

Decision, Risk & Operations
Working Papers Series

Structural Estimation of the Effect of Out-of-Stocks

A. Musalem, M. Olivares, E.T. Bradlow, C. Terwiesch and
D. Corsten

August 2008
DRO-2008-12

Structural Estimation of the Effect of Out-of-Stocks

Andrés Musalem

Marcelo Olivares

Eric T. Bradlow

Christian Terwiesch

Daniel Corsten*

August 6, 2008

* Andrés Musalem is Assistant Professor of Marketing at Duke University's Fuqua School of Business. Marcelo Olivares is Assistant Professor of Decision, Risk and Operations at Columbia University's Business School. Eric T. Bradlow is The K.P. Chao Professor, Professor of Marketing, Statistics and Education and Co-Director of the Wharton Interactive Media Initiative at The Wharton School of the University of Pennsylvania. Christian Terwiesch is Associate Professor of Operations and Information Management at The Wharton School of the University of Pennsylvania. Daniel Corsten is Professor of Operations and Technology Management at Instituto de Empresa Business School. We would like to thank seminar participants at the University of Chicago's Graduate School of Business and the 2008 Marketing Science Conference hosted by the University of British Columbia. In particular, we would like to thank J.P. Dubé, Günter Hitsch, Peter Rossi and Naufel Vilcassim for providing very detailed and helpful comments about this paper. We would also like to extend our gratitude to the Jay H. Baker Retailing Initiative, The Wharton School, University of Pennsylvania for providing partial financial support. Please address all correspondence on this manuscript to Andrés Musalem, 1 Towerview Drive, Durham, NC 27708, Phone: (919) 660-7827, Fax: (919) 681-6245, amusalem@duke.edu.

Structural Estimation of the Effect of Out-of-Stocks

Abstract

We develop a structural demand model that captures the effect of out-of-stocks on customer choice. Our estimation method uses store-level data on sales and partial information on product availability. Our model allows for flexible substitution patterns which are based on utility maximization principles and can accommodate categorical and continuous product characteristics. The methodology can be applied to data from multiple markets and in categories with a relatively large number of alternatives, slow moving products and frequent out-of-stocks. We estimate our model using sales data from multiple stores for twenty four items in the shampoo product category. In addition, we illustrate how the model can be used to assist the decisions of a retailer in two ways. First, we show how to quantify the lost sales induced by out-of-stock products. Second, we provide insights on the financial consequences of out-of-stocks and suggest simple policies that can be used to help mitigate the negative economic impact of out-of-stocks.

Keywords: choice models; retailing; out-of-stocks; inventory management; Bayesian methods; data augmentation; aggregate demand estimation.

1 Introduction

In retailing, inventory decisions have direct implications on product availability to customers. When making decisions related to the inventory levels of a product category, store managers need to balance the costs of holding and replenishing inventory versus the costs of out-of-stocks. The cost of holding inventory can be calculated directly using financial measures available to store managers. In contrast, evaluating the cost of an out-of-stock requires estimating its impact on customers' buying behavior. The lack of precise measures of the costs of out-of-stocks has been cited as one of the root causes for the slow adoption of quantitative models in inventory management by practitioners (Zipkin (2000)).

In terms of the magnitude and prevalence of this problem, out-of-stocks are certainly not uncommon in retailing. The average out-of-stock rate in the U.S. and Europe is about 8% and the costs associated to out-of-stocks vary across product categories and can be substantial in some cases.¹ In order to quantify the financial consequences of out-of-stocks, it is useful to analyze the choices that a customer facing an out-of-stock could make. First, a customer encountering an out-of-stock may choose to defer its purchase until the product becomes available. Second, the customer may choose to purchase a substitute product in the category. Third, a customer may decide not to purchase any products – that is, the out-of-stock leads to a *lost sale*– which has the largest negative financial impact for the retailer. Our main objective is to develop a methodology to quantify the effect of the two latter scenarios, product substitution and lost sales, that can be implemented using information commonly available to a store manager.

A major challenge in estimating the impact of out-of-stocks on retail demand is the lack of precise data on product availability. One would expect the extensive adoption of perpetual inventory management systems in retailing, which monitor inventory in real-time, to alleviate this data limitation problem. However, there are two reasons why these inventory systems have not entirely solved this problem. First, in many cases these systems do not distinguish between inventory on the shelf or in the backroom and, as a consequence, the system might report availability for the store, while the shelf may be empty. Second, recent work by DeHoratius and Raman (2008) has shown large discrepancies between the actual inventory and system-recorded inventory. Consequently, audits

¹See Gruen et al. (2002) for a detailed study on the incidence and consequences of out-of-stocks across different product categories and geographies.

need to be conducted periodically in order to reconcile actual inventory with what is kept in the inventory system. Therefore, in practice, many retailers operate with a *periodic inventory review system*, where inventory is observed precisely only at specific time epochs and cannot be perfectly inferred at other points in time. For these reasons, we design our methods to work with sparse information of product availability, as provided by periodic inventory review systems.

Our estimation approach is based on an extension of the methodology in Chen and Yang (2007), Musalem et al. (2007) and Musalem et al. (2008) for demand estimation from aggregate data. Accordingly, we treat the sequence of individual purchases in a given time period (i.e., the order in which individual purchases were made) as missing data and simulate these sequences from their posterior distribution. Combining these simulated sequences of purchases with periodic inventory information, we estimate the evolution of product inventory on the shelf and make inferences about the demand model taking into account the occurrence of out-of-stocks. Consequently, we explicitly consider the *endogenous* variation in product availability, modeling the set of products available to a customer in a given time period as a function of the initial inventory *and* the sequence of choices made by customers. Moreover, when used in categories where out-of-stocks occur frequently, our data augmentation simplifies substantially the estimation of the model parameters relative to the Expectation-Maximization (EM) approach used in previous work (Anupindi et al. (1998) and Conlon and Mortimer (2007)). Furthermore, this structural approach to model product availability allows us to perform policy experiments that can be useful to estimate lost sales and evaluate the impact of policies designed to mitigate the consequences of out-of-stocks.

We estimate our model using sales data from multiple stores for twenty four items in the shampoo product category. In addition, we illustrate how the model can be used to assist the decisions of a store manager in two ways. First, we estimate lost sales induced by product out-of-stocks, which is used to assign a financial tag to the cost of out-of-stocks. Second, we use the model to evaluate the financial consequences of temporary price promotions that can help alleviate the costs of out-of-stocks by recapturing a fraction of the lost sales.

To summarize, our work makes four important contributions. First, we develop a methodology that can be applied in fairly general settings, including categories with a large number of products, some of which may be slow-moving products that exhibit zero sales in some periods. Second, our methodology explicitly considers the endogenous changes in product availability triggered by

customer choices. Third, our structural approach allows us to perform policy experiments that can be useful to estimate lost sales and evaluate the impact of policies designed to reduce the consequences of out-of-stocks. Fourth, the use of data augmentation greatly simplifies the estimation of the model parameters, especially when compared with EM or maximum simulated likelihood approaches.

The rest of this article is structured as follows. Section 2 relates our work to the existing literature. In section 3 we describe the demand model, estimation methodology and how the methodology can be used to estimate lost sales. Section 4 presents an empirical application of the methodology using a data set with shampoo purchases. Section 5 estimates the costs of out-of-stocks and illustrates how the methodology can be used to assess the impact of different policies aimed at mitigating their financial consequences. Finally, Section 6 concludes this paper with a discussion of interesting avenues for future research.

2 Literature Review

Our work is related to research streams in the operations management, marketing and economics literatures. For instance, our work is related to analytical models of inventory management and assortment planning developed in the operations management literature (e.g. Zipkin (2000), Smith and Agrawal (2000)). These models typically assume a specific demand formulation that is incorporated into an optimization problem; but in general they provide little guidance on how to determine the input parameters of the demand model. This limits the applicability of this work in two ways. First, the demand model specification needs to be validated through empirical data in order for the prescribed decisions to be relevant in practice. Second, a naïve estimation of a correctly specified demand model which neglects the effect of out-of-stocks on sales can lead to biased estimates and incorrect model inputs. A notable exception in this literature is Kök and Fisher (2007), who estimate a demand model that captures the effect of *permanent* changes in an assortment on sales, which is then used to choose the number of facings of products to be included in an optimal assortment. Our focus is different in that we measure the effect of *temporary* changes of product availability on sales. Another difference is that they use ad-hoc functional forms for the substitution patterns,

while in our approach product substitutions are based on utility maximization principles.² Recent work by Vulcano et al. (2008) estimates substitution effects induced by out-of-stocks. Their method is relatively simple to implement but requires precise information on product availability for every purchase. In contrast, our methodology can be used with partial information on product availability, but is more computationally intensive.

In the context of the marketing literature, models that use individual data (e.g. Fader and Hardie (1996); Rossi et al. (1996)) or aggregate data (e.g., Besanko et al. (2003), Jiang et al. (2007)) to estimate consumer demand are ubiquitous; but most of this work does not take into account the effect of product availability on demand. Studies that use customer panel-data to estimate the effect of out-of-stocks (e.g., Campo et al 2003; Swait and Erdem 2002) usually infer product availability based on sales data.³ A simple method to identify an out-of-stock event from sales information is to *assume* that if no sales of an item are observed during a certain time period then the product was not available. This rule is not always appropriate because: (1) zero sales are not always caused by out-of-stocks, especially for slow moving items; and (2) out-of-stocks do not always lead to zero sales (i.e., an out-of-stock may occur in the middle of a time period after the remaining units of an item were sold). For these reasons, it is important to incorporate more direct measures of product availability when estimating customers' preferences and reactions to stock-outs.

To our knowledge, Anupindi et al. (1998) are the first to estimate the effect of out-of-stocks on customer demand using actual measures of product availability. Their model accounts for lost sales and product substitution effects, but unless further restrictions are imposed to this model, it is necessary to estimate a different set of arrival rates for every possible set of available alternatives faced by a customer. Therefore, the number of parameters rapidly grows with the number of alternatives and in order to fully characterize customers' propensity to buy each alternative, it is necessary to have observations for every possible choice set.⁴ In order to address this estimation issue, Anupindi et al. (1998) focus on the case of *one-stage substitution*: the demand for an out-of-stock product is partially transferred to another product, but if this product is also out-of-stock, then this demand is lost. This restriction substantially reduces the number of parameters that need

²Kök and Fisher (2007) also develop a model to estimate stock-out based substitution (temporary changes in availability), which they do not estimate with their empirical data.

³An exception is Anderson et al. 2006, who use sales data from a catalogue retailer.

⁴Denoting by J the number of alternatives, the number of non-empty choice sets is equal to $2^J - 1$.

to be estimated, but it may not be appropriate in some applications.

A different approach is proposed by Kalyanam et al. (2007), who use a reduced-form model to capture cross-item substitutions through a finite number of categorical variables. A limitation of this approach is that the number of parameters to be estimated grows with the square of the number of options in each categorical variable. In addition, the model does not incorporate time-varying or continuous product characteristics such as price and therefore it cannot be used to estimate price-elasticities of demand, which are needed to evaluate alternative policies to manage a product category or to mitigate the costs of out-of-stocks.

In addition, our work is also related to the literature on consideration sets (e.g., Hauser and Wernerfelt (1990); Roberts and Lattin (1991); Andrews and Srinivasan (1995)). This body of research focuses on modeling the set of alternatives that are *considered* by a consumer when estimating consumer preferences. The estimation of consideration sets is usually implemented by inferring them from choice data and parametric assumptions about the consideration and choice processes. This presents an important estimation challenge, given that data from a consumer who never chooses a given product can be rationalized in at least two ways: i) the product does not enter the consumer's consideration set or ii) the product is considered by the consumer, but she assigns a large disutility to this product. This highlights the importance of using *auxiliary* information to separately identify consumer preferences and consideration sets (e.g., asking consumers which alternatives they considered as in Roberts and Lattin (1991)). In our paper, we face a similar problem as we are interested in estimating the underlying set of alternatives that are available to each customer. We estimate this set of alternatives combining sales data with information on product inventory. This enables us to characterize the evolution of the choice set as a function of variations in product availability.

In the economics literature there has been an extensive development of methods to estimate demand based on random utility maximization models (RUM) using *market-level* sales data (e.g. Berry (1994) and Berry et al. (1995)). RUM models can be very effective at providing a parsimonious characterization of consumer preferences reducing the number of parameters to be estimated. In the context of out-of-stocks, Bruno and Vilcassim (2008) extend the methodology in Berry et al. (1995) to incorporate external information about product availability, showing that neglecting the effects of out-of-stocks leads to substantial biases in the estimation. Obtaining measures of availability at the market level is difficult; for example, Bruno and Vilcassim use the all commodity volume (ACV)

weighted distribution information, which may not be a good proxy of local availability to consumers (Dolan and Hayes (2005)). We overcome this problem by using *store* data as opposed to *market level* information and, therefore, we are able to obtain more precise measures of product availability to exploit variations in the choice set due to out-of-stocks. In addition, given the complexity of the methodology proposed by Bruno and Vilcassim (2008), these authors assume that the probability of one product being sold-out is independent of the availability of other products. This assumption, may not be appropriate for demand models that do not satisfy the independence of irrelevant alternatives assumption.⁵ We explicitly address this issue by not imposing restrictive assumptions on the joint probability of different products being sold-out. In addition and as mentioned by Chintagunta and Dubé (2005), a limitation of the basic methodology developed in Berry (1994) and most of its extensions (including Bruno and Vilcassim (2008)) is that it cannot be used when some of the products have zero sales, which is not uncommon for store-level data of slow-moving categories.

Closest to our work is Conlon and Mortimer (2007), who estimate the substitution effects induced by stock-outs using a random utility model and partial data on product availability. They use an EM algorithm to account for the missing data on product availability faced by each customer. However, the expectation step becomes difficult to implement when multiple products are simultaneously out-of-stock (as in the case of our empirical application). As we will explain in the next section, our approach can be applied to cases where a large number of products become out-of-stock without increasing the complexity of the estimation method.

In terms of Bayesian methods to estimate demand using aggregate data, and as we mentioned in the introduction, our estimation approach presented here is based on an extension of the methodology in Chen and Yang (2007), Musalem et al. (2007) and Musalem et al. (2008). In contrast to prior work in this area (e.g., Musalem et al. (2008)), we not only use aggregate sales data, but also use periodic information about product inventory. Conditioning on both types of information, we *jointly* estimate the distribution of consumer preferences and product availability, which enables us to estimate the impact of out-of-stocks on consumer choices.

Finally, we note that other authors have used controlled laboratory experiments (e.g. Fitzsimons

⁵For example, it seems more likely for an out-of-stock of caffeine-free diet Coke to trigger a subsequent out-of-stock for a similar product (e.g., caffeine-free diet Pepsi) than one for a dissimilar one (e.g., Sprite).

(2000)), field experiments (e.g., Anderson et al. (2006)), and questionnaires (e.g., Campo et al. (2003b)) to estimate customer response to out-of-stocks. We focus instead on developing methods that use data routinely collected by store managers.

3 Model and Methodology

This section describes the customer choice model and the estimation method. The demand model is based on utility maximization principles and can be estimated with sales data from multiple stores combined with (partial) information about product availability. We also discuss identification and endogeneity issues and show results from a simulation experiment to test and validate our methodology. Finally, we show how the model can be used to estimate lost sales.

3.1 Customer Choice model

We start by specifying a random-coefficients multinomial logit (MNL) model for product choice within a category, a demand specification widely used in the economics and marketing literature (e.g. Train (2003), Chintagunta et al. (1991)). Consider a customer i that visits store m during time period t and chooses to buy a single unit from among the alternatives in the set $\mathcal{J} = \{1, \dots, J\}$ or chooses not to purchase (no-purchase option). We specify the utility of purchasing product $j \in \mathcal{J}$ as follows:

$$\begin{aligned} U_{ijtm} &= V_{ijtm} + \varepsilon_{ijtm} \\ &= \beta'_{itm} X_{jtm} + \xi_{jtm} + \varepsilon_{ijtm} \end{aligned} \tag{1}$$

where X_{jtm} is a vector of covariates that may include product characteristics, price and other marketing variables. The vector of random coefficients β_{itm} , which varies across customers, describes individual preferences and is assumed to be distributed according to a multivariate normal with mean $\bar{\theta}_m$ and covariance matrix Σ . The mean of individual preferences for a given store m is specified as:

$$\bar{\theta}_m = \theta' Z_m \tag{2}$$

where Z_m is a vector of characteristics for store m , which enables a researcher to capture *observed* heterogeneity across customers in different markets and θ is a matrix of coefficients. The variance-covariance matrix Σ captures *unobserved* customer heterogeneity, which is assumed for simplicity and parsimony to be constant across stores. In addition, ξ_{jtm} is a common demand shock that is introduced to capture market factors unobserved to the researcher that change the utility of all customers for alternative j in period t and market m . The random vector $\xi_{tm} = (\xi_{1tm} \dots \xi_{Jtm})$ is assumed to follow a multivariate normal distribution with zero mean and covariance matrix Σ_ξ . For simplicity, we assume that $\Sigma_\xi = \sigma_\xi^2 \cdot I_J$ (where I is the identity matrix with J rows and columns). The inclusion of these demand shocks helps to prevent overfitting problems when aggregate data are used to estimate the model (Berry (1994)). Finally, ε_{ijtm} is an individual-specific demand shock, modeled as an i.i.d. random variable from an extreme value distribution.

The probability that a customer chooses a given brand depends on the set of alternatives available to a customer. In particular, due to the occurrence of out-of-stocks, the choice set of a customer may not include all the products in the set \mathcal{J} . Product availability is characterized in our model by the vector $a_{itm} = (a_{itm}^1 \dots a_{itm}^J)$, where $a_{itm}^j = 1$ if customer i visiting store m in period t finds product j available, and $a_{itm}^j = 0$ otherwise. Without any loss of generality, we index customers in each period by their order of arrival to the store.

Let y_{itm} denote the product chosen by customer i in period t and market m and let $U_{i0tm} = \varepsilon_{i0tm}$ be the utility of the no-purchase option (the subscript 0 denotes the no-purchase option). The probability that a customer i facing availability a_{itm} purchases product j is given by:

$$p_j(\beta_{itm}, a_{itm}, \xi_{tm}) \equiv \Pr(y_{itm} = j | \beta_{itm}, a_{itm}, \xi_{tm}, X) = \frac{a_{itm}^j \cdot \exp(\beta_{itm} X_{jtm} + \xi_{jtm})}{1 + \sum_{k \in \mathcal{J}} a_{itm}^k \cdot \exp(\beta_{itm} X_{ktm} + \xi_{ktm})}, \quad (3)$$

while the probability of choosing the no purchase option is given by $p_0(\beta_{itm}, a_{itm}, \xi_{tm}) = 1 - \sum_{j \in \mathcal{J}} p_j(\beta_{itm}, a_{itm}, \xi_{tm})$. Note that this customer model is similar to the one used by Bruno and Vilcassim (2008).⁶ In addition, we assume that aggregate store-level sales data, S_{jtm} , are observed. Given an observed number of customers making purchase decisions in each period, N_{tm} , the number of customers choosing the no purchase option is given by $S_{0tm} = N_{tm} - \sum_{j=1}^J S_{jtm}$.

⁶Campo et al. (2003a) make use of a similar extension to the MNL model, but do not incorporate random coefficients.

As is evident from equation (3), the estimation of the model requires information about the set of products available to each customer visiting the store (a_{itm}). As we mentioned in the introduction, this information is rarely available in practice. Instead, we assume that product availability is observed at the beginning and at the end of each period.⁷ We use the notation $I_{tm} = (I_{tm}^1, \dots, I_{tm}^J)$ to denote inventory at the beginning of the period; \hat{I}_{tm} is used for product inventory at the end of the period. If there are replenishments, the definition of the time periods is such that they occur just before the beginning of the period and are therefore accounted in I_{tm} . When a store m runs out-of-stock for some product j during period t , we observe that $I_{tm}^j > 0$ and $\hat{I}_{tm}^j = 0$ but we do not know the *exact time* of the out-of-stock. This missing piece of information is important as it determines how many customers were exposed to the out-of-stock. For example, if the first customer visiting the store purchased the last unit of product j , then all other customers were exposed to this out-of-stock. In contrast, if the last customer visiting the store purchased the last unit of product j , then no customers were affected by this out-of-stock. Because this information is not directly observable, we don't know a priori whether customers buying product $k \neq j$ during that period chose this product because it was their most preferred item or because it was their second preferred option when product j , the preferred product, was not available. Consequently, given that the purchasing probability depends on the availability vector a , the likelihood function for this problem cannot be expressed only in terms of sales data and the coefficients of the utility function of each customer, as in standard applications of choice models to purchase data.

These difficulties caused by sparse product availability data were recognized by Anupindi et al. (1998). They develop a model where the time epoch of each out-of-stock is treated as missing data, and use an EM algorithm for estimation. The EM algorithm facilitates the search for the maximum likelihood estimators. However, the closed-form expressions used in their expectation step become complicated as the number of out-of-stocks in a single period increases. In fact, they derive expressions for no more than two out-of-stocks taking place in the same period. The Bayesian estimation method described in the next section circumvents this limitation, allowing for simultaneous out-of-stocks for an arbitrary number of products in each period and for a much more general class of problems that do not require closed-form expressions.

⁷Note that store sales are usually monitored more frequently than product availability. In those cases, our method can still be applied after aggregating sales to form a sales series with the same frequency as the availability data.

3.2 Bayesian estimation

The proposed methodology uses a data augmentation approach to overcome the difficulties generated by incomplete data on product availability. Let w_{ijtm} be a choice indicator equal to 1 if $y_{itm} = j$ and equal to zero otherwise and denote by W_{tm} the set of choice indicators w_{ijtm} for all customers and alternatives corresponding to period t and market m . There are many configurations of W_{tm} which could have generated the observed data. Denoting by Ω_{tm} the set of all the configurations of choice indicators consistent with the aggregate sales data for period t and market m , $W_{tm} \in \Omega_{tm}$ if and only if:

$$\sum_{i=1}^{N_{tm}} w_{ijtm} = S_{jtm}, \text{ for all } j \in \mathcal{J}. \quad (4)$$

Our data augmentation approach treats each choice (y_{itm}) and choice indicator (w_{ijtm}) as missing data. The main advantage of using these augmented data is that they determine unambiguously the availability observed by each customer, as illustrated by the following example. Suppose that two products, B and C, are available at the beginning of a given time period, but only product C is available at the end of the period. We observe total period sales for each product: $S_B = 1$ and $S_C = 2$. Therefore, we know that initially there was exactly 1 unit of B and at least 3 units of C in inventory. Furthermore, suppose we also observe that $N = 6$ customers visited the store during this period; therefore, $N - S_B - S_C = 3$ customers chose the no-purchase option denoted by 0. We index customers by their order of arrival and the choice of each customer by $y_i \in \{B, C, 0\}$. The vector $y = (y_1, \dots, y_6)$ characterizes the order of purchases during the period. If $y = (C, 0, C, 0, 0, B)$, then all 6 customers found both products available. If we consider instead a different value of the vector of individual choices by swapping the choices of customers 1 and 6, we obtain $y = (B, 0, C, 0, 0, C)$. In this case, customers 2 through 6 found product B out-of-stock. Also note that had we swapped instead the choices of customers 1 and 2, $y = (0, C, C, 0, 0, B)$, the set of products available to each customer would be unchanged.

The relationship between individual choices and availability can be expressed, in general terms, as follows. Given w and the observed data (sales on each period, S_{jtm} , and inventory information at the beginning and end of each period, I_{tm} and \hat{I}_{tm}), product availability for each customer can

be determined as:

$$\alpha_{itm}^j = \begin{cases} 1 & \text{if } I_{tm}^j > 0, \hat{I}_{tm}^j > 0, \\ 0 & \text{if } I_{tm}^j = \hat{I}_{tm}^j = 0, \\ 1 \left\{ \sum_{k=1}^{i-1} w_{kjt} < S_{jtm} \right\} & \text{if } I_{tm}^j > 0, \hat{I}_{tm}^j = 0. \end{cases} \quad (5)$$

where the summation over customers $1, \dots, i-1$ includes all customers arriving prior to customer i .

Using the augmented data, the likelihood function for store m and period t is given by:

$$L_{tm}(\theta, \Sigma, \Sigma_\xi) = \sum_{W_{tm} \in \Omega_{tm}} \int_{\xi_{tm}} \prod_{i=1}^{N_m} \left[\int_{\beta_{itm}} \prod_{j=0}^J p_j(\beta_{itm}, a_{itm}, \xi_{tm})^{w_{ijt}} \phi(\beta_{itm}; \theta' Z_m, \Sigma) d\beta_{itm} \right] \phi(\xi_{tm}; 0, \Sigma_\xi) d\xi_{tm}, \quad (6)$$

where $\phi(x; \mu, \Sigma)$ denotes the multivariate normal density with mean μ and variance covariance matrix Σ . The overall likelihood (including all stores and periods) is given by $L = \prod_{m=1}^M \prod_{t=1}^T L_{tm}(\theta, \Sigma, \Sigma_\xi)$. This likelihood function is difficult to compute due to the summation over all possible configurations of choice indicators (W_{tm}) consistent with the aggregate data (see equation 6).⁸ Instead of using a Maximum Likelihood (ML) approach, our strategy is to apply data augmentation together with Markov Chain Monte-Carlo methods (MCMC) to facilitate the estimation of the parameters of the model. This approach is described in the next paragraphs.

In contrast to ML, which seeks to find a point estimate for the parameters of interest $(\theta, \Sigma, \Sigma_\xi)$, a Bayesian approach seeks to estimate the posterior distribution for the parameters, given the observed data. Towards this objective, we define the hyper-prior distribution $\pi(\theta, \Sigma, \Sigma_\xi)$ which formalizes the researcher's prior beliefs about the demand parameters. We also write the likelihood of the augmented data for customer i in store m and period t as:

$$L_{itm}(\beta_{itm}, \xi_{tm} | a, w) = \prod_{j=0}^J [p_j(\beta_{itm}, a_{itm}, \xi_{tm})]^{w_{ijt}}. \quad (7)$$

The posterior density of $(\theta, \Sigma, \Sigma_\xi)$, $\{\beta_{itm}\}$, and the augmented data (a, w) given the *observed* data

⁸Another complication of doing maximum likelihood is that the integrals in (6) cannot be expressed in closed analytical form and have to be approximated through simulation. The resulting objective function may be highly non-linear and the search for a global optimum becomes challenging.

is proportional to:

$$\begin{aligned}
p(\theta, \Sigma, \Sigma_\xi, a, w, \{\beta_{itm}\} | S) &\propto \pi(\theta, \Sigma, \Sigma_\xi) \prod_t \prod_m 1 \left\{ \sum_{i=1}^{N_{tm}} w_{ijtm} = S_{jtm}; j \in \mathcal{J} \right\} \phi(\xi_{tm}; 0, \Sigma_\xi) \\
&\times \prod_{i=1}^{N_{tm}} \phi(\beta_{itm}; \theta' Z_m, \Sigma) \cdot L_{itm}(\beta_{itm}, \xi_{tm} | a, w)
\end{aligned} \tag{8}$$

where the indicator function incorporates constraint (4) into the probability model.

We seek to estimate the posterior (8) using MCMC methods. Conditional on the augmented variables (a, w) , the parameters $\{\beta_{itm}\}$, θ , Σ and Σ_ξ can be sampled using existing Bayesian methods for individual level data (e.g. Allenby and Rossi (2003)). In particular, we combine these existing methods with a sampling scheme for the missing data a and w . Following Musalem et al. (2007), posterior samples for these variables can be obtained by deriving the full-conditional distribution of the augmented choices and availability data. Note that the full-conditional distributions have a simple structure if we partition the set of customers into pairs and consider the distribution of the choices of customers in a given pair, holding constant all other choices.⁹ This enables us to define a mechanism to sequentially sample a and w from their posterior distribution for each pair of customers. We illustrate this sampling scheme with an example (see Table 1).

To simplify notation, we suppress time (t) and store (m) indexes for this example. Consider the set of products formed by $\{B, C\}$ plus the no-purchase option 0. The observed data are $S_B = 1$, $S_C = 2$, $N = 6$, $I^B = 1$, $I^C = 3$, $\hat{I}^B = 0$ and $\hat{I}^C = 1$. We start with an initial assignment of choices given by y (see column 2 in Table 1), which is consistent with the observed sales data. We now show how to sample a new vector y' . We randomly generate a partition of the set of customers into 3 pairs: for example, $\{1, 3\}$, $\{2, 4\}$, $\{5, 6\}$. We are interested in calculating the full conditional distribution of the choices of customers in one of these pairs, say $\{2, 4\}$. Conditioning on the choices of all other customers $\{1, 3, 5, 6\}$, denoted by y_{-24} , and the observed data, there are only two vectors of choices that are consistent with the sales data: the current choices (y) and the choice vector y^* (see column 5 in Table 1) where the choices of customers 2 and 4 are swapped.

Note that when exchanging the choices of customers 2 and 4, the sets of available products

⁹As mentioned in Musalem et al. (2007), one could also consider partitioning the set of customers into larger groups (e.g., triplets or quadruplets). This may provide a more efficient approach to estimate the posterior distribution of the model parameters, but it increases the computational cost of each simulation.

to customers arriving between customers 2 and 4 may endogenously change, but the choice set of all other customers remains unchanged. In the example, customers 2, 3 and 4 find alternative B available under the current configuration of choices (y , see column 3 in Table 1), while only customer 2 finds this alternative available under the new configuration (y^* , see column 6 in Table 1). In addition, product availability for all other customers (1, 5 and 6) remains constant. Therefore, swapping the choices of customers 2 and 4 can only affect the likelihood of customers 2, 3 and 4. When sampling from this full conditional distribution, the probability of generating a draw where the choices of 2 and 4 are exchanged is:

$$\Pr(y^*|y_{-24}) = \frac{\prod_{i=2}^4 L_i(\beta_i, \xi|a^*, w^*)}{\prod_{i=2}^4 L_i(\beta_i, \xi|a^*, w^*) + \prod_{i=2}^4 L_i(\beta_i, \xi|a, w)}$$

where a^*, w^* is the augmented data implied by the vector of choices y^* . The new draw is equal to the previous draw y with probability $1 - \Pr(y^*|y_{-24})$.

Based on this intuition, the Gibbs sampling scheme for the missing data (a, w) of each store-period can be formalized as follows:

Step 1: In a given iteration r , randomly generate a partition of customer-pairs, denoted P_{tm} , of the N_{tm} customers.

Step 2: For a given pair $(l, k) \in P_{tm}$, exchange the choices of customers l and k . Generate the new choice vector y^* and binary variables a^* and w^* by swapping the choices of customers in pair (l, k) . The new variables are accepted with probability:

$$\Pr(\text{accept}) = \frac{\prod_{i=l}^k L_i(\beta_i, \xi|a^*, w^*)}{\prod_{i=l}^k L_i(\beta_i, \xi|a^*, w^*) + \prod_{i=l}^k L_i(\beta_i, \xi|a, w)}$$

If accepted, update $y_{tm}^{(r+1)} \leftarrow y_{tm}^*$, $a^{(r+1)} \leftarrow a^*$ and $w^{(r+1)} \leftarrow w^*$, otherwise set $y_{tm}^{(r+1)} \leftarrow y_{tm}^{(r)}$, $a^{(r+1)} \leftarrow a^{(r)}$ and $w^{(r+1)} \leftarrow w^{(r)}$. Repeat Step 2 for a new pair in P_{tm} until no pairs are left.

The Gibbs sampling scheme ensures that every feasible choice vector y can be sampled with positive probability. Therefore, when combined with a proper sampling scheme for $\{\beta_{itm}\}$, $\{\xi_{tm}\}$, θ , and Σ , it generates an irreducible Markov chain with stationary probability distribution equal to the posterior distribution of the parameters of the model.

3.3 Identification

In this subsection we discuss issues related to identification of our model. In the economics and marketing literature, variations in prices, promotions and other marketing variables are usually exploited to identify and estimate demand models (e.g., Villas-Boas and Winer (1999)). In the case of this paper, another important source of identification arises from changes in product availability. In particular, observing customer behavior under different choice sets provides useful information to uncover the patterns of substitution among different alternatives. One difficulty is that the variation in product availability is not perfectly observed and this may lead to some challenges in the identification of the model parameters.

Consider the following example. A few units of product A are available at the beginning of a period, and at the end of period the product is out-of-stock. There are two possible explanations that fit these data. In the first one, customers have a low utility for product A (relative to other products), which generates just enough demand for all the units initially available. In this case, the out-of-stock of product A occurs towards the end of the period. A second explanation is that customers place a high utility on product A, and the first customers visiting the store purchase product A, generating an out-of-stock early in the period. Without additional information, we cannot distinguish between these two alternative explanations. However, we may also observe data for : (1) the sales for other products with similar characteristics to product A, which were available throughout the period; and (2) sales during other periods where product A had full availability. This additional information is useful to separately identify the utility of product A and the latent distribution of product availability.¹⁰

An additional challenge that we face in the estimation of our model is related to the endogeneity of inventory. A retailer is likely to choose inventory levels based on demand projections, which we do not entirely observe in the data set of our application (the data set is described in Section 4). Hence, inventories could be correlated with the demand shocks (ξ_{jtm}), which could lead to biases in our estimation. To avoid this type of endogeneity, our methodology requires the demand shocks ξ_{jtm}

¹⁰The demand model could also be estimated by considering only the time periods for which the set of available alternatives does not change, i.e. periods in which the set of products available at the beginning and end of the period is the same. Although this would certainly reduce the complexity of the estimation method, this approach exhibits two important disadvantages. First, one would be forced to discard valuable information about the customer demand and this approach would only work if there are enough time periods that meet this requirement. In particular, for longer time periods, it will be less likely to observe constant availability sets.

to be unobservable to the retailer. This is not an unreasonable assumption when the model is used by a store manager, since she is likely to have access to the information used to determine inventory levels and this information could also be used to estimate the demand model. For our empirical application, we include covariates that capture demand seasonality, which we believe capture most of the demand changes anticipated by the store manager that may have occurred during the study.

Furthermore, note that our model uses information on availability (a_{itm}) and does not use data on inventory levels directly. Hence, endogeneity would only be an issue if the demand shocks ξ_{tm} affect the distribution of availability. Models in inventory management suggest that the probability of an out-of-stock of a product does not depend on its specific demand characteristics (e.g. mean or standard deviation) but rather on the relative costs of holding inventory versus the costs of having an out-of-stock. The cost of holding inventory depends mainly on the cost of capital, depreciation and storage, which are unrelated to demand. The cost of an out-of-stock is affected, among other things, by the product’s margin, which may be related to the demand through price. Therefore, this implies that it is important to include prices in the utility specification not only to be able to estimate price elasticities, but also to avoid this type of endogeneity bias.

Another potential endogeneity problem may arise from serially correlated demand shocks (ξ_{jtm}). Today’s inventory depends on yesterday’s demand, which induces a correlation between current inventory and previous demand shocks. In addition, if demand shocks are autocorrelated, current demand shocks will be correlated with previous shocks that are in turn correlated with today’s inventory, leading to an endogeneity problem. As we previously mentioned, we do not expect to observe autocorrelation in demand shocks in our application once we control for seasonality. Nevertheless, the methodology presented here could be extended to allow for demand shock persistence, which would eliminate this form of endogeneity.

3.4 Numerical Experiment

In this subsection, we test the proposed methodology using simulated data. We generated sales and inventory data for $J = 10$ purchase alternatives and a no-purchase option available in $M = 12$ markets and $T = 15$ periods. We included four variables (x_1, \dots, x_4) in the utility function. We use random coefficients for x_3 and x_4 and assume fixed (i.e., constant across customers in a given market) coefficients for x_1 and x_2 . The first two variables correspond to two brands dummies, one

for alternatives 1, 2 and 3 (x_1), and another for alternatives 4, 5 and 6 (x_2). The third variable is a dummy variable equal to 1 for all purchase alternatives (x_3), while the fourth variable is generated from a normal distribution (x_4). In order to replicate some of the features of the data set used in our empirical application, the continuous variable is generated so that its values for a given brand and market are the same across all time periods. Accordingly, the values of x_4 for each brand and market are generated from a normal distribution with mean equal to 2 and variance equal to 1.

Based on these four explanatory variables, customer coefficients for a given market m are generated from a 4-dimensional multivariate normal distribution with mean $\theta_m = \theta' Z_m$ as in equation (2) and variance Σ , where Z_m represents a 2-dimensional vector of demographic variables. In terms of the demographics, the first variable z_{1m} is equal to 1 for all markets (intercept), while the second variable z_{2m} is generated from a uniform distribution in the interval $[-1.5, 1.5]$. The true value of θ corresponds to:

$$\theta = \begin{bmatrix} 2.0 & 1.5 & -3.0 & -2.5 \\ 0.5 & -0.5 & 0.0 & 0.7 \end{bmatrix}.$$

The variance of the random coefficients (Σ) is equal to a diagonal matrix with elements equal to 0.8 and 2.0 for x_3 and x_4 respectively and elements equal to zero for the two variables with fixed coefficients (x_1 and x_2). The size of each market (N_m) is generated by taking the integer part of a uniform random variable defined on the interval $[0, 300]$. In addition, common demand shocks for each alternative in each period and market (ξ_{jtm}) are generated from a normal random variable with zero mean and variance σ_ξ^2 equal to 0.5. Finally, initial inventory levels for each alternative are generated by taking the integer part of a uniform random variable in the interval $[0, I]$. We use three different values of I (10, 60 and 400) which lead to different inventory service levels. In the first case ($I = 10$), each alternative stocks out on average in 28.9% of the time periods across all markets, while this fraction corresponds to 8.7% and 1.3% for the second and third cases, respectively.

For each of these three data sets ($I = 10$, $I = 60$ and $I = 400$) we estimated θ , Σ and σ_ξ^2 using the method described in the previous section and based on the aggregate data available for each period (sales and initial inventory). We also estimated these parameters ignoring the occurrence of out-of-stocks (assuming all products were available for every single period and market). In both

cases, we used the following weakly informative prior distributions: $\theta \sim N(0, 100 \mathbf{I}_8)$, $\Sigma_{jj} \sim$ scaled inverse chi-square ($df = 3$, scale= 1) and $\sigma_{\xi}^2 \sim$ scaled inverse chi-square ($df = -1$, scale= 0.01). The results are presented in Table 2 and they are based on a single run of 100,000 iterations from a Markov Chain Monte Carlo (MCMC) sampler, where the last 50,000 iterations are used for parameter estimation.

From the results in Table 2 we observe important differences in terms of parameter inference comparing the case where out-of-stocks are modeled with the one in which out-of-stocks are ignored. Considering the results in the first three blocks of Table 2, it is evident that when out-of-stocks are ignored, our inferences about the model parameters are biased, especially when out-of-stocks are more frequent (see first and second blocks of Table 2). Specifically, when $I = 10$ the 95% posterior probability intervals for most of the components of θ do not cover the corresponding true values. For example, θ_{z_1, x_3} which is used to derive the mean utility of all the purchase options is underestimated when out-of-stocks are ignored (true value= -3 , 95% posterior probability interval: $[-4.60, -3.78]$). In addition, the heterogeneity in the random coefficients for the continuous variable (Σ_{x_4, x_4} , true value= 2, 95% posterior probability interval: $[0.54, 1.46]$) is underestimated. As expected, the estimation results improve as out-of-stocks become less frequent. In particular, in the third case, where out-of-stocks are on-average observed only in 1.3% of the time periods for each alternative ($I = 400$), the results are very similar across both models (see third and sixth blocks in Table 2).

In terms of the full model, the results in the last three blocks in Table 2 show that the method recovers well the original parameters under each of the three scenarios. In fact, the posterior means are close to their true values (on average within 0.8 posterior standard deviations from the corresponding true values). In addition, in all but four cases (θ_{z_1, x_3} and Σ_{x_3, x_3} in the fifth block and θ_{z_1, x_2} and Σ_{ξ} in the sixth block of Table 2), the true values are contained within the 95% posterior probability intervals. Finally, it is important to mention that we also conducted a simulation study by generating 50 data sets for the case where $I = 10$. We estimated the model parameters using the proposed method and also ignoring the occurrence of out-of-stocks. The results of this simulation study, confirm the basic findings discussed in this subsection (please refer to Online Appendix A).

Table 1: Gibbs sampling example

Customer i	y_i	a_i^B	a_i^C	y_i^*	a_i^{B*}	a_i^{C*}
1	C	1	1	C	1	1
2	0	1	1	B	1	1
3	C	1	1	C	0	1
4	B	1	1	0	0	1
5	0	0	1	0	0	1
6	0	0	1	0	0	1

Table 2: Results: Estimated posterior mean, standard deviation and quantiles for θ , Σ and σ_ξ^2 ignoring (benchmark model) and accounting for (full model) the occurrence of out-of-stocks.

Model		θ_{z_1, x_1}	θ_{z_1, x_2}	θ_{z_1, x_3}	θ_{z_1, x_4}	θ_{z_2, x_1}	θ_{z_2, x_2}	θ_{z_2, x_3}	θ_{z_2, x_4}	Σ_{x_3, x_3}	Σ_{x_4, x_4}	σ_ξ^2
Benchmark I=10	mean	1.45	1.30	-4.13	-1.52	0.41	-0.39	0.00	0.53	0.66	0.93	0.48
	std.dev.	0.09	0.09	0.22	0.22	0.09	0.10	0.10	0.06	0.39	0.24	0.06
	2.5%	1.28	1.14	-4.60	-1.96	0.23	-0.59	-0.19	0.43	0.24	0.54	0.36
	50.0%	1.45	1.30	-4.11	-1.49	0.41	-0.39	0.01	0.53	0.49	0.90	0.48
	97.5%	1.62	1.48	-3.78	-1.13	0.60	-0.18	0.19	0.65	1.58	1.46	0.60
Benchmark I=60	mean	1.82	1.37	-3.39	-2.14	0.41	-0.50	-0.02	0.70	0.85	1.56	0.61
	std.dev.	0.07	0.08	0.09	0.12	0.10	0.10	0.09	0.06	0.24	0.19	0.05
	2.5%	1.68	1.22	-3.58	-2.38	0.22	-0.71	-0.18	0.59	0.47	1.23	0.52
	50.0%	1.82	1.37	-3.40	-2.13	0.41	-0.50	-0.02	0.70	0.82	1.54	0.61
	97.5%	1.96	1.54	-3.21	-1.94	0.59	-0.30	0.14	0.81	1.41	1.95	0.72
Benchmark I=400	mean	1.94	1.70	-3.09	-2.54	0.38	-0.62	0.06	0.71	0.52	2.04	0.60
	std.dev.	0.09	0.09	0.11	0.11	0.09	0.09	0.08	0.04	0.22	0.17	0.05
	2.5%	1.77	1.52	-3.32	-2.76	0.21	-0.80	-0.12	0.62	0.20	1.75	0.51
	50.0%	1.95	1.70	-3.09	-2.54	0.38	-0.62	0.06	0.71	0.49	2.03	0.60
	97.5%	2.12	1.88	-2.89	-2.35	0.57	-0.45	0.21	0.79	1.06	2.41	0.69
Full I=10	mean	1.92	1.70	-3.24	-2.41	0.50	-0.53	-0.08	0.79	0.91	1.91	0.61
	std.dev.	0.11	0.11	0.25	0.27	0.10	0.11	0.11	0.08	0.39	0.40	0.08
	2.5%	1.70	1.49	-3.75	-3.03	0.29	-0.76	-0.32	0.64	0.41	1.25	0.47
	50.0%	1.92	1.69	-3.24	-2.38	0.50	-0.53	-0.08	0.79	0.80	1.85	0.61
	97.5%	2.12	1.91	-2.71	-1.94	0.70	-0.31	0.12	0.95	1.85	2.87	0.78
Full I=60	mean	1.97	1.48	-3.36	-2.36	0.53	-0.51	0.03	0.73	1.76	1.87	0.55
	std.dev.	0.07	0.09	0.14	0.13	0.10	0.11	0.10	0.06	0.52	0.22	0.05
	2.5%	1.82	1.31	-3.64	-2.62	0.33	-0.72	-0.14	0.61	0.99	1.47	0.45
	50.0%	1.97	1.48	-3.35	-2.36	0.53	-0.51	0.03	0.73	1.67	1.86	0.54
	97.5%	2.11	1.66	-3.12	-2.12	0.71	-0.30	0.23	0.84	2.88	2.33	0.65
Full I=400	mean	1.97	1.73	-3.15	-2.51	0.41	-0.62	0.06	0.70	0.67	2.00	0.59
	std.dev.	0.08	0.09	0.09	0.11	0.09	0.10	0.10	0.04	0.31	0.18	0.05
	2.5%	1.81	1.56	-3.32	-2.72	0.22	-0.82	-0.14	0.62	0.19	1.69	0.51
	50.0%	1.97	1.73	-3.15	-2.51	0.41	-0.61	0.07	0.70	0.64	1.99	0.59
	97.5%	2.13	1.89	-2.98	-2.32	0.59	-0.41	0.27	0.78	1.37	2.37	0.69
	true	2.00	1.50	-3.00	-2.50	0.50	-0.50	0.00	0.70	0.80	2.00	0.50

3.5 Estimating Lost Sales

An important factor that determines inventory levels in retailing is the cost of shortage. This cost is closely related to the behavior of customers that encounter an out-of-stock. The cost of an out-of-stock increases with: (i) the markup of the product that sells out; and (ii) the fraction of customers that choose not to purchase after experiencing an out-of-stock.¹¹ The former is known by the store manager, but the latter is not directly observable. In what follows, we show how to use the model to estimate lost sales - the fraction of customers that chose not to purchase but would have purchased if some of the out-of-stock products had been available. This estimate can be used together with markup information to assign a dollar value to the cost of shortage.

Consider a customer i facing a set of available products A who chose the no-purchase option. We drop the time (t) and market (m) subscripts for ease of exposition. The set of out-of-stock products, denoted by $A^c = \mathcal{J} \setminus A$, includes product k . The probability that customer i would have purchased product k had it been available is given by:

$$\begin{aligned} \Pr(\text{choose } k \mid \text{choose 0 from } A) &= \int \Pr(\text{choose } k \mid \text{choose 0 from } A, \beta_i) f(\beta_i \mid \text{choose 0 from } A) d\beta_i \\ &= \int \frac{e^{V_{ik}}}{1 + \sum_{j \in A \cup \{k\}} e^{V_{ij}}} f(\beta_i \mid \text{choose 0 from } A) d\beta_i \end{aligned} \quad (9)$$

where β_i is the (random) “taste” coefficient of the customer, $f(\beta_i|\cdot)$ its conditional density and V_{ij} is the implied deterministic utility of this customer from purchasing product j . The first equality comes from conditioning on the customers preferences (β_i). The second equality comes from standard properties of the multinomial logit choice model.¹² Similarly, the probability of that the customer would have chosen any of the missing alternatives had they been available is given by:

$$\Pr(\text{choose } A^c \mid \text{choose 0 from } A) = \int \frac{\sum_{k \in A^c} e^{V_{ik}}}{1 + \sum_{j \in \mathcal{J}} e^{V_{ij}}} f(\beta_i \mid \text{choose 0 from } A) d\beta_i \quad (10)$$

¹¹Also note that the cost of an out-of-stock decreases with the markup of the products that capture the demand for the non-available products.

¹²Conditioning on choosing the no-purchase option, the random utility U_{i0} is distributed according to an extreme value distribution with mode (location) equal to $\ln(1 + \sum_{j \in A^c} e^{V_{ij}})$ and scale equal to 1. Hence, the probability of

preferring k over the no-purchase option is $e^{V_{ik}} \cdot \left(1 + e^{V_{ik}} + \sum_{j \in A} e^{V_{ij}}\right)^{-1}$.

The main challenge of computing lost sales via equations (9) and (10) is that it requires integration over those customers who were exposed to out-of-stocks and chose the no-purchase option. Our data augmentation strategy enables us to easily identify these customers and compute these integrals through simulation without further computational effort.

Specifically, in the context of MCMC estimation, in any given iteration we can use the draws for the augmented data on choices (y_i) and availability (a_i) to identify the set O_A of all customers facing availability A who chose the no-purchase option.¹³ Using these draws, we can estimate the lost sales (LS) induced by all of the out-of-stock products as:

$$LS = \sum_{i \in O_A} \frac{\sum_{k \in A^c} e^{V_{ik}}}{1 + \sum_{j \in \mathcal{J}} e^{V_{ij}}}. \quad (11)$$

Finally, it is important to note that one could erroneously consider an alternative approach (on the surface correct) to estimate the impact of out-of-stocks where the sales of each product are estimated assuming all products were available and computing choice probabilities for *every* customer using equation (3). This approach, however, is not accurate as it ignores the choices implied by the observed sales data. For example, suppose only product k was out-of-stock and 20% of the customers decided not to buy any products (no-purchase option). If product k had been available, then, from utility maximization, the only option these customers could have switched to is product k . This is determined by the fact that the customers' observed choices require the no-purchase option to dominate all other options (all options but product k). Consequently, it is necessary to condition on the observed sales data when estimating the consequences of out-of-stocks. As we previously mentioned, this conditional inference is facilitated by the augmentation of individual choices and availability data that we employ to estimate the parameters of the model.

The next section describes an application of the methodology based on data from the shampoo product category.

¹³Mathematically, this set can be expressed as $O_A = \{i : y_i = 0, a_i^j = 1 \forall j \in A, a_i^j = 0 \forall j \in A^c\}$.

4 Empirical Application: Demand Estimation in the Shampoo Product Category

We use data on shampoo purchases from six supermarket stores located in different regions of Spain to illustrate the methodology. The stores are owned by a major supermarket chain, with more than 400 stores in this country. The data set was collected from 6 of these stores located in different regions of Spain, including daily sales during 15 consecutive days (excluding two Sundays in which stores are closed) between May 13th and May 29th in 1999. The product category includes 24 different stock keeping units (SKUs), but not all SKUs were offered in each store during the study period. In addition to the sales data, information about product availability on the shelf was recorded at the beginning and end of each day. In total, 291 out-of-stocks were recorded across all days, stores and products.

The data also contain price information for every product on each store-day. From the 24 products in the data set, 16 products exhibited price variation across stores, but very few products exhibited temporal price variation within a given store during the study period. Table 3 displays the brands of each product and shows summary statistics for each of the products in terms of daily unit sales, prices (in Spanish Pesetas), availability (OOS) and number of stores where each product is offered (Stores).¹⁴

4.1 Model specification

We include eight covariates (X_{jtm}) in the utility function of each customer. Two of these covariates have random coefficients: price and purchase option. The latter is a dummy variable equal to 1 in every period for all options except the no-purchase option. This variable captures total category demand for shampoo products as higher coefficients associated to this variable raise the utility of all purchase alternatives in relation to the utility of the no-purchase option. In addition, we include dummy variables for the following brands: Pantene, Herbal Essence, Head & Shoulders (H&S), Cabello Sano and Timotei. Finally, we also include a dummy variable (Weekend) to control for seasonality (changes in demand during weekends). This variable is equal to 1 for all purchase alternatives for every time period corresponding to a Saturday and is equal to 0 otherwise (recall

¹⁴Availability is measured as the fraction of days in which an out-of-stock for a given product was observed.

Table 3: Brands and summary statistics for the shampoo data: brands, mean and standard deviation of daily sales and prices, fraction of out-of-stocks (OOS) and number of stores for each product.

Product	Brand	Daily Unit Sales		Price		OOS	Stores
		M	SD	M	SD		
1	Other	0.08	0.27	3.99	0.00	0.00	6
2	Other	0.00	0.00	4.97	0.04	0.11	3
3	Other	0.07	0.25	4.97	0.04	0.12	4
4	Timotei	0.12	0.39	6.03	0.10	0.27	6
5	Timotei	0.12	0.33	5.99	0.27	0.10	6
6	Timotei	0.17	0.43	6.03	0.10	0.26	6
7	Timotei	0.06	0.23	5.79	0.41	0.16	5
8	Other	0.03	0.18	4.72	0.32	0.38	4
9	Other	0.06	0.23	4.75	0.00	0.20	6
10	Pantene	0.17	0.37	4.50	0.11	0.33	6
11	Pantene	0.28	0.97	4.50	0.16	0.40	5
12	Pantene	0.29	0.89	4.50	0.11	0.39	6
13	Pantene	0.28	0.73	4.45	0.15	0.39	6
14	Other	0.07	0.29	4.99	0.00	0.17	6
15	Other	0.24	0.53	2.07	0.03	0.11	5
16	Other	0.24	0.53	4.85	0.00	0.03	6
17	Herbal Essence	1.03	1.47	4.55	0.19	0.08	6
18	H&S	0.32	0.68	5.20	0.22	0.36	6
19	H&S	0.21	0.63	5.17	0.19	0.31	6
20	Other	0.19	0.45	4.99	0.00	0.11	6
21	Cabello Sano	0.11	0.32	3.45	0.00	0.12	5
22	Cabello Sano	0.27	0.61	3.45	0.00	0.08	6
23	Cabello Sano	0.30	0.64	3.47	0.04	0.42	6
24	Other	0.06	0.23	4.25	0.00	0.17	5

that the stores are closed on Sundays). We note that this variable also enables us to explicitly consider in the utility function changes in demand that might be anticipated by the store manager.

We use demographic information to capture observed preference heterogeneity across stores (unobserved heterogeneity is captured through random coefficients). Specifically, we collected information on average declared income for each market. Accordingly, the vector of demographic variables for a given store m (Z_m) has two components: an indicator equal to 1 for all stores (intercept) and the standardized natural logarithm of income¹⁵. In addition, the size of the market for each store was estimated combining data on population and total consumption of shampoo in Spain¹⁶.

4.2 Results

Using the method introduced in Section 3 we estimated the model based on the covariates and demographic variables previously described. We also estimated a benchmark model that ignores out-of-stocks. In this model, every product offered in a given store is assumed to be available to all customers in every time period. However, products that were never available in a particular store were excluded from the choice set in the benchmark model. We computed the log-marginal likelihood for the full and benchmark models and obtained a log-Bayes Factor equal to 77.1, which gives very strong empirical support for the full model (Kass and Raftery (1995)). Table 4 reports the estimation results for the hyper-parameters θ , Σ and σ_ξ^2 under both model specifications.

From the results of the full model we note that products from the brands Pantene, Herbal Essence and H&S are on average more demanded than all other products (see the results for $\theta_{Intercept}$ in Table 4). In particular and consistent with the summary statistics in Table 3, Herbal Essence products are those that attract more customers on average. Furthermore, the results for θ_{Income} show lower intrinsic demand for Herbal Essence products in stores reaching customers with higher income levels, while the opposite is observed for H&S products. As expected, the weekend and price effect (see $\theta_{Intercept,Weekend}$ and $\theta_{Intercept,Price}$) show higher levels of demand during Saturdays and significant disutility for higher prices. In addition to the observed heterogeneity across stores

¹⁵We also experimented using additional demographic information (e.g., age), but the results and the fit of the model do not substantially change.

¹⁶The population data were downloaded from <http://www.ine.es>, while the consumption data were obtained from Euromonitor International. We also experimented with an alternative definition of the market size (approximately twice the size of the one used in this section) and obtained very similar results.

Table 4: Empirical Results: Estimated posterior mean, standard deviation, 25% and 75% quantiles for θ (mean of the demand coefficients), Σ and σ_ξ (iterations 200,001-400,000).

Variable		Full Model				Ignoring OOS			
		mean	(s.d.)	25%	75%	mean	(s.d.)	25%	75%
$\theta_{Intercept}$	Pantene	1.19	0.20	1.05	1.33	0.87	0.21	0.73	1.01
	Herbal Essence	2.14	0.26	1.97	2.32	2.22	0.26	2.04	2.40
	H&S	1.38	0.26	1.20	1.55	1.13	0.26	0.95	1.30
	Cabello Sano	0.65	0.24	0.49	0.80	0.45	0.24	0.28	0.61
	Timotei	0.53	0.31	0.33	0.74	0.41	0.31	0.21	0.62
	Weekend	1.08	0.22	0.93	1.22	1.17	0.25	1.00	1.32
	Purchase Option	-4.88	0.62	-5.21	-4.47	-4.92	0.63	-5.32	-4.53
	Price	-0.91	0.19	-1.02	-0.81	-0.99	0.24	-1.13	-0.82
θ_{Income}	Pantene	0.02	0.21	-0.12	0.16	0.05	0.21	-0.08	0.19
	Herbal Essence	-0.52	0.28	-0.70	-0.33	-0.57	0.27	-0.75	-0.39
	H&S	0.56	0.26	0.39	0.74	0.62	0.27	0.44	0.80
	Cabello Sano	0.12	0.24	-0.04	0.28	0.00	0.24	-0.16	0.16
	Timotei	-0.06	0.31	-0.27	0.15	0.04	0.31	-0.16	0.24
	Weekend	-0.26	0.21	-0.39	-0.12	-0.24	0.21	-0.38	-0.09
	Purchase Option	-0.54	0.50	-0.88	-0.19	-0.73	0.47	-0.97	-0.44
	Price	-0.01	0.12	-0.09	0.07	0.02	0.11	-0.05	0.08
Σ	Purchase Option	0.70	0.48	0.39	0.84	0.74	0.51	0.42	0.87
	Price	0.24	0.09	0.18	0.27	0.28	0.13	0.19	0.33
σ_ξ^2		1.10	0.22	0.95	1.24	1.25	0.26	1.07	1.41

captured by θ_{Income} , the results for Σ show the magnitude of the unobserved heterogeneity within stores in terms of total demand (Purchase option) and price sensitivity (Price). The variances for these random coefficients are estimated as 0.70 and 0.24, respectively. The results for the variance of the demand shocks (σ_{ξ}^2) suggest that these unobserved demand effects are substantial (the posterior mean of σ_{ξ}^2 is estimated as 1.10).

Finally, when comparing the full and benchmark models, we note that the posterior means of $\theta_{Intercept}$ for the brands Pantene, H&S and Cabello Sano are smaller under the benchmark model than in the full model (these posterior means are approximately one posterior standard deviation from each other). These differences do not exhibit a strong level of statistical significance, but their direction suggests that ignoring out-of-stocks leads to lower estimates of demand for these brands. In contrast, the corresponding posterior mean for Herbal Essence is very similar across both models, which is consistent with the high availability of this product on the shelf (92%, see fraction of out-of-stocks for product 17 in Table 3).

5 Estimating and Mitigating the Costs of Out-Of-Stocks

Using the expressions derived in Subsection 3.5 (see equation 11), we estimated lost sales for every time period in each store. Figure 1 shows the estimated lost sales as a fraction of the total sales that would have been observed if all products had been available (referred to as *full-availability sales*).¹⁷ The estimates suggest that the percentage of lost sales is large and varies considerably within and across stores. In some instances, the lost sales reach levels close to 100%, as in days 3 and 9 in store 5. Both days correspond to Saturdays with a substantial fraction of out-of-stock products (10 out of the 24 SKUs were out-of-stock).

To better understand the effect of product availability on lost sales, Figure 2 shows the estimated average lost sales at store-days with different number of products out-of-stock. The figure reveals that lost sales increases non-linearly in the number of products not available. When five or fewer products are out-of-stock the average lost sales are equal to 9.5% of the full availability sales. But when six or more products are out-of-stock, lost sales grows by more than 3 times, up to 30.1% of the full-availability sales. This suggests that the first products that stock-out have a smaller impact

¹⁷The full-availability scenario assumes that the retailer has enough shelf space to satisfy the full-availability demand. Accordingly, full-availability sales are defined as the sum between the observed and lost sales.




Figure 1: Percentage of lost purchases by day and store as a fraction of full-availability sales.




Figure 2: Percentage of lost purchases as a fraction of full-availability sales versus number of out-of-stock products.

on category sales, but lost sales increase more rapidly as more products become unavailable. This pattern has important implications for assortment planning and inventory replenishment decisions. In supermarkets, there are substantial fixed costs of replenishing products on the shelf. Hence, Figure 2 suggests that replenishments could be postponed until several products become out-of-stock without substantially affecting category sales, although the profitability consequences of such a policy will depend on the retail margins of each of the sold-out products.

In addition to estimating lost sales, we also evaluate the effect of a strategy that seeks to mitigate the effect of out-of-stocks: conducting a temporary price reduction for a single product to recapture a fraction of lost sales. Price discounts increase the attractiveness of substitute products, inducing some customers whose preferred product is not available to substitute their intended purchase with another product rather than choosing the no-purchase option. In what follows, we show how to use our model to quantify the expected fraction of lost sales that would be recaptured through this temporary price promotions. Consider the set O_A of customer who chose not to purchase when assortment A was available (previously defined in section 3.5). Suppose that the price of product

$k \in A$ is discounted by some fraction of the original price. For each customer $i \in O_A$, we define two events: (1) E_{i1} is the event that customer i would have purchased some of the unavailable products (i.e. products in A^c) had they been available; (2) E_{i2} is the event that the customer would purchase the discounted product k when only products in A are available. All customers in O_A satisfying E_{i1} are counted as lost sales; those who in addition satisfy E_{i2} count as recaptured lost sales. Therefore, the fraction of lost sales that is recaptured can be calculated as:

$$\% \text{ Reduction in Lost Sales} = \frac{\sum_{i \in O_A} \Pr(E_{i1}, E_{i2} | i \in O_A)}{LS},$$

where LS is defined in equation (11) and details on how to calculate the numerator of the expression above using draws from the MCMC simulation are shown in the Online Appendix B.

Accordingly, the model was used to measure the effectiveness of implementing price reductions of 20% for different store-days. We consider scenarios where only one product is discounted at a time and then compare the impact among products. As an illustration, we report results for two store-days with very different levels of availability: (i) day 3 in market 5, where 10 SKUs are out-of-stock; and (ii) day 15 in market 2, where only one product is out-of-stock—SKU 15 (Pantene). In the case of day 3 in market 5, a price promotion on product 17 (Herbal Essence) reduces lost sales by 2.5%; all other promotions reduce lost sales by less than 1%. However, in the case of day 15 in market 2, it is more effective to discount product 13 (Pantene), which leads to a lost-sales reduction of 4.5%. In this case, a promotion of product 17 leads to a 3.2% reduction in lost sales.

The comparison of the results from these two store-days provides insights on the effectiveness of price promotions at mitigating the consequences of out-of-stocks. In the case of day 15 in market 2 where only one product was out-of-stock, a price promotion on a product with similar characteristics to the missing product is more effective to recapture lost sales. In contrast, when many products are missing, it is more effective to use a price promotion on a popular product (SKU 17).

These results show that price promotions can be useful to reduce lost sales and improve customer service. However, promotions may also have negative consequences, including: (1) a reduction in the margin of the discounted product; (2) cannibalization of sales from other products with higher margins. These consequences can also be easily estimated with our model if data on the cost of goods sold by SKU is available. Using gross margin data from the supermarket chain under study,

we evaluated the net change in category profits accounting for the effect of increased sales, reduced markups and cannibalization. In the case of day 3 in market 5, the promotion on product 17 increased total profits by 14.7%. In contrast, in the case of day 15 in market 2, the price promotion on products 13 and 17 *reduces* profits by 33% and 8%, respectively.¹⁸

In summary, this section illustrates how the methodology can be used to assess the consequences of policies aimed at mitigating the costs of out-of-stocks (e.g., lost sales reduction, impact on category profits). Many additional policy experiments could be also performed, including discounts on multiple products and experimenting with different magnitudes for price discounts. These additional policy experiments can be easily implemented using the estimated model.

6 Conclusions

In this paper, we have proposed a method to capture the effects of out-of-stocks on customer behavior using data available to a store manager. Lack of precise availability data in real-time, which is common in practical retail settings, introduces a major challenge in identifying the effect of out-of-stocks on sales. To overcome these difficulties, our methodology simulates the transition of the inventory on the shelf conditioning on snapshots of information about product inventory, which could be obtained through a periodic inventory review system. These data are combined with daily sales and pricing data to estimate a structural model of demand for a product category.

The method has several attractive features. First, the model can be estimated using data from multiple stores and markets. Second, our method is applicable in categories with many products and with slow moving products exhibiting frequent out-of-stocks. Previous methods in this area were limited by the number of products, the type of substitutions and the number of products that could be simultaneously out-of-stock. Third, our structural model allows the evaluation of policies to mitigate the consequences of out-of-stocks (such as temporary price promotions). Fourth, the methodology explicitly models the distribution of the timing of out-of-stocks (which are unobservable in periodic inventory systems), providing useful information on product availability.

In terms of future research, several interesting generalizations can be identified. In terms of the demand model, alternative specifications or behavioral assumptions could be considered (e.g.,

¹⁸Promotions on other products, such as SKU 10, reduced lost sales by about 1% without substantially changing profits.

probit model, complexity of choice decisions). In addition, the model could be extended to explicitly consider the possibility of some customers returning to the store in future periods if they did not find a product in a given time period. Finally, since our methodology is based on estimating the joint distribution of availability and sales, a possible extension could be developed for the detection of out-of-stocks based on point-of-sales data.

In summary, we hope that the methodology presented in this paper may provide useful tools for researchers and managers to estimate the consequences of out-of-stocks and assess the impact of policies to mitigate their costs. This issue is certainly relevant for audiences in the operations management, marketing and economics disciplines. We hope this paper may help to stimulate further research aimed at addressing problems that cross the boundaries between these fields.

References

- Allenby, G. M., P. E. Rossi. 2003. Bayesian statistics and marketing. *Marketing Science* **22** 304–328.
- Anderson, E., G. Fitzsimons, D. Simester. 2006. Measuring and mitigating the cost of stockouts. *Management Science* **52**(11) 1751–1763.
- Andrews, R. L., T. C. Srinivasan. 1995. Studying consideration effects in empirical choice models using scanner panel data. *Journal of Marketing Research* **32** 30–41.
- Anupindi, R., M. Dada, S. Gupta. 1998. Estimation of consumer demand with stock-out based substitution: An application to vending machine products. *Marketing Science* **17**(4) 406–423.
- Ben-Akiva, M. E., S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Berry, S. 1994. Estimating discrete choice models of product differentiation. *Rand Journal of Economics* **25**(2) 242–262.
- Berry, S., J. Levinson, A. Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **963**(4) 841–890.
- Besanko, D., J.-P. Dubé, S. Gupta. 2003. Competitive price discrimination strategies in a vertical channel using aggregate data. *Management Science* **49** 1121–1138.

- Bruno, H., N. Vilcassim. 2008. Structural demand estimation with varying product availability. *Marketing Science* forthcoming.
- Campo, K., E. Gijsbrechts, P. Nisol. 2003a. The impact of retailer stockouts on whether, how much, and what to buy. *International Journal of Research in Marketing* **20**(3) 273–286.
- Campo, K., E. Gijsbrechts, P. Nison. 2003b. Towards understanding consumer response to stockouts. *Journal of Retailing* **72**(2).
- Chen, Y., S. Yang. 2007. Estimating disaggregate models using aggregate data via augmentation of individual choice. *Journal of Marketing Research* **4** 596–613.
- Chintagunta, P. K., J.-P. Dubé. 2005. Estimating a stockkeeping-unit-level brand choice model that combines household panel data and store data 368-79. *Journal of Marketing Research* **42** 368–379.
- Chintagunta, P. K., D. Jain, N. J. Vilcassim. 1991. Investigating heterogeneity in brand preferences in logit models for panel data. *Journal of Marketing Research* **28** 417–428.
- Conlon, C., J. Holland Mortimer. 2007. Demand estimation under incomplete product availability. Working paper, Harvard University.
- DeHoratius, N., A. Raman. 2008. Inventory record inaccuracy: an empirical analysis. *Management Science* **54**(4) 627–641.
- Dolan, R., T. Hayes. 2005. The myths and dangers of acv. Tech. rep., The New England Consulting Group.
- Fader, P., B. Hardie. 1996. Modeling consumer choice among skus. *Journal of Marketing Research* **33** 442–452.
- Fitzsimons, G. J. 2000. Consumer response to stockouts. *Journal of Consumer Research* **27** 249–266.
- Gruen, T.W., D.S. Corsten, S. Bharadwaj. 2002. *Retail Out-of-stocks: A Worldwide Examination of Extent, Causes and Consumer Responses*. Grocery Manufacturers of America.
- Hauser, J., B. Wernerfelt. 1990. An evaluation cost model of consideration sets. *Journal of Consumer Research* **16** 393–408.

- Jiang, R., P. Manchanda, P. Rossi. 2007. Bayesian analysis of random coefficient logit models using aggregate data. Working paper, University of Chicago.
- Kalyanam, K., S. Borle, P. Boatwright. 2007. Deconstructing each item's category contribution. *Marketing Science* **26** 327–341.
- Kass, R. E., A. E. Raftery. 1995. Bayes factors. *Journal of the American Statistical Association* **Vol. 90** 773–795.
- Kök, A. G., M. L. Fisher. 2007. Demand estimation and assortment optimization under substitution: Methodology and application. *Operations Research* **55**(6) 1001.
- Musalem, A., E. T. Bradlow, J. S. Raju. 2007. Bayesian estimation of random-coefficients choice models using aggregate data. *forthcoming, Journal of Applied Econometrics* .
- Musalem, A., E. T. Bradlow, J. S. Raju. 2008. Who's got the coupon? estimating consumer preferences and coupon usage from aggregate information. *Journal of Marketing Research* forthcoming.
- Roberts, J. H., J. M. Lattin. 1991. Development and testing of a model of consideration set composition. *Journal of Marketing Research* **28** 429–440.
- Rossi, P. E., R. McCulloch, G. M. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Science* **15** 321–340.
- Smith, S. A., N. Agrawal. 2000. Management of Multi-Item Retail Inventory Systems with Demand Substitution. *Operations Research* **48**(1) 50–64.
- Train, K. E. 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Villas-Boas, J. M., R. S. Winer. 1999. Endogeneity in brand choice models. *Management Science* **45** 1324–1338.
- Vulcano, G., G. van Ryzin, R. Ratliff. 2008. Estimating primary demand for substitutable products from sales transaction data. Working paper.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill.

Online Appendix A: Simulation Study

In this Appendix we describe the results of a simulation study designed to show the statistical properties of the method proposed in this paper. Using the same parameter values as in Subsection 3.4, we generated 50 data sets and estimated the model parameters using the proposed method. We also repeated this process estimating the model parameters ignoring the occurrence of out-of-stocks. Table A1 shows the results of this simulation study. In each block, we display the mean, standard deviation, median, mean squared error (mse) and bias of the posterior mean of each parameter under each estimation approach. We also report the fraction of replications for which each posterior mean is below its true value (p_{below}). Ideally, we would expect this percentage to take values close to 0.5.

From these results we observe an important bias of the posterior means under the model that ignores out-of-stocks. In fact, for 7 out of 10 parameters, the posterior means are *almost always* below or above the truth. For example, in the case of θ_{z_1, x_1} which is used to determine the mean of the coefficients for brand 1, the posterior mean is below the true value for all replications ($p_{\text{below}} = 1.0$ for θ_{z_1, x_1}). In contrast, under the full model, we see a much more even distribution of the posterior mean around the truth, with values of p_{below} much closer to 0.5 ($p_{\text{below}} = 0.56$ for θ_{z_1, x_1}). In addition, we observe an improvement in the bias and mean squared error for most parameters under the full model when compared with the results of the model that ignores out-of-stocks.

Online Appendix B: Calculating the reduction in lost sales due to a price promotion

This section shows the details on how to calculate the reduction of lost sales after applying a discount on a single product. Define:

- A : the set of products available.
- A^c : the set of products out of stock.
- j' : index of the discounted product.
- Δ : the change (increase) in the utility of product j' after applying the discount.

The set of customer that choose not to purchase under assortment A is $O_A = \{i/u_{i0} > u_{ij}, \forall j \in A\}$.

The events E_{i1} and E_{i2} are defined as:

$$E_{i1} = \{u_{i0} < \max_{k \in A^c} u_{ik}\}$$

$$E_{i2} = \{u_{i0} < u_{ij'} + \Delta\}.$$

We also define the following subsets of consumers:

$$O_{-j'} = \{i : u_{i0} > u_{ij}, \forall j \in A \setminus j'\},$$

$$O_{j'} = \{i : u_{i0} > u_{ij'}\}.$$

Note that $O_A = O_{-j'} \cap O_{j'}$. Our objective is to compute $\Pr(E_{i1}, E_{i2} | i \in O_A)$, the reduction in lost sales when discounting product j' . To facilitate notation, we denote the event $i \in O_A$ more succinctly as O_A (similarly with $O_{-j'}$ and $O_{j'}$) and suppress all i indexes. Note that:

$$\Pr(E_1, E_2, O_A) = \Pr(E_2, O_A) - \Pr(E_1^c, E_2, O_A). \quad (12)$$

In our derivation, we will use two standard results of multinomial logit choice models (see Ben-Akiva and Lerman (1985)):

1 The distribution of $X = \max_{k \in P} u_k$ is Gumbel with mode $\log \sum_{k \in P} \exp(V_k)$.

2 Given two independent X and Y Gumbel random variables with modes V_x and V_y (respectively), the difference $X - Y$ is logistically distributed: $F_{X-Y}(z) = \frac{1}{1 + \exp(V_y - V_x - z)}$.

We first calculate $P(O_A, E_2)$:

$$\begin{aligned} \Pr(O_A, E_2) &= \Pr(O_{j'}, E_2, O_{-j'}) \\ &= \Pr(O_{-j'}) \cdot \Pr(O_{j'}, E_2 | O_{-j'}) \\ &= \Pr(O_{-j'}) \cdot \Pr(0 < u_0 - u_{j'} < \Delta | u_0 > u_k, \forall k \in A \setminus j'). \end{aligned}$$

The distribution of u_0 conditional on $u_0 > u_k, \forall k \in A \setminus j'$ is a Gumbel distribution with mode $\log \sum_{k \in A \setminus j'} \exp(V_k)$. Let $Y_{-j'}$ denote a random variable with this distribution. Moreover, $Y_{-j'}$ and $u_{j'}$ are independent. Using results 1 and 2 above, $Y_{-j'} - u_{j'}$ has a logistic distribution and we obtain:

$$\begin{aligned} \Pr(O_{j'}, E_2 | O_{-j'}) &= \Pr(0 < Y_{-j'} - u_{j'} < \Delta) \\ &= \frac{\exp(V_{j'} + \Delta)}{1 + \exp(V_{j'} + \Delta) + \sum_{k \in A \setminus j'} \exp(V_k)} - \frac{\sum_{k \in A \setminus j'} \exp(V_k)}{1 + \sum_{k \in A} \exp(V_k)}. \end{aligned}$$

In addition, $\Pr(O_{-j'}) = \left(1 + \sum_{k \in A \setminus j'} \exp(V_k)\right)^{-1}$, which completes the calculation of $P(O_A, E_2)$.

We now turn to the calculation of $\Pr(O_A, E_2, E_1^c)$. We have:

$$\begin{aligned} \Pr(O_A, E_2, E_1^c) &= \Pr(O_{j'}, O_{-j'}, E_1^c, E_2) \\ &= \Pr(O_{j'}, E_2 | O_{-j'}, E_1^c) \cdot \Pr(O_{-j'}, E_1^c) \\ &= \Pr(0 < u_0 - u_{j'} < \Delta | u_0 > u_k, \forall k \in \mathcal{J} \setminus j') \cdot \frac{1}{1 + \sum_{k \in \mathcal{J} \setminus j'} \exp(V_k)}. \end{aligned}$$

The distribution of u_0 conditional on $u_0 > u_k, \forall k \in \mathcal{J} \setminus j'$ is a Gumbel distribution with mode $\log \sum_{k \in \mathcal{J} \setminus j'} \exp(V_k)$. Let Z denote a random variable with this distribution. The random variable $Z - u_{j'}$ has a logistic distribution and we obtain:

$$\begin{aligned}
\Pr(0 < u_0 - u_{j'} < \Delta | u_0 > u_k, \forall k \in \mathcal{J} \setminus j') &= \Pr(0 < Z - u_{j'} < \Delta) \\
&= \frac{\exp(V_{j'} + \Delta)}{1 + \exp(V_{j'} + \Delta) + \sum_{k \in \mathcal{J} \setminus j'} \exp(V_k)} \\
&= \frac{\sum_{k \in \mathcal{J} \setminus j'} \exp(V_k)}{1 + \sum_{k \in \mathcal{J} \setminus j'} \exp(V_k)}
\end{aligned}$$

This completes the calculation of $\Pr(O, E_2, E_1^c)$. Replacing $\Pr(O, E_2, E_1^c)$ and $P(O, E_2)$ in (12) enabling us to estimate the reduction in lost sales.

Table A1

Results: Summary statistics for the posterior mean of each parameter across 50 replications ignoring out-of-stocks (first block) and accounting for out-of-stocks (second block).

Ignoring OOS	θ_{z_1,x_1}	θ_{z_1,x_2}	θ_{z_1,x_3}	θ_{z_1,x_4}	θ_{z_2,x_1}	θ_{z_2,x_2}	θ_{z_2,x_3}	θ_{z_2,x_4}	Σ_{x_3,x_3}	Σ_{x_4,x_4}	σ_ξ^2
mean	1.172	1.009	-4.469	-0.969	0.015	-0.430	0.245	0.332	0.665	0.462	0.476
sd	0.172	0.118	0.322	0.296	0.242	0.159	0.177	0.091	0.322	0.250	0.076
median	1.170	1.002	-4.493	-0.906	-0.033	-0.432	0.246	0.335	0.539	0.406	0.469
mse	0.716	0.255	2.261	2.432	0.294	0.030	0.091	0.144	0.122	2.427	0.006
bias	-0.828	-0.491	-1.469	1.531	-0.485	0.070	0.245	-0.368	-0.135	-1.538	-0.024
p_{below}	1.000	1.000	1.000	0.000	0.980	0.380	0.040	1.000	0.800	1.000	0.600
true	2.000	1.500	-3.000	-2.500	0.500	-0.500	0.000	0.700	0.800	2.000	0.500

Full Model	θ_{z_1,x_1}	θ_{z_1,x_2}	θ_{z_1,x_3}	θ_{z_1,x_4}	θ_{z_2,x_1}	θ_{z_2,x_2}	θ_{z_2,x_3}	θ_{z_2,x_4}	Σ_{x_3,x_3}	Σ_{x_4,x_4}	σ_ξ^2
mean	1.994	1.501	-3.148	-2.429	0.496	-0.484	0.070	0.666	1.107	1.949	0.505
sd	0.106	0.079	0.321	0.294	0.146	0.122	0.139	0.092	0.781	0.413	0.092
median	1.983	1.502	-3.072	-2.443	0.511	-0.504	0.048	0.670	0.828	1.891	0.514
mse	0.011	0.006	0.125	0.092	0.021	0.015	0.024	0.010	0.704	0.173	0.009
bias	-0.006	0.001	-0.148	0.071	-0.004	0.016	0.070	-0.034	0.307	-0.051	0.005
p_{below}	0.560	0.480	0.660	0.420	0.480	0.520	0.280	0.660	0.480	0.560	0.440
true	2.000	1.500	-3.000	-2.500	0.500	-0.500	0.000	0.700	0.800	2.000	0.500