcal{N}(\mathbf{P}_0)$ for all $n \in \mathbb{Z}_+$ and

$$s_n(\mathbf{D}_{0,[1,T]}) \xrightarrow{P} \mathbf{D}_{0,[1,T]}, \text{ and} \tag{EC.8}$$

$$\max_{b \in B_t(n)} \left| \frac{p_{bt}(n)}{P(\hat{\Xi}^{bt}(n))} - 1 \right| \rightarrow 0, \forall 1 \leq t \leq T, \tag{EC.9}$$

as n tends to infinity.

- *Condition 2*:

- The function $f : \mathbb{R}^{T(S+2)} \times \Xi_0 \rightarrow (-\infty, \infty]$ given by (EC.4) is lower semi-continuous in the first argument and has the lower compactness property of Ioffe (1977).
- For every feasible $\mathbf{z} \in \mathbb{R}^{T(S+2)}$, there is a uniformly bounded sequence $\mathbf{y}_\mu \rightarrow \mathbf{z}$ of non-anticipative, P -a.s. continuous functions $\mathbf{y}_\mu : \Xi_0 \rightarrow \mathbb{R}^{T(S+2)}$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{\hat{P}_n} f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \leq \mathbb{E}_P f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \quad \forall \mu \in \mathbb{Z}_+$$

$$\limsup_{\mu \rightarrow \infty} \mathbb{E}_P f(\mathbf{y}_\mu(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}) \leq \mathbb{E}_P f(\mathbf{z}(\mathbf{D}_{0,[1,T]}); \mathbf{D}_{0,[1,T]}),$$

where $f(\cdot, \cdot)$ is the function defined in (EC.4).

In words, Condition 1 states necessary conditions on the data-driven estimate of the stochastic program (PP). Specifically, (EC.8) is a statement regarding the asymptotic consistency of the data-driven estimate of the demand process and (EC.9) is a statement regarding the consistency of the residual histogram tree. Separately, Theorem 2 of Pennanen (2009) shows that if the underlying stochasticity of a multi-stage stochastic program is a time series model with uniform innovations as is the case with our problem by Lemma EC.4, then the sequence $(\hat{\mathbf{P}}_n)_{n=1}^\infty$ satisfies Condition 1 if the following conditions hold for all $t = 1, \dots, T$:

- C1. $g_t(D_{00}, \dots, D_{0,t-1}, \cdot)$ is a bijection for every $\mathbf{D}_{0,[1,T]} \in \Xi_0$, where $g_t(\cdot)$ is the function defined in (EC.5) of Lemma EC.4,
- C2. $g_t(\cdot)$ and the function $(D_{01}, \dots, D_{0t}) \mapsto g_t(D_{00}, \dots, D_{0,t-1}, \cdot)^{-1}(D_{0t})$ are Borel-measurable,
- C3. $g_t(\cdot)$ is $P_{01} \times \dots \times P_{0,t-1} \times \omega_t$ -almost-surely continuous, and
- C4. $\hat{\omega}_t(n) \rightarrow \omega_t$ weakly as $n \rightarrow \infty$, where $\hat{\omega}_t(n)$ is the random variable with the discrete uniform distribution $F_{\hat{\varepsilon}_t(n)}(\hat{\varepsilon}_t(n))$, where $\hat{\varepsilon}_t(n)$ is a random variable equal to ε_{kt} $k = 1, \dots, n$ with probability $1/n$ (recall ε_{kt} 's are the residuals from Step 1 (a) of the Residual Tree Algorithm (Least-squares Regression version)), and $F_{\hat{\varepsilon}_t(n)}(\cdot)$ denotes its cdf.

Below we show the conditions C1–C4 are satisfied by our problem.

- PC1. By the model assumption, ε_t has an inverse cdf, so $g_t(D_{00}, \dots, D_{0T}, \cdot)$ is a bijection for every $\mathbf{D}_{0,[1,T]} \in \Xi_0$.
- PC2. The Borel-measurability of $g_t(\cdot)$ and of the function $(D_{00}, \dots, D_{0,t-1}) \mapsto g_t(D_{00}, \dots, D_{0,t-1}, \cdot)^{-1}(D_{0t})$ follows from the measurability of ω_t and ε_t .
- PC3. It is clear that the function $g_t(\cdot)$ is continuous in all its arguments.
- PC4. By Lemma 2 of Pennanen (2009), this is equivalent to (EC.8) and (EC.9) being satisfied. (EC.8) is satisfied by the consistency of the least-squares regression model under the assumptions made in Sec. 3.1 [see, for example, Chapter 4.4.1 of Greene (2011)], and (EC.9) is satisfied by the asymptotic equivalence of the histogram distribution to the empirical distribution (spelled out in (6) of Step 2. of Residual Tree Algorithm (Least-squares Regression version)), the latter of which is uniformly consistent via the Glivenko-Cantelli Theorem. Note the conditions (EC.8) and (EC.9) hold even when there are time-series components in the covariate vector which may need to be estimated from data in earlier periods, which can be shown recursively from $t = 1$ because any time series covariate components are themselves estimated via least-squares regression in a consistent manner.

Condition 2 concerns the structure of the original problem itself, and we have shown that our problem satisfies these conditions by Lemmas EC.2 and EC.3. Thus our problem satisfies the assumptions of Corollary 3.1. of Pennanen (2005) and the conclusions follow. \square

Proof of Theorem EC.1. It suffices to check that the map $s_n^{Lasso} : \Xi_0 \mapsto \Xi_0$ defined by

$$s_n^{Lasso}(\mathbf{d}_{0,[1,T]}) := \bar{\mathbf{d}}_{0,[1,T]}^{Lasso}(\mathbf{b}),$$

where $\mathbf{b} = [b_1, \dots, b_T]$ is such that d_{0t} belongs to the b_t -th histogram bin $\hat{\Xi}^{b_t}$ at time t , $\bar{\mathbf{d}}_{0,[1,T]}^{Lasso}(\mathbf{b}) := [\bar{d}_{01}^{Lasso}(b_1), \dots, \bar{d}_{0T}^{Lasso}(b_T)]$, where $\bar{d}_{0t}^{Lasso}(\cdot)$, $t = 1, \dots, T$ are as defined in (EC.3), converges to the identity map as $n \rightarrow \infty$. Knight and Fu (2000) shows that this holds if $\lambda_n/n = o(1)$, i.e. λ_n decreases faster than $O(n)$ rate. \square

Appendix D: Standard Errors for Tables 1 and 2

This section provides standard errors corresponding to the tables of results in Sections 6.1 and 6.2.

B	Known coefficients			Estimated coefficients			Intercept-only		
	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$
1	1.28	0.61	0.40	1.24	0.64	0.40	1.15	0.59	0.39
2	0.53	0.26	0.20	0.52	0.26	0.20	0.58	0.31	0.21
3	0.30	0.16	0.11	0.30	0.16	0.11	0.39	0.20	0.14
5	0.15	0.08	0.06	0.17	0.08	0.06	0.22	0.11	0.08
10	0.07	0.03	0.02	0.07	0.03	0.02	0.10	0.04	0.04
25	0.04	0.01	0.01	0.05	0.01	0.01	0.13	0.06	0.05
50	0.05	0.01	0.00	0.06	0.01	0.00	0.13	0.06	0.05

Table EC.1: Standard errors for the performance ratio results presented in Table 1, estimated using bootstrapping.

B	Known coefficients			Estimated coefficients			Ignoring static covariates			Ignoring all covariates		
	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$	$n = 50$	$n = 200$	$n = 1000$
1	13.87	13.98	13.86	15.18	14.74	13.90	15.53	15.56	15.73	15.72	15.52	15.49
2	6.36	6.33	6.39	7.82	6.73	6.59	7.88	8.60	8.70	9.71	10.03	9.90
3	3.56	3.74	3.85	4.58	4.02	4.08	5.77	6.16	5.91	7.71	7.94	7.74
5	1.87	1.99	1.70	3.02	2.36	1.98	3.55	3.96	3.68	5.52	6.04	5.94
7	1.21	1.25	1.09	2.41	1.68	0.96	3.03	3.11	2.77	4.46	5.34	5.05
10	1.10	0.44	0.00	2.15	1.05	0.47	2.13	2.26	2.02	3.54	4.56	4.42

Table EC.2: Standard errors for the performance ratio results presented in Table 2, estimated using bootstrapping.