



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Technical Note—An Expectation-Maximization Method to Estimate a Rank-Based Choice Model of Demand

Garrett van Ryzin, Gustavo Vulcano

To cite this article:

Garrett van Ryzin, Gustavo Vulcano (2017) Technical Note—An Expectation-Maximization Method to Estimate a Rank-Based Choice Model of Demand. *Operations Research* 65(2):396-407. <http://dx.doi.org/10.1287/opre.2016.1559>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Technical Note—An Expectation-Maximization Method to Estimate a Rank-Based Choice Model of Demand

Garrett van Ryzin,^a Gustavo Vulcano^{b,c}

^a Graduate School of Business, Columbia University, New York, New York 10027; ^b Leonard N. Stern School of Business, New York University, New York, New York 10012; ^c School of Business, Universidad Torcuato di Tella, Buenos Aires, Argentina

Contact: gjv1@columbia.edu (GvR); gvulcano@stern.nyu.edu (GV)

Received: April 3, 2015

Revised: January 27, 2016; June 20, 2016

Accepted: August 5, 2016

Published Online in Articles in Advance: February 24, 2017

Subject Classifications: marketing: choice models; statistic: estimation, nonparametric

Area of Review: Operations Management

<https://doi.org/10.1287/opre.2016.1559>

Copyright: © 2017 INFORMS

Abstract. We propose an expectation-maximization (EM) method to estimate customer preferences for a category of products using only sales transaction and product availability data. The demand model combines a general, rank-based discrete choice model of preferences with a Bernoulli process of customer arrivals over time. The discrete choice model is defined by a probability mass function (pmf) on a given set of preference rankings of alternatives, including the no-purchase alternative. Each customer is represented by a preference list, and when faced with a given choice set is assumed to either purchase the available option that ranks highest in her preference list, or not purchase at all if no available product ranks higher than the no-purchase alternative.

We apply the EM method to jointly estimate the arrival rate of customers and the pmf of the rank-based choice model, and show that it leads to a remarkably simple and highly efficient estimation procedure. All limit points of the procedure are provably stationary points of the associated incomplete data log-likelihood function, and the output produced are maximum likelihood estimates (MLEs). Our numerical experiments confirm the practical potential of the proposal.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2016.1559>.

Keywords: demand estimation • demand untruncation • choice behavior • EM method

1. Introduction

Demand estimation is a fundamental task in retail operations and revenue management, providing the necessary input data for inventory control and price optimization models. The task is particularly difficult when product availability varies over time and customers may substitute.

In this paper, we propose a remarkably simple estimation algorithm for a nonparametric model of choice-based demand that has been introduced in the operations management related literature by Mahajan and van Ryzin (2001b, a). Therein, consumers are characterized by a rank ordering of all alternatives of a given product category (along with the no-purchase option). When faced with a choice, a consumer is assumed to either purchase the available product that ranks highest in her preference list, or not purchase at all if the no-purchase alternative ranks higher than any available product. In this way, the market can be segmented into a finite number of customer types defined by their corresponding rank orders. The preferences of a random customer drawn from a given population are then described by a discrete probability mass function (pmf) on the set of customer types. Several common demand processes studied in the marketing, economics, and operations literature can be modeled as special cases of this rank-based choice

model (e.g., independent demand, multinomial logit (MNL), nested logit, Markovian second choice, universal backup, and Lancaster demand). To complete the demand model, we assume customers arrive over time according to a Bernoulli process with constant rate, which must be estimated along with the customer type pmf.

Rank-based choice models of demand have received increasing attention in the retail operations and (airline) revenue management literature, focusing largely on optimization rather than estimation problems. Examples of the former (in addition to the two aforementioned papers) are Smith et al. (2009); Honhon et al. (2010, 2015); Jagabathula and Rusmevichientong (2014); Jagabathula (2014), and Bertsimas and Mišić (2015). In the context of airline revenue management, Zhang and Cooper (2006), van Ryzin and Vulcano (2008), Chen and de Mello (2010), Chaneton and Vulcano (2011), and Kunnumkal (2014) also use this rank-based choice model. With the increasing trend towards nonparametric, data-driven approaches, we believe rank-based models of choice will continue their expansion in the academic literature and industry practice.¹

However, with few exceptions, relatively little work has been done on estimation methods for rank-based choice models. The main challenge on the estimation

side is derived from a major drawback of this choice model: The potential number of customer types is factorial in the number of alternatives. Farias et al. (2013) propose a constraint sampling procedure to implement a robust estimation approach, finding the distribution over customer types that produces the worst-case revenue for a given fixed assortment, compatible with the observed data. Then they show that the demand model constructed by the robust approach is approximately the *sparsest-choice model*, where sparsity is measured by the number of customer types that occur with positive probability in the population. By contrast, we find maximum likelihood estimates (MLEs) of such proportions in the confines of a pre-specified set of customer types. The only required inputs (in addition to the set of customer types) are observed historical sales and product availability data. While the maximum likelihood (ML) estimation problem can be solved via standard nonlinear optimization techniques, doing so is computationally intensive, especially when a large number of customer types is considered. The objective of this technical note is to develop a problem-specific and greatly simplified alternative to general nonlinear optimization techniques.

Our paper van Ryzin and Vulcano (2015) addresses the dimensionality challenge using an iterative, column-generation-based procedure to enrich an initial, parsimonious set of customer types (e.g., preference lists with one element followed by the no-purchase option, describing the independent demand model). Yet even in each iteration of that procedure, for a limited set of customer types, one must still solve the problem of efficiently estimating the pmf and arrival rate via ML estimation.

Using a different alternative to the uninformed market structure assumed by Farias et al. (2013) and our previous paper van Ryzin and Vulcano (2015), one can attempt to limit the set of customer types based on a preliminary market description. For instance, Yunes et al. (2007, Section 4) describe how they use conjoint survey results at John Deere & Co. to identify consumer types and their ranking lists based on the part-worth utilities calculated from the survey for two of its product lines. In the context of assortment planning, Honhon et al. (2012) show that when all the products can be mapped onto a hierarchical ordering system such as a branched (outtree), vertical (one-way substitution), and horizontal (locational) order, the assortment optimization methods run in polynomial time and consist of simple algorithms based on dynamic programming or the solution to a shortest path problem on a properly defined directed graph. Nonetheless, the number of types for the three hierarchical ordering systems is $O(n^2)$, where n is the number of products in the category, which leads to a number of types in the

order of hundreds, confirming the need for an efficient estimation procedure.²

We explore two alternative formulations of the rank-based estimation problem. In the first, we assume that the modeler can distinguish a period with no arrival from a period where an arrival ended in a no-purchase, e.g., as in online retailing where customer visits, purchases, and no-purchases are tracked. The unobservable data are only the type of an arrival in a given period. The second variant is aimed at brick-and-mortar retailers, where customer visits are not tracked and therefore arriving customers who did not purchase are not recorded. (The same problem arises in online retailers that record transactions rather than visits.) In this case, our procedure corrects for both sources of incompleteness in the data: (a) the arrival (or not) of a customer, and (b) if an arrival occurs, the type of the customer.

One widely used method to correct for realized transactions affected by stockouts and substitution effects is the expectation-maximization (EM) algorithm (e.g., Anupindi et al. 1998, K ok and Fisher 2007, Conlon and Mortimer 2013, Stefanescu 2009). The EM method was proposed by Dempster et al. (1977). It is a generic procedure that works using alternating steps: (1) computing the conditional expected value of the complete-data log-likelihood based on the current parameter estimates and the observable data (the E-step), and (2) maximizing this function to obtain improved parameter estimates (the M-step). Its practical effectiveness depends on the computational efficiency of both steps. Our contribution is to specialize the EM method to our two formulations and show that it is trivial to implement in any simple procedural language or numerical computing environment (e.g., Matlab) via iterations involving closed-form expressions. Moreover, even though it is well acknowledged that the convergence of the EM algorithm could be slow, particularly with large numbers of parameters or high degrees of censoring, by running an exhaustive set of numerical experiments, we show that in our case it is between twice and six times faster than direct MLE methods, and that the quality of the estimates is essentially equivalent.

Haensel and Koole (2011) also present an EM-based procedure to estimate a nonparametric choice model that is quite similar to ours, consisting also of a set of customer types that are defined by preference orderings. However, they do not explicitly take into account the incompleteness of the data with respect to arrival and no-arrival periods and instead use a heuristic extrapolation to account for the lack of observations in a given time interval. Also different from our work, their model allows for multiple sales per period, which leads to a discrete Poisson model of demand

that requires an approximate rounding procedure in the E-step.

In terms of theoretical convergence, our EM procedure benefits from one of the strongest results derived for EM variants: All limit points of the algorithm are provably stationary points of the associated incomplete data log-likelihood function. Since the demand model under consideration is non identifiable in general (e.g., see van Ryzin and Vulcano 2015, Section 3.3 for a discussion), alternative solutions could lead to the same log-likelihood value. However, in cases where the underlying demand model were identifiable in the confines of a given set of customer types, given that our estimates are MLEs, they inherit the desirable statistical properties of maximum likelihood estimators, i.e., they are consistent, asymptotically unbiased, and asymptotically efficient (i.e., their variance tends, asymptotically, to the Cramer-Rao lower bound).

The remainder of this note is organized as follows. In Section 2 we introduce the demand model and the maximum likelihood estimation problem. Section 3 presents our EM algorithm for the observed arrivals case, and Section 4 extends the approach to the censored demand case. Section 5 shows our numerical experiments, and we present our conclusions in Section 6.

2. Model

A sequence of offer sets $(S_1, \dots, S_t, \dots, S_T)$ is available over T purchase periods. Each offer set is a subset of the products in $\mathcal{N} = \{0, 1, \dots, n\}$, where 0 stands for the outside or no-purchase option so that $0 \in S_t$, for all t . We assume customers arrive according to a discrete-time, homogeneous Bernoulli arrival process, with rate $\lambda < 1$. There is at most one arrival per period, and the parameter λ must be estimated from data. If $\lambda \ll 1$, this arrival process can be considered a discrete-time approximation to a Poisson arrival process. In cases where the arrival rate may not be homogeneous throughout the selling horizon, one can partition time periods into multiple windows in such a way that in each time window the arrival rate is constant. For ease of exposition, however, we will consider only the case where the arrival rate is constant over time periods.

Customers are assumed to have a rank-based preference for products. That is, each customer has a preference list (or total ranking), σ , of the products in \mathcal{N} . Each preference list σ defines a *customer type*. A customer of type σ prefers product h to product j if and only if $\sigma(h) < \sigma(j)$. The market, in turn, is assumed to be composed of a fixed, pre-defined set of customer types $\sigma = \{\sigma^{(1)}, \dots, \sigma^{(N)}\}$. This set could be defined based on judgment and knowledge of the market or on surveys.³ Nevertheless, since the set of types used in the estimation model may differ from the *true* set of types present in the market, the set σ is a potential source

of specification error in our model. In van Ryzin and Vulcano (2015) we develop a procedure to discover new customer types that improve the likelihood function, which makes it possible to sequentially augment an initial candidate set of types σ . Yet, in each iteration of that procedure, we still have the problem of computing MLEs for the candidate set, which can be addressed via our EM.

Note that the meaningful rankings are only those truncated at the position of product 0. That is, a preference list σ will have as many elements as $\sigma(0)$, where we will assume that $\sigma(0) \geq 2$, since a customer type whose first preference is not to purchase anything is irrelevant. Arriving customers are assumed to be of type $\sigma^{(i)}$ with probability $x_i = \mathbb{P}(\sigma^{(i)})$, $i = 1, \dots, N$. This pmf for customer types is denoted $\mathbf{x} = \mathbb{P}(\sigma)$, and it must be estimated from data. Customer types are assumed to be independent across periods.

Upon arrival, each customer chooses from a set of products S_t , with $|S_t| \geq 2$, so that at least one product other than the no-purchase option is available in each period.⁴ We assume S_t is observable. An arriving customer chooses her most preferred product among those available in S_t . In symbols, if there is an arrival of type σ in period t , that customer chooses $j_t = \arg \min_{j \in S_t} \sigma(j)$.

Below we consider two variants of this problem. For both variants, the data needed is minimal, i.e., product availability information (not even inventory levels, just binary indicators for availability), and sales transaction data for each period. The difference between the variants rests on the observability of the no-arrivals and of the no-purchase outcome of the arrivals. In the uncensored demand setting (Section 3), we assume that the arrivals and no-arrivals are observed, and that the no-purchase outcome of an arrival is also observed. This corresponds, for example, to the case of online retailers who record shopping data (i.e., customer visits to the website that may end up in a transaction or a no-purchase). Next, in Section 4, we consider the effect of censoring that occurs when no-purchase outcomes are unobservable, which is typical in brick-and-mortar retail settings. The extra difficulty here is that the modeler cannot distinguish a period with no arrival from a period with an arrival that did not end in a transaction.

For each variant, our objective is to compute MLEs for the parameters $\theta = (\mathbf{x}, \lambda)$, where $\mathbf{x} = \mathbb{P}(\sigma)$ for a given set of rankings σ . Since customers do not explicitly reveal their types upon arrival, all we can determine is that for each period with an observed purchase, the customer's choice is consistent with some subset of the possible types (i.e., those types who rank the purchased product higher than all other available products). In addition, in the censored demand case, for each period with no observed purchase, wherein we can only claim that there was no arrival or if there was an

arrival, the customer preferred the no-purchase alternative to any of the products offered. Hence, sales data provide only incomplete observations of the choice model we wish to estimate.

The incompleteness in the data with respect to the nonobservability of customer types is captured by the following definition: For $j_t \in S_t$ consider the *compatible-type set*:

$$\mathcal{M}_t(j_t, S_t) = \{i: \sigma^{(i)}(j_t) < \sigma^{(i)}(k), \forall k \in S_t, k \neq j_t\}.$$

In words, $\mathcal{M}_t(j_t, S_t)$ is the set of customer type indexes for which the selected product j_t is ranked highest among the available products in S_t . The definition also covers the case $j_t = 0$, and the corresponding set $\mathcal{M}_t(0, S_t)$. That is, $\mathcal{M}_t(j_t, S_t)$ represents the customer types that are consistent with the observed transaction j_t and the offer set S_t . We further assume that the full set of customer types σ is consistent with the observed transactions and availability data, in the sense that all transactions and no-transactions observed could be explained by σ . To guarantee this, it is enough to include customer types of the form $(j, 0)$, spanning all products $j \in \mathcal{N}$.

Given a probability distribution \mathbf{x} over the set of types, the probability that a random arrival chooses product j_t given offer set S_t is given by

$$\begin{aligned} \mathbb{P}(j_t | S_t) &= \sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i, \quad \text{if } j_t \in S_t, \quad \text{and} \\ \mathbb{P}(j_t | S_t) &= 0, \quad \text{if } j_t \notin S_t. \end{aligned} \quad (1)$$

3. Estimation: Uncensored Demand Case

We start by assuming that the modeler can observe all arrivals and their purchase incidence: Either the arriving customer purchased an available product or chose not to purchase. The only source of incompleteness is the type of the arrival.

3.1. Problem Formulation

Let \mathcal{P} be the set of periods with purchases. Let $\bar{\mathcal{P}}_\lambda$ be the set of periods with arrivals that trigger no purchases, and let $\tilde{\mathcal{P}}_\lambda$ be the set of periods with no arrivals, with $T = |\mathcal{P}| + |\tilde{\mathcal{P}}_\lambda| + |\bar{\mathcal{P}}_\lambda|$. Given a fixed set of types $\sigma = \{\sigma^{(1)}, \dots, \sigma^{(N)}\}$, the incomplete data log-likelihood function is formulated as:

$$\begin{aligned} \mathcal{L}_I(\mathbf{x}, \lambda) &= \sum_{t \in \mathcal{P}} (\log \lambda + \log \mathbb{P}(j_t | S_t)) \\ &\quad + \sum_{t \in \bar{\mathcal{P}}_\lambda} (\log \lambda + \log \mathbb{P}(0 | S_t)) + \sum_{t \in \tilde{\mathcal{P}}_\lambda} \log(1 - \lambda) \\ &= \sum_{t \in \mathcal{P}} \left(\log \lambda + \log \left(\sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i \right) \right) \\ &\quad + \sum_{t \in \bar{\mathcal{P}}_\lambda} \left(\log \lambda + \log \left(\sum_{i \in \mathcal{M}_t(0, S_t)} x_i \right) \right) \\ &\quad + \sum_{t \in \tilde{\mathcal{P}}_\lambda} \log(1 - \lambda). \end{aligned}$$

The first term accounts for the likelihood of the observed transactions in the periods with purchases; the second term accounts for the no-purchase periods, where an arriving customer preferred not to buy; and the third term accounts for the periods with no arrivals. A more compact representation of the log-likelihood function follows:

$$\begin{aligned} \mathcal{L}_I(\mathbf{x}, \lambda) &= \sum_{t \in \mathcal{P}} \log \left(\sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i \right) + \sum_{t \in \bar{\mathcal{P}}_\lambda} \log \left(\sum_{i \in \mathcal{M}_t(0, S_t)} x_i \right) \\ &\quad + (|\mathcal{P}| + |\bar{\mathcal{P}}_\lambda|) \log \lambda + |\tilde{\mathcal{P}}_\lambda| \log(1 - \lambda). \end{aligned} \quad (2)$$

The function is separable in \mathbf{x} and λ , and globally concave in (\mathbf{x}, λ) . The maximizer λ^* has a closed form given by $\lambda^* = (|\mathcal{P}| + |\bar{\mathcal{P}}_\lambda|)/T$. To simplify notation, we define $\hat{T} = |\mathcal{P}| + |\bar{\mathcal{P}}_\lambda|$, the number of periods where arrivals occurred.

Next we assume that the arrival rate λ^* has already been established. The ML estimation problem can be formulated as follows:

$$\max_{\mathbf{x} > 0} \mathcal{L}_I(\mathbf{x}) \quad \text{s.t.} \quad \sum_{i=1}^N x_i = 1. \quad (3)$$

Formulation (3) is a concave, constrained optimization problem, defined over the open set $\mathbf{x} \geq 0$, with at least one global optimum x_1^*, \dots, x_N^* . One way to solve it is by using standard nonlinear optimization methods. However, the challenge is the potentially high dimensionality of the model. To simplify its solution, we use the EM method of Dempster et al. (1977) as described below.

3.2. Theoretical Properties of MLEs

Before proceeding with the estimation algorithm, we derive a few properties of the MLEs of the model parameters. We introduce new variables \mathbf{y} defined as $y_t = \sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i$, representing the aggregate likelihood of the customer types that would pick alternative j_t in period t , i.e., $y_t = \mathbb{P}(j_t | S_t)$ according to (1). Define the matrix $A \in \{0, 1\}^{\hat{T} \times N}$, with elements $a_{ti} = 1$ if $i \in \mathcal{M}_t(j_t, S_t)$ (here, j_t could also be zero), and $a_{ti} = 0$ otherwise. The matrix A has one row per period, and one column per customer type, with $a_{ti} = 1$ when customer type $\sigma^{(i)}$ is compatible with the transaction observed in period t . Problem (3) can be reformulated in terms of the aggregate likelihoods y_t , for $\mathcal{L}_I(\mathbf{y}) = \sum_{t \in \mathcal{P}} \log(y_t) + \sum_{t \in \bar{\mathcal{P}}_\lambda} \log(y_t)$:

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{y} > 0} \mathcal{L}_I(\mathbf{y}) \\ \text{s.t.} \quad \sum_{i=1}^N x_i - 1 = 0, \\ A\mathbf{x} - \mathbf{y} = 0. \end{aligned} \quad (4)$$

As discussed in van Ryzin and Vulcano (2015, Section 3.3), this choice model is nonidentifiable in general⁵ (e.g., when we consider the full set of $O(n!)$ types).

This downside is somewhat mitigated by the next result, which follows from an adaptation of van Ryzin and Vulcano (2015, Proposition 2).

Proposition 1. *If for each offer set S that appears in the selling horizon, and for each product $j \in S$, there is at least one observed transaction (j, S) in the data set, then there is a unique vector \mathbf{y}^* that solves problem (4).*

Assuming that for the ground-truth choice model, $\mathbb{P}(j_t | S_t) > 0$ for all $j_t \in S_t$, the condition about the existence of an observation (j_t, S_t) for all $j_t \in S_t$ is mild under a large enough sample. In essence, although there may not be a unique pmf \mathbf{x}^* , Proposition 1 states that the likelihood-maximizing probabilities \mathbf{y}^* of observing the sequence $\{(j_t, S_t): t = 1, \dots, T\}$ are unique, implying that the choice model is (at least) *partially identifiable*. The proof follows from splitting the likelihood function in a collection of disjoint, strictly concave programs with a single linear constraint.

Corollary 1, based on the uniqueness of the solution of the system $A\mathbf{x} = \mathbf{y}^*$, follows easily.

Corollary 1. *If $\text{rank}(A) = N$, then there is a unique solution \mathbf{x}^* that solves (3) given a unique solution \mathbf{y}^* .*

Hence, if in addition matrix A is full rank, then the model would also be identifiable with respect to the true proportions \mathbf{x} . The next result formalizes this observation.

Corollary 2. *Suppose the sequence of data grows such that for each set S observed in the sequence of offer sets, $\lim_{T \rightarrow \infty} (\sum_{t=1}^T \mathbf{1}\{S_t = S\} / T) = q_S > 0$ (a.s.); that is, each set observed appears infinitely often, in a positive fraction of the intervals (a.s.). Then as the sample size $T \rightarrow \infty$, $\lambda^*(T)$ satisfies $\lambda^*(T) \rightarrow \lambda$ w.p.1. In addition, for each observed offer set S , $\mathbf{y}_i^*(T) \rightarrow \mathbb{P}(j_i | S_t)$ w.p.1, and any optimal solution $\mathbf{x}^*(T)$ to problem (3) satisfies $A_t \mathbf{x}^*(T) \rightarrow \mathbb{P}(j_t | S_t)$ w.p.1, where A_t is the t th row of A . Furthermore, if $\text{rank}(A) = N$, then $\mathbf{x}^*(T) \rightarrow \mathbf{x}$ w.p.1.*

The identifiability of the choice model with respect to the arrival rate λ and purchase probabilities $\mathbb{P}(j_t | S_t)$ (and potentially, also with respect to the proportions \mathbf{x}), make λ^* and \mathbf{y}^* (and, if applicable, \mathbf{x}^*) inherit the desirable statistical properties of ML estimators: They are consistent, asymptotically unbiased, and asymptotically efficient (i.e., they attain the Cramer-Rao lower bound for the variance, asymptotically).

3.3. EM Method

The building block for our simplified approach to solve the estimation problem (3) is to consider the *complete data log-likelihood function*; that is, the likelihood function we would get if we could directly observe the precise type of each customer. This results in a trivial MLE problem. The EM algorithm exploits this simplification by making use of the complete data

log-likelihood, but replacing the complete data with beliefs (expectations) about their values conditional on the observed data and current parameter estimates.

3.3.1. The Complete Data Log-Likelihood Function.

Define σ_t as the arrival type in period t . The complete data log-likelihood function associated with problem (3) is given by:

$$\begin{aligned} \mathcal{L}_C(\mathbf{x}) &= \sum_{t \in \mathcal{P}} \sum_{i \in \mathcal{M}_t(j_t, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \\ &\quad + \sum_{t \in \mathcal{P}_\lambda} \sum_{i \in \mathcal{M}_t(0, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \\ &= \sum_{i=1}^N \left(\sum_{t \in \mathcal{P}} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(j_t, S_t)\} \right. \\ &\quad \left. + \sum_{t \in \mathcal{P}_\lambda} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(0, S_t)\} \right) \log x_i \\ &= \sum_{i=1}^N m_i \log x_i, \end{aligned} \quad (5)$$

where $\mathbf{I}\{\cdot\}$ is the indicator function,

$$\begin{aligned} m_i &= \sum_{t \in \mathcal{P}} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(j_t, S_t)\} \\ &\quad + \sum_{t \in \mathcal{P}_\lambda} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(0, S_t)\}. \end{aligned} \quad (6)$$

That is, m_i counts the number of occurrences of a type $\sigma^{(i)}$ arrival, in periods with and without purchases. For $m_i > 0, i = 1, \dots, N$, the problem of maximizing $\mathcal{L}_C(\mathbf{x})$ can be posed as follows:

$$\max_{\mathbf{x} > 0} \sum_{i=1}^N m_i \log x_i, \quad \text{s.t.} \quad \sum_{i=1}^N x_i = 1. \quad (7)$$

This is a concave program for which the Karush-Kuhn-Tucker (KKT) conditions give unique maximizers $x_i^* = m_i / \sum_{h=1}^N m_h$, i.e., x_i^* is simply the number of type $\sigma^{(i)}$ customers observed divided by the total number of customers observed.

3.3.2. The Two Main Steps of the EM Algorithm.

The EM method works by starting with arbitrary initial estimates $\hat{\mathbf{x}}$ of the parameters. These estimates are then used to compute the conditional expected value of \mathcal{L}_C : $E[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}]$ (the ‘‘E,’’ expectation, step). Effectively, this replaces the missing data (i.e., the number of occurrences of each preference list) by their expected values conditioned on the current estimates. The resulting expected log-likelihood function is then maximized to generate new estimates $\hat{\mathbf{x}}$ (the ‘‘M,’’ maximization, step), and the procedure is repeated so that a sequence $\{\hat{\mathbf{x}}^{(k)}, k = 1, 2, \dots\}$ is generated. We next describe the two steps of each iteration. Section A1.1 in the online appendix provides the pseudocode.

The E-step. In the E-step, the unknown data values are the values of the indicators in (6). However, given current estimates $\hat{\mathbf{x}}$, we can determine their expected values. First, we use Bayes' theorem to update the probability mass function over the set of rankings:

$$\begin{aligned} \mathbb{P}(\sigma^{(i)} | j_t, S_t, \hat{\mathbf{x}}) &= \frac{\mathbb{P}(j_t | \sigma^{(i)}, S_t, \hat{\mathbf{x}}) \mathbb{P}(\sigma^{(i)} | \hat{\mathbf{x}})}{\mathbb{P}(j_t | S_t, \hat{\mathbf{x}})} \quad (8) \\ &= \frac{\mathbf{I}\{i \in \mathcal{M}_t(j_t, S_t)\} \hat{x}_i}{\sum_{h \in \mathcal{M}_t(j_t, S_t)} \hat{x}_h}, \end{aligned}$$

where $\mathbb{P}(j_t | \sigma^{(i)}, S_t, \hat{\mathbf{x}})$ in the numerator of (8) stands for the conditional probability of choosing product $j_t \in \mathcal{N}$. In particular, the case $j_t = 0$ implements the update for the no-purchase probability.

From (6) and (8), we obtain the estimates

$$\begin{aligned} \hat{m}_i &= \mathbb{E}[m_i | \hat{\mathbf{x}}] \\ &= \sum_{t \in \mathcal{P}} \mathbb{P}(\sigma^{(i)} | j_t, S_t, \hat{\mathbf{x}}) + \sum_{t \in \mathcal{P}_\lambda} \mathbb{P}(\sigma^{(i)} | 0, S_t, \hat{\mathbf{x}}). \quad (9) \end{aligned}$$

In words, \hat{m}_i represents the conditional expected number of occurrences of the preference list $\sigma^{(i)}$. Substituting \hat{m}_i into (5), we obtain the expected, complete data log-likelihood function

$$\mathbb{E}[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}] = \sum_{i=1}^N \hat{m}_i \log x_i. \quad (10)$$

This is the function to be maximized in the current iteration.

The M-step. To determine a maximizer $\hat{\mathbf{x}}^*$ of (10), we use the results from formulation (7). The function $\mathbb{E}[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}]$ is globally concave with a unique, closed-form maximizer given by:

$$\hat{x}_i^* = \frac{\hat{m}_i}{\sum_{h=1}^N \hat{m}_h}, \quad i = 1, \dots, N. \quad (11)$$

The simplicity of this maximization step is the most appealing feature of the EM algorithm.

3.3.3. Convergence. The concavity of the complete data log-likelihood function guarantees that our procedure is an EM algorithm, a special instance of the so-called Generalized EM algorithm (GEM). In the case of GEM, the M-step requires only that we generate an improved set of estimates over the current ones (i.e., it requires to find improved estimates \mathbf{x} such that $\mathbb{E}[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}] \geq \mathbb{E}[\mathcal{L}_C(\hat{\mathbf{x}}) | \hat{\mathbf{x}}]$), and the conditions for convergence are more stringent (e.g., see McLachlan and Krishnan 1996, Chapter 3 for further discussion.)

The following result argues that since our EM method satisfies a mild regularity condition, the sequence of likelihoods converges to a stationary value of the incomplete-data log-likelihood function (2).

Proposition 2. *All limit points of any instance $\{\hat{\mathbf{x}}^{(k)}, k = 1, 2, \dots\}$ of the EM algorithm are stationary points of the corresponding incomplete-data log-likelihood function $\mathcal{L}_I(\mathbf{x})$, and*

$\{\mathcal{L}_I(\hat{\mathbf{x}}^{(k)}), k = 1, 2, \dots\}$ converges monotonically to a value $\mathcal{L}_I^ = \mathcal{L}_I(\mathbf{x}^*)$, for some stationary point \mathbf{x}^* .*

Proof. The conditional expected value $\mathbb{E}[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}]$ in (10) is continuous in $x > 0$ and $\hat{x} > 0$. The result follows from the fact that \hat{m}_i in (9) is continuous in $\hat{\mathbf{x}}$ according to Equation (8). Clearly, $\mathbb{E}[\mathcal{L}_C(\mathbf{x}) | \hat{\mathbf{x}}]$ is also continuous in \mathbf{x} . Therefore, from Wu (1983, Theorem 2) (see also McLachlan and Krishnan 1996, Theorem 3.2) the EM algorithm, given the unique maximizer found in the M-step, generates an implied sequence $\{\mathcal{L}_I(\hat{\mathbf{x}}^{(k)}), k = 1, 2, \dots\}$ whose limit point is a stationary point. \square

As pointed out by Wu (1983, Section 2.2), the convergence of $\{\mathcal{L}_I(\hat{\mathbf{x}}^{(k)})\}$ to $\mathcal{L}_I(\mathbf{x}^*)$, for some stationary point \mathbf{x}^* , does not automatically imply the convergence of $\{\hat{\mathbf{x}}^{(k)}\}$ to a point \mathbf{x}^* . Convergence of the EM estimator in the later sense usually requires more stringent regularity conditions. Boyles (1983) further investigates these requirements. In his Theorem 2, he identifies sufficient conditions under which the sequence of iterates $\hat{\mathbf{x}}^{(k)}$ converges to a compact, connected component of stationary points of the incomplete data log-likelihood function with value \mathcal{L}_I^* . Most of the conditions are rather mild and are satisfied by our algorithm. The stringent condition is the one that requires

$$\|\hat{\mathbf{x}}^{(k+1)} - \hat{\mathbf{x}}^{(k)}\| \rightarrow 0, \quad \text{as } k \rightarrow \infty. \quad (12)$$

As Boyles (1983) comments, it is difficult to verify (12) in general. Yet under this condition, the sequence $\{\hat{\mathbf{x}}^{(k)}\}$ will seek an isolated plateau of stationary points, and for k sufficiently large will not leap over valleys to neighboring plateaux. See also Wu (1983, Theorem 5).

Nevertheless, as a practical matter, the convergence of the sequence of points $\{\hat{\mathbf{x}}^{(k)}\}$ can be checked numerically as part of the EM procedure. In our experiments reported in Section 5, we consistently observed that the sequence of points visited by EM converges to a limit point.

4. Estimation: Censored Demand Case

4.1. Problem Formulation

When, in addition to the nonobservability of the customer types, we cannot distinguish a period with no arrival from a period with an arrival who did not purchase, the incomplete data log-likelihood function is given by:

$$\begin{aligned} \mathcal{L}_I(\mathbf{x}, \lambda) &= \sum_{t \in \mathcal{P}} \left(\log \lambda + \log \left(\sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i \right) \right) \\ &\quad + \sum_{\substack{t \in \mathcal{P}, \\ \mathcal{M}_t(0, S_t) \neq \emptyset}} \log \left(\lambda \sum_{i \in \mathcal{M}_t(0, S_t)} x_i + (1 - \lambda) \right) \\ &\quad + \sum_{\substack{t \in \mathcal{P}, \\ \mathcal{M}_t(0, S_t) = \emptyset}} \log(1 - \lambda), \quad (13) \end{aligned}$$

where $\bar{\mathcal{P}} = \bar{\mathcal{P}}_\lambda \cup \bar{\mathcal{P}}_{\bar{\lambda}}$. The first term accounts for the likelihood of the observed transactions in the periods with purchases. The second term accounts for the no-purchase periods, where an arriving customer preferred not to buy (i.e., customer types in $\mathcal{M}_t(0, S_t)$) or no customer arrived at all. The third term accounts for the periods where none of the customer types would have picked a product from the available assortment, and hence $|\mathcal{M}_t(0, S_t)| = 0$. This last case indicates with certainty that no arrival occurred.

Consider the problem:

$$\max_{\mathbf{x} > 0, \lambda > 0} \mathcal{L}_I(\mathbf{x}, \lambda) \quad \text{s.t.} \quad \sum_{i=1}^N x_i = 1, \quad \lambda < 1. \quad (14)$$

The function $\mathcal{L}_I(\mathbf{x}, \lambda)$ is known not to be quasi-concave in general (see van Ryzin and Vulcano 2015, Proposition A1). The model is not even partially identifiable with respect to the aggregate probabilities $y_t = \sum_{i \in \mathcal{M}_t(j_t, S_t)} x_i$, and multiple local optima may exist. One can attempt to solve it using a standard nonlinear optimization package, but the incompleteness in the data creates a challenging estimation problem that again can be greatly simplified using a problem-specific EM method.

4.2. EM Method

As in Section 3, we first introduce the *complete data log-likelihood function* and then describe the EM algorithm.

4.2.1. The Complete Data Log-Likelihood Function.

Initially, assume that we can overcome the first source of incompleteness by distinguishing the periods with arrivals from the periods with no-arrivals. For the periods with no observed transactions, let $a_t = 1$ if there is an arrival in the period, and $a_t = 0$ otherwise. Note that $a_t = 1$ accounts for an arrival that buys an outside option or does not buy at all.

Next, suppose we can also distinguish the type of the arrival (in case there is such an arrival), and denote $\sigma_t \in \{\sigma^{(1)}, \dots, \sigma^{(N)}\}$ the arrival type in period t . The complete data log-likelihood function is

$$\begin{aligned} \mathcal{L}_C(\mathbf{x}, \lambda) = & \sum_{t \in \mathcal{P}} \left(\log \lambda + \sum_{i \in \mathcal{M}_t(j_t, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \right) \\ & + \sum_{t \in \bar{\mathcal{P}}} a_t \left(\log \lambda + \sum_{i \in \mathcal{M}_t(0, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \right) \\ & + \sum_{t \in \bar{\mathcal{P}}} (1 - a_t) \log(1 - \lambda). \end{aligned} \quad (15)$$

The first term in (15) accounts for the observed purchases; there is one term for every period in which there is a sale. The second term accounts for customer arrivals that do not purchase at all. The third term accounts for periods with no arrivals. Note that the function \mathcal{L}_C is separable in \mathbf{x} and λ , and jointly concave.

Define the function $F(\lambda)$ as

$$\begin{aligned} F(\lambda) = & |\mathcal{P}| \log \lambda + \log \lambda \sum_{t \in \bar{\mathcal{P}}} a_t + |\bar{\mathcal{P}}| \log(1 - \lambda) \\ & - \sum_{t \in \bar{\mathcal{P}}} a_t \log(1 - \lambda). \end{aligned}$$

This function is globally concave in λ . Taking derivative and setting it equal to zero, we obtain the unique maximizer

$$\lambda^* = \frac{|\mathcal{P}| + \sum_{t \in \bar{\mathcal{P}}} a_t}{|\mathcal{P}| + |\bar{\mathcal{P}}|} = \frac{|\mathcal{P}| + \sum_{t \in \bar{\mathcal{P}}} a_t}{T},$$

which is simply the number of arrivals divided by the total number of periods. Clearly, λ^* satisfies $0 < \lambda^* < 1$.

Now, define

$$\begin{aligned} H(\mathbf{x}) = & \sum_{t \in \mathcal{P}} \sum_{i \in \mathcal{M}_t(j_t, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \\ & + \sum_{t \in \bar{\mathcal{P}}} a_t \sum_{i \in \mathcal{M}_t(0, S_t)} \mathbf{I}\{\sigma_t = \sigma^{(i)}\} \log x_i \\ = & \sum_{i=1}^N \left(\sum_{t \in \mathcal{P}} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(j_t, S_t)\} \right. \\ & \left. + \sum_{t \in \bar{\mathcal{P}}} a_t \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(0, S_t)\} \right) \log x_i \\ = & \sum_{i=1}^N m_i \log x_i, \end{aligned} \quad (16)$$

where

$$\begin{aligned} m_i = & \sum_{t \in \mathcal{P}} \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(j_t, S_t)\} \\ & + \sum_{t \in \bar{\mathcal{P}}} a_t \mathbf{I}\{\sigma_t = \sigma^{(i)}, i \in \mathcal{M}_t(0, S_t)\}. \end{aligned} \quad (17)$$

That is, m_i counts the number of occurrences of a type $\sigma^{(i)}$ arrival, in periods with and without purchases. For $m_i > 0$, $i = 1, \dots, N$, the problem of maximizing $H(\mathbf{x})$ can be posed as follows:

$$\max_{\mathbf{x} > 0} \sum_{i=1}^N m_i \log x_i, \quad \text{s.t.} \quad \sum_{i=1}^N x_i = 1.$$

The objective function is globally concave, with unique maximizers $x_i^* = m_i / \sum_{h=1}^N m_h$, which is simply the number of type $\sigma^{(i)}$ customers observed divided by the total number of customers observed.

4.2.2. The Two Main Steps of the EM Algorithm. As before, the EM method starts with arbitrary initial estimates of the parameters $\hat{\theta} = (\hat{\mathbf{x}}, \hat{\lambda})$. These estimates are then used to compute the conditional expected value of \mathcal{L}_C : $E[\mathcal{L}_C(\theta) | \hat{\theta}]$ (the E-step). The resulting expected log-likelihood function is then maximized to generate new estimates $\hat{\theta}$ (the M-step), and the procedure is repeated until convergence. The two steps of each iteration are described below; the pseudocode is given in the appendix (Section A1.2).

The E-step. In the E-step, the unknown data values are the a_t for all $t \in \bar{\mathcal{P}}$. However, given current estimates $\hat{\theta} = (\hat{\mathbf{x}}, \hat{\lambda})$, we can determine the expected values of these arrival indicators. Again using Bayes' theorem, we get the update for \hat{a}_t . As it is implicit in (15), if $t \in \mathcal{P}$, then $\hat{a}_t = 1$. For $t \in \bar{\mathcal{P}}$,

$$\begin{aligned} \hat{a}_t &= \mathbb{E}[a_t | t \in \bar{\mathcal{P}}, \hat{\theta}] = \mathbb{P}(a_t = 1 | t \in \bar{\mathcal{P}}, \hat{\theta}) \\ &= \frac{\mathbb{P}(t \in \bar{\mathcal{P}} | a_t = 1, \hat{\theta}) \mathbb{P}(a_t = 1 | \hat{\theta})}{\mathbb{P}(t \in \bar{\mathcal{P}} | \hat{\theta})} \\ &= \frac{\mathbb{P}(0 | t, S_t, \hat{\theta}) \hat{\lambda}}{\hat{\lambda} \mathbb{P}(0 | t, S_t, \hat{\theta}) + (1 - \hat{\lambda})}, \end{aligned} \quad (18)$$

where

$$\mathbb{P}(0 | t, S_t, \hat{\theta}) = \sum_{i \in \mathcal{M}_t(0, S_t)} \hat{x}_i$$

represents the conditional probability of no-purchasing. If $\mathcal{M}_t(0, S_t) = \emptyset$, then $\mathbb{P}(0 | t, S_t, \hat{\theta}) = 0$, and therefore $\hat{a}_t = 0$. This is the case when all products are offered and no transaction occurs, which reveals that no arrival occurs.

From (17), we get the estimates

$$\hat{m}_i = \mathbb{E}[m_i | \hat{\theta}] = \sum_{t \in \mathcal{P}} \mathbb{P}(\sigma^{(i)} | j_t, S_t, \hat{\theta}) + \sum_{t \in \bar{\mathcal{P}}} \hat{a}_t \mathbb{P}(\sigma^{(i)} | 0, S_t, \hat{\theta}),$$

where $\mathbb{P}(\sigma^{(i)} | j_t, S_t, \hat{\theta})$ is defined as in (8), potentially with $j_t = 0$. In words, \hat{m}_i represents the conditional expected number of occurrences of the preference list $\sigma^{(i)}$. Note that the following balance equation holds: $\sum_{i=1}^N \hat{m}_i = \sum_{t=1}^T \hat{a}_t$.

Substituting \hat{a}_t and \hat{m}_i into (15), we obtain the expected, complete data log-likelihood function

$$\begin{aligned} \mathbb{E}[\mathcal{L}_C(\mathbf{x}, \lambda) | \hat{\theta}] &= \sum_{i=1}^N \hat{m}_i \log x_i + \left(|\mathcal{P}| + \sum_{t \in \bar{\mathcal{P}}} \hat{a}_t \right) \log \lambda \\ &\quad + \left(|\bar{\mathcal{P}}| - \sum_{t \in \bar{\mathcal{P}}} \hat{a}_t \right) \log(1 - \lambda). \end{aligned} \quad (19)$$

This is the function to be maximized in the current iteration.

The M-step. To determine a maximizer $\hat{\theta}^*$ of (19), we use the results from Equation (15). The function $\mathbb{E}[\mathcal{L}_C(\mathbf{x}, \lambda) | \hat{\theta}]$ is globally concave in (\mathbf{x}, λ) , separable in both variables, and has unique, closed-form maximizers:

$$\hat{\lambda}^* = \frac{|\mathcal{P}| + \sum_{t \in \bar{\mathcal{P}}} \hat{a}_t}{T}, \quad \text{and} \quad \hat{x}_i^* = \frac{\hat{m}_i}{\sum_{k=1}^N \hat{m}_k}, \quad i = 1, \dots, N.$$

Again, closed form expressions for the maximization step are a very appealing feature of this EM algorithm.

Following the arguments for the uncensored demand case, the convergence properties of Proposition 2 also hold for the censored demand case.

5. Numerical Examples

Next we present a set of numerical examples based on synthetic and real-world data sets. All our experiments were conducted using MATLAB,⁶ in which our EM proposal is extremely simple to code. In all the examples, we set the stopping criteria for EM based on the difference between the points produced by two consecutive iterations, halting the procedures as soon as the norm of the difference vector was smaller than 1e-5. In fact, in all the experiments below, we verified the numerical convergence of the EM iterates to a limit point, though such convergence in theory is not guaranteed a priori. We also set a maximum number of iterations at 1e7.

5.1. Experiments Based on Synthetic Data

For experiments based on synthetic data, we assume complete market information in the sense that the modeler knows the description of the existing customer types, but does not know their proportions. The objective of these experiments is to assess the performance of our EM method relative to direct maximization of the incomplete data log-likelihood function (labeled *Direct Max* below) for uncensored and censored demand scenarios, with respect to computational speed and quality of the estimates.

We consider two versions of Direct Max, V1 and V2, based on different implementations of the built-in MATLAB function “fmincon,” which finds a constrained minimum of a function of several variables, in our case subject to linear constraints and 0–1 bounds for the decision variables (\mathbf{x}, λ) (see (3) and (14)). Direct Max V1 uses sequential quadratic programming (SQP), closely mimicking Newton's method for constrained optimization just as is done for unconstrained optimization. At each major iteration, an approximation is made of the Hessian of the Lagrangian function using a quasi-Newton updating method. This is then used to generate a quadratic programming subproblem whose solution is used to form a search direction for a line search procedure. Direct Max V2 implements a large-scale, interior point, barrier-type algorithm that exploits the sparsity structure of the problem. For V1 and V2 we set the same tolerance 1e-5 as for EM; similarly, we establish the iteration and function evaluation limits at 1e7.

The data generation proceeds as follows: We fixed the number of products at $n = 15$. For the uncensored demand case, we made the simplification $\lambda = 1$.⁷ For the censored demand case we consider two demand scenarios, i.e., a low volume scenario (with $\lambda = 0.2$) and a high volume scenario (with $\lambda = 0.8$). We simulated 30 instances of transaction data for different combinations of length of selling season $T \in \{5,000, 10,000, 50,000, 100,000\}$, and number of customer types $N \in \{10, 30, 50, 100\}$. The types themselves were generated

by sampling random permutations of the n products along with the no-purchase alternative, and the true proportions were sampled from the $\text{Unif}(0,1)$ distribution. For any instance, in every period, a random number of available products between 2 and 10 was generated.

The starting point for EM and Direct Max was the same, i.e., $x_i = 1/N$ (and $\lambda = 0.5$ for the censored demand case). The quality of the estimates is evaluated on a hold-out data sample of the same size T as the generated data.

5.1.1. Uncensored Demand. We note that starting from the same initial points, and setting the same tolerance level, EM and Direct Max (V1 and V2) produce estimates of very similar quality in terms of relative errors with respect to the true underlying proportions (as reported in Figure A1 in the appendix) and the log-likelihood values reached, where we observe differences of order $1e-4$.

The key distinction lies in the computational times. Table 1 shows 95% confidence intervals (CI) of the mean time for estimating the parameters of a single instance. When comparing both versions of Direct Max we observe that V2 is around 20% faster than V1. Yet EM clearly dominates Direct Max, by factors between 1.6 and 6.1, with an average factor of 3.8. The difference is quite significant for practical applications where the consumer choice estimation must be performed at the consideration set level. For instance, a major carrier in an airline revenue management setting must estimate hundreds of thousands of O-D pairs on a daily or even more frequent basis. The computational advantage is also critical in the context of the market discovery algorithm that we propose in van Ryzin and Vulcano (2015), where the MLE procedure is repeatedly called as a subroutine.

5.1.2. Censored Demand. Next we consider a similar data setting for the censored demand case. Since arrivals that ended in a no-purchase are nonobservable, the modeler now must jointly estimate the arrival rate λ and the customer type proportions. Figure A2 in the appendix reports estimation errors (i.e., relative differences between true and estimated parameters (x, λ)) for the case where the underlying true model is generated with $\lambda = 0.2$; Figure A3 does it for a ground truth model with $\lambda = 0.8$. Again, we observe a very similar quality of fit for Direct Max (V1 and V2) and EM methods. Compared to the uncensored demand case, for a given value of T , we see a slightly bigger dispersion of errors, particularly for the low volume demand case where more no-arrivals occur and therefore the accuracy in distinguishing a no-arrival from a no-purchase diminishes. In addition, as in the uncensored demand case, we observe the same log-likelihood values from both approaches (up to 3 decimals).

Table 2 reports the mean computational times for both methods, for each of the demand volume cases. We still observe a significant difference in favor of EM, by factors up to 3.2 for the low volume demand case (with an average factor of 1.9), and by factors between 1.2 and 5.8 for the high volume demand case (with an average factor of 2.9). Interestingly, when the demand volume increases from $\lambda = 0.2$ to $\lambda = 0.8$ and the other parameters remain constant, the computational time of Direct Max increases by 22% whereas EM reduces by around 10%, on average.

5.2. Experiment Based on Real Data

Next we present results applying our EM method to a publicly-available hotel data set (see Bodea et al. 2009), also used in our paper van Ryzin and Vulcano (2015). The booking records correspond to transient customers

Table 1. Computational times (seconds) for a single instance of the uncensored demand case, in a market with $n = 15$ products: 95% CI for the mean

T	N	Direct Max V1	Direct Max V2	EM
5,000	10	0.3601 ± 0.0439	0.3177 ± 0.0204	0.0708 ± 0.0179
	30	2.4732 ± 0.0386	2.1081 ± 0.0894	0.6960 ± 0.0463
	50	7.1881 ± 0.0992	5.9516 ± 0.1465	1.6865 ± 0.0938
	100	31.8209 ± 0.5510	29.5816 ± 0.7943	5.2106 ± 0.1295
10,000	10	0.5920 ± 0.0141	0.4752 ± 0.0215	0.1264 ± 0.0216
	30	4.4655 ± 0.0639	3.6583 ± 0.0948	1.4516 ± 0.1006
	50	13.2047 ± 0.1623	11.7561 ± 0.3861	3.2601 ± 0.1403
	100	68.3841 ± 0.8304	63.6594 ± 1.4255	12.2334 ± 0.2886
50,000	10	2.7910 ± 0.0595	2.0296 ± 0.0790	0.6485 ± 0.1135
	30	23.7017 ± 0.3311	18.1254 ± 0.4670	9.9305 ± 0.8306
	50	69.7365 ± 0.8066	58.8844 ± 1.2662	21.8394 ± 0.9300
	100	396.1877 ± 2.5411	365.2032 ± 11.8077	68.1128 ± 1.4271
100,000	10	5.7643 ± 0.1033	4.0734 ± 0.3515	1.2615 ± 0.2525
	30	47.2118 ± 0.8721	36.8851 ± 1.4604	23.0733 ± 1.4921
	50	141.6512 ± 1.5072	115.7827 ± 4.1309	58.4436 ± 3.6771
	100	814.0861 ± 7.1345	709.8073 ± 24.7901	199.9066 ± 4.2604

Table 2. Computational times (seconds) for a single instance of the censored demand case, in a market with $n = 15$ products: 95% CI for the mean

T	N	$\lambda = 0.2$			$\lambda = 0.8$		
		Direct Max V1	Direct Max V2	EM	Direct Max V1	Direct Max V2	EM
5,000	10	0.3837 ± 0.0121	0.4133 ± 0.0346	0.1868 ± 0.0522	0.4512 ± 0.0083	0.3987 ± 0.0218	0.1427 ± 0.0204
	30	2.1157 ± 0.0586	2.0468 ± 0.1181	1.3003 ± 0.1337	2.8049 ± 0.0493	2.5299 ± 0.1186	1.2388 ± 0.0782
	50	4.9033 ± 0.1144	4.4996 ± 0.1591	2.5067 ± 0.1641	7.3641 ± 0.1264	6.3656 ± 0.2142	2.5465 ± 0.1175
	100	16.1336 ± 0.3224	15.1294 ± 0.3665	5.8789 ± 0.3568	30.4199 ± 0.4774	27.9002 ± 0.5891	6.3916 ± 0.2917
10,000	10	0.7037 ± 0.0163	0.6578 ± 0.0438	0.3172 ± 0.0736	0.8273 ± 0.0127	0.6631 ± 0.0262	0.2261 ± 0.0320
	30	4.0955 ± 0.1035	3.9321 ± 0.2440	2.7104 ± 0.2368	5.3964 ± 0.0779	4.8466 ± 0.2211	2.5177 ± 0.1140
	50	10.1392 ± 0.2600	9.4767 ± 0.2998	5.7965 ± 0.3970	14.6129 ± 0.1806	12.5659 ± 0.4598	5.2899 ± 0.2733
	100	36.1058 ± 0.7274	34.5616 ± 0.8115	14.2153 ± 0.6690	65.5522 ± 0.9853	59.4967 ± 1.2140	13.4244 ± 0.6002
50,000	10	3.2092 ± 0.0555	2.8275 ± 0.1974	1.6003 ± 0.4094	4.0803 ± 0.0864	3.1421 ± 0.2071	1.6240 ± 0.3513
	30	22.3262 ± 0.4279	21.5134 ± 1.2864	20.1628 ± 1.8622	28.6463 ± 0.4011	23.7095 ± 1.0711	20.4073 ± 2.0080
	50	58.8535 ± 1.0970	58.6661 ± 2.6471	34.9383 ± 2.9590	79.1596 ± 1.1829	65.9929 ± 2.1056	33.1007 ± 1.6596
	100	259.8091 ± 5.3822	261.5783 ± 7.4833	82.1049 ± 4.8107	399.2866 ± 3.7449	360.0141 ± 7.2055	68.7775 ± 2.7535
100,000	10	6.9600 ± 0.1839	5.8617 ± 0.4361	3.1176 ± 0.7368	8.2890 ± 0.1578	6.0308 ± 0.2463	2.6801 ± 0.3904
	30	47.7381 ± 0.9713	48.0468 ± 3.4149	54.2606 ± 5.3650	59.1728 ± 0.6496	48.2268 ± 2.5505	41.7799 ± 3.1860
	50	124.8093 ± 1.7413	124.1907 ± 6.8489	102.3223 ± 7.5898	161.5316 ± 1.5883	134.1275 ± 4.2230	88.1674 ± 5.5703
	100	553.3401 ± 8.7139	541.8066 ± 27.8692	237.9211 ± 14.5738	832.2295 ± 7.4274	718.1243 ± 21.5686	203.4727 ± 9.5043

(predominately business travelers) with check-in dates between March 12, 2007, and April 15, 2007 in one of five continental U.S. hotels. For every hotel, a minimum booking horizon of four weeks for each check-in date was considered. Rate and room type availability information present at the time of booking were recorded for reservations made via the hotel or customer relationship officers (CROs), the hotels' web sites, and off-line travel agencies. The data was preprocessed to comply with the model assumptions (e.g., there is at least one transaction per product, and observed transactions must come from available options).

We define a *product* as a room type (e.g., king non-smoking, queen smoking, suite type 1, etc.), and a period as a (booking date, check-in date) pair. The original data set corresponds only to booking records, and therefore we assume an uncensored demand case with arrival rate $\lambda = 1$. Table 3 summarizes further details and the estimation results of a rank-based model capturing buy-ups relative to the independent demand

model. The buy-up model was built assuming that customers are price sensitive, and that for a similar price, they prefer the following room type order: (1) suite, (2) king, and (3) queen. If the price is not the same across product types, we assume that customers tend to buy-up within the same room type. For each hotel, we sifted through the data to check the price order among the different room types, and defined a set of preference lists capturing this simple behavior. We provide the transaction and availability data after our preprocessing, and the defined customer types, as a supplemental e-companion to this manuscript.

Given the limited amount of data for two of the hotels, our goodness-of-fit measures are all in-sample. We report the log-likelihood values (since they both correspond to the same log-likelihood function given in (2)) and the root mean square error (RMSE) between the predicted and the observed bookings aggregated over the selling horizon. The RMSE is an absolute measure of fit between predicted and observed

Table 3. Estimation results for the hotel example

Feature	Hotel 1	Hotel 2	Hotel 3	Hotel 4	Hotel 5
Number of products after preprocessing, n	10	12	8	9	8
Number of periods, \hat{T}	1,315	211	1,147	288	245
Availability over selling horizon (%)	64	74	88	77	86
Indep. demand					
Log-value	-2,621	-378	-1,569	-416	-403
RMSE	40.91	6.10	47.90	7.37	7.25
AIC _c	5,262	782	3,155	851	822
Buy-up					
Number of types	10	23	14	15	14
Log-value	-2,194	-318	-1,376	-375	-371
RMSE	20.77	3.06	15.21	4.43	4.44
AIC _c	4,408	687	2,780	782	772

purchases of the n products, defined as $\text{RMSE} = (\sum_{j=1}^n (\sum_{t=1}^{\hat{T}} (\mathbb{P}(j|S_t) - \mathbf{I}\{j_t = j\}))^2 / n)^{1/2}$. A zero value indicates perfect fit. One downside of the RMSE and the log-likelihood value is that they do not include any penalty for model complexity, which increases with more types in the choice model. To account for this, we also report the corrected Akaike information criterion, defined as: $\text{AIC}_c = 2(N - \mathcal{L}_1(\mathbf{x}) + N(N + 1)/(\hat{T} - N - 1))$, where N is the number of parameters (i.e., customer types) in the model, $\mathcal{L}_1(\mathbf{x})$ is the maximized value of the log-likelihood function, and \hat{T} is the sample size (i.e., number of periods). The AIC_c measure rewards the log-likelihood value but also penalizes model complexity (captured by the number of parameters) to control for over fitting. Lower values of AIC_c indicate a better fit.

In the five cases, the types defined based on the simple buy-up principle outperform the traditional independent demand assumption with respect to log-likelihood values, RMSE, and AIC_c . In particular, the latter indicate that the buy-up model provides a good compromise between explanatory power and complexity. In our previous paper van Ryzin and Vulcano (2015), we show the performance of a market discovery procedure that automates the inference of new types over the same data set. In fact, we used EM in each estimation step of that procedure. Not surprisingly, the performance is not better than the *rank-based* performance in van Ryzin and Vulcano (2015) with respect to log-likelihood and AIC_c values (see Table 3 therein),⁸ but the purpose of the experiment is to demonstrate that even a simple judgmental rule can provide a significant improvement over the standard independent demand model, and that our EM method is an effective means to estimate such a model.

6. Conclusions

In this paper, we propose an EM algorithm for estimating rank-based preferences for a given category of products using only sales transaction and product availability data. The demand model is quite general and compatible with any random utility model. The algorithm we propose is easy to implement in any procedural language, and its performance in terms of quality of the derived estimates is comparable to estimates obtained by maximizing the incomplete data log-likelihood function using standard numerical techniques. However, our EM method is computationally remarkably faster than competitive nonlinear optimization procedures by time factors up to six. The difference is even more pronounced when the procedures are executed on censored demand cases, which require the joint estimation of demand volumes and customer type proportions. This combination of simplicity and computational performance makes our EM method

a very appealing procedure for real-world estimation problems under the rank-based choice model of demand.

Endnotes

- ¹To our knowledge, Celect's assortment optimization platform (Celect, Inc. 2014) is one of the first commercial packages to implement this choice model.
- ²For instance, for dairy products, n could typically range between 20–60 (e.g., Honhon et al. 2012).
- ³For instance, in an airline revenue management setting, a type could be defined by price sensitive customers (e.g., customers who just prefer the low fare versus others who may be willing to buy-up to a medium or high fare), other types may be defined by time sensitive customers (e.g., customers who prefer morning flights to afternoon flights, and vice versa).
- ⁴Otherwise, if $S_t = \{0\}$, then the period is not informative and can be disregarded.
- ⁵We say that the parameters \mathbf{x} corresponding to the distribution over types σ is *identified* if for any alternative parameters \mathbf{x}' , $\mathbf{x} \neq \mathbf{x}'$, for some data (j_t, S_t) , $t = 1, \dots, \hat{T}$, we have $\mathcal{L}_1(\mathbf{x}) \neq \mathcal{L}_1(\mathbf{x}')$.
- ⁶We used MATLAB version 8.1 (R2013a) for Mac OS X on a CPU with Intel Core i7 processor and 8 GB of RAM.
- ⁷Recall that for the general case, the estimate for λ^* can be easily computed via the formula $\lambda^* = (|\mathcal{P}| + |\mathcal{P}'|)/T$.
- ⁸The RMSEs are calculated slightly differently in van Ryzin and Vulcano (2015), since there the no-purchase alternative is considered an explicit outside option, leading to $n + 1$ terms in the RMSE formula.

References

- Anupindi R, Dada M, Gupta S (1998) Estimation of consumer demand with stock-out based substitution: An application to vending machine products. *Marketing Sci.* 17(4):406–423.
- Bertsimas D, Mišić V (2015) Data-driven assortment optimization. Working paper, MIT Sloan School of Management, Cambridge, MA.
- Bodea TD, Ferguson ME, Garrow L (2009) Choice-based revenue management: Data from a major hotel chain. *Manufacturing Service Oper. Management* 11(2):356–361.
- Boyles R (1983) On the convergence of the EM algorithm. *J. Roy. Statist. Soc. (Series B)* 45(1):47–50.
- Chaneton J, Vulcano G (2011) Computing bid-prices for revenue management under customer choice behavior. *Manufacturing Service Oper. Management* 13(4):452–470.
- Chen L, Homem-de-Melo T (2010) Mathematical programming models for revenue management under customer choice. *Eur. J. Oper. Res.* 203(2):294–305.
- Conlon C, Mortimer J (2013) Demand estimation under incomplete product availability. *Amer. Econom. J.-Microeconomics* 5(4):1–30.
- Dempster A, Laird N, Rubin D (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. B* 39(1):1–38.
- Farias V, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Sci.* 59(2):305–322.
- Haensel A, Koole G (2011) Estimating unconstrained demand rate functions using customer choice sets. *J. Revenue Pricing Management* 10(5):438–454.
- Honhon D, Gaur V, Seshadri S (2010) Assortment planning and inventory decisions under stockout-based substitution. *Oper. Res.* 58(5):1364–1379.
- Honhon D, Jonnalagedda S, Pan X (2012) Optimal algorithms for assortment selection under ranking-based consumer choice models. *M&SOM* 12(2):279–289.

- Honhon D, Pan X, Jonnalagedda S (2015) Heuristics for the assortment planning problem under ranking-based consumer choice models. Working paper, Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX.
- Jagabathula S (2014) Assortment optimization under general choice. Working paper, New York University, New York.
- Jagabathula S, Rusmevichientong P (2016) A nonparametric joint assortment and price choice model. *Management Sci.*, ePub ahead of print July 22, <https://doi.org/10.1287/mnsc.2016.2491>.
- Kök G, Fisher M (2007) Demand estimation and assortment optimization under substitution: Methodology and application. *Oper. Res.* 55(6): 1001–1021.
- Kunnumkal S (2014) Randomization approaches for network revenue management with customer choice behavior. *Production Oper. Management* 23(9):1617–1633.
- Mahajan S, van Ryzin G (2001a) Inventory competition under inventory competition under dynamic consumer choice. *Oper. Res.* 49(5): 646–657.
- Mahajan S, van Ryzin G (2001b) Stocking retail assortments under dynamic consumer substitution. *Oper. Res.* 49(3):334–351.
- McLachlan G, Krishnan T (1996) *The EM Algorithm and Extensions* (John Wiley & Sons, New York).
- Smith J, Lim C, Alptekinoglu A (2009) New product introduction against a predator: A bilevel mixed-integer programming approach. *Naval Res. Logist.* 56(8):714–729.
- Stefanescu C (2009) Multivariate customer demand: Modeling and estimation from censored sales. Working paper, European School of Management and Technology, Berlin.
- van Ryzin G, Vulcano G (2008) Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing Service Oper. Management* 10(3):448–467.
- van Ryzin G, Vulcano G (2015) A market discovery algorithm to estimate a general class of nonparametric choice model. *Management Sci.* 61(2):281–300.
- Wu CF (1983) On the convergence properties of the EM algorithm. *Ann. Statist.* 11(1):95–103.
- Yunes T, Napolitano D, Scheller-Wolf A, Tayur S (2007) Building efficient product portfolios at John Deere and building efficient product portfolios at John Deere and company. *Oper. Res.* 55(4):615–629.
- Zhang D, Cooper W (2006) Revenue management for parallel flights with customer-choice behavior. *Oper. Res.* 53(3):415–431.

Garrett van Ryzin is the Paul M. Montrone Professor of Decision, Risk, and Operations at the Columbia University Graduate School of Business. His research interests include analytical pricing, stochastic modeling, and operations management.

Gustavo Vulcano is professor at the School of Business, Universidad Torcuato di Tella, Buenos Aires, Argentina, and associate professor in the Department of Information, Operations, and Management Sciences, Leonard N. Stern School of Business, New York University. His research interests include analytical pricing, stochastic modeling, and operations management.