Technical Note: An Expectation-Maximization Algorithm to Estimate the Parameters of the Markov Chain Choice Model

A. Serdar Şimşek[1], Huseyin Topaloglu[2]

August 1, 2017

Abstract

We develop an expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. In this choice model, a customer arrives into the system to purchase a certain product. If this product is available for purchase, then the customer purchases it. Otherwise, the customer transitions between the products according to a transition probability matrix until she reaches an available one and purchases this product. The parameters of the Markov chain choice model are the probability that the customer arrives into the system to purchase each one of the products and the entries of the transition probability matrix. In our expectation-maximization algorithm, we treat the path that a customer follows in the Markov chain as the missing piece of the data. Conditional on the final purchase decision of a customer, we show how to compute the probability that the customer arrives into the system to purchase a certain product and the expected number of times that the customer transitions from a certain product to another one. These results allow us to execute the expectation step of our algorithm. Also, we show how to solve the optimization problem that appears in the maximization step of our algorithm. Our computational experiments show that the Markov chain choice model, coupled with our expectation-maximization algorithm, can yield better predictions of customer choice behavior when compared with other commonly used alternatives.

1 Introduction

Incorporating customer choice behavior into revenue management models has been seeing considerable attention. Traditional revenue management models capture the customer demand for a product through an exogenous random variable whose distribution does not depend on what other products are available. In reality, however, many customers choose and substitute among the available products in a particular product category, either because the customer arrives with no specific product in mind and makes a choice among the offered products or because the customer arrives with a specific product in mind and this product is not available. Field experiments,

[1]Naveen Jindal School of Management, The University of Texas at Dallas, Richardson, TX 75080, USA, serdar.simsek@utdallas.edu

[2]School of Operations Research and Information Engineering, Cornell Tech, New York, NY 10011, USA, topaloglu@orie.cornell.edu

customer surveys and controlled studies in Zinn and Liu (2001), Chong et al. (2001), Campo et al. (2003), Sloot et al. (2005) and Guadagni and Little (2008) indicate that customers indeed make a choice among the offered products after comparing them with respect to various features and substitute for another product when the product that they originally have in mind is not available. When customers choose and substitute, the demand for a particular product depends on what other products are available. Discrete choice models become useful to model the customer demand when the customers choose and substitute among the available products. Although discrete choice models can provide a more realistic model of the customer demand when compared with using an exogenous random variable, estimating the parameters of a discrete choice model can be challenging.

In this paper, we consider the problem of estimating the parameters of the Markov chain choice model. In this choice model, with a certain probability, a customer arriving into the system is interested in purchasing a certain product. If this product is available for purchase, then the customer purchases it and leaves the system. Otherwise, the customer transitions into another product with a certain transition probability and checks the availability of this next product. In this way, the customer transitions between the products until she reaches an available one. Therefore, the parameters of the Markov chain choice model are the probability that a customer arrives into the system to purchase each one of the products and the probability that a customer transitions from the current product to the next one when the current product is not available. We develop an expectation-maximization algorithm to estimate the parameters of the Markov chain choice model from the past purchase history of the customers.

The expectation-maximization algorithm dates back to Dempster et al. (1977) and it is useful for solving parameter estimation problems when the data available for estimation has a missing piece. In this algorithm, we start from some initial parameter estimates and iterate between the expectation and maximization steps. We focus on the so-called complete log-likelihood function, which is constructed under the assumption that we have access to the missing piece of the data. In the expectation step, we compute the expectation of the complete log-likelihood function, when the distribution of the missing piece of the data is driven by the current parameter estimates. In the maximization step, we maximize the expectation of the complete log-likelihood function to obtain new parameter estimates and repeat the process starting from the new parameter estimates. In parameter estimation problems, our goal is to find a maximizer of the so-called incomplete log-likelihood function, which is constructed under the assumption that we do not have access to the missing piece of the data. Dempster et al. (1977) show that the successive parameter estimates from the expectation-maximization algorithm monotonically improve the value of the incomplete log-likelihood function. Wu (1983) and Nettleton (1999) give regularity conditions to ensure convergence to a local maximum of the incomplete log-likelihood function.

Main Contributions. The data for estimating the parameters of the Markov chain choice model are the set of products made available to each customer and the product purchased by each customer. In our expectation-maximization algorithm, we treat the path that a customer follows in the Markov chain choice model as the missing piece of the data. In the expectation step, we need to compute two quantities, when the distribution of the missing piece of the data is driven by the current parameter estimates. The first quantity is the probability that the customer arrives into the system to purchase a certain product, conditional on the final purchase decision of the customer in the data. The second quantity is the expected number of times that a customer transitions from a certain product to another one, also conditional on the final purchase decision of the customer in the data. We show how to compute these quantities by solving linear systems of equations. Also, we show that the optimization problem in the maximization step has a closed-form solution.

We give a convergence result for our expectation-maximization algorithm. In particular, we show that the value of the incomplete log-likelihood function at the parameter estimates generated by our algorithm monotonically increases and converges to the value at a local maximizer, under the assumption that the parameters of the Markov chain choice model that we are trying to estimate are bounded away from zero. This assumption is arguably mild since we can put a small but strictly positive lower bound on the parameters with a negligible effect on the choice probabilities. In our computational experiments, we fit a Markov chain choice model to different data sets by using our expectation-maximization algorithm and compare the fitted Markov chain choice model with a benchmark that captures the choice process of the customers by using the multinomial logit model. The out-of-sample log-likelihoods of the Markov chain choice model can improve those of the benchmark by as much as 2%.

Related Literature. The Markov chain choice model was proposed by Blanchet et al. (2013). The authors study assortment problems, where there is a fixed revenue associated with each product and the goal is to find a set of products to offer to maximize the expected revenue from a customer. They give a polynomial-time algorithm to solve the assortment problem under the Markov chain choice model exactly. Feldman and Topaloglu (2014) focus on a deterministic approximation for network revenue management problems, where the decision variables correspond to the durations of time during which different subsets of products are offered to the customers. Thus, the number of decision variables increases exponentially with the number of products. The authors show that if the customers choose under the Markov chain choice model, then the number of decision variables in the deterministic approximation increases linearly with the number of products. Desir et al. (2015) study assortment problems under the Markov chain choice model with a constraint on the total space consumption of the offered products. They give a constant factor approximation algorithm and show that it is NP-hard to approximate the problem better than a fixed constant factor.

The expectation-maximization algorithm is used to estimate the parameters of various choice models. Vulcano et al. (2012) focus on estimating the parameters of the multinomial logit model when the demand is censored so that the customers who do not make a purchase are not recorded in the data. Following their work, we can also deal with demand censorship, as discussed in our conclusions section. Farias et al. (2013), van Ryzin and Vulcano (2015), van Ryzin and Vulcano (2016), Jagabathula and Vulcano (2016) and Jagabathula and Rusmevichientong (2016) consider the ranking-based choice model, where each customer has a ranked list of products in mind and she purchases the most preferred available product. The authors focus on estimating the parameters and coming up with ranked lists supported by the data. Chong et al. (2001), Kok and Fisher (2007), Misra (2008), Vulcano et al. (2010) and Dai et al. (2014) use real data to quantify the revenue improvements when one accounts for the customer choice process in assortment decisions.

Outline. In Section 2, we describe the Markov chain choice model. In Section 3, we provide the incomplete and complete likelihood functions. In Section 4, we give our expectation-maximization algorithm, show how to execute the expectation and maximization steps and discuss convergence. In Section 5, we give computational experiments. In Section 6, we conclude.

2 Markov Chain Choice Model

In the Markov chain choice model, we have n products indexed by $N = \{1, \ldots, n\}$. A customer arriving into the system is interested in purchasing product i with probability λ_i. If this product is available for purchase, then the customer purchases it and leaves the system. If this product is not available for purchase, then the customer transitions from product i to product j with probability $\rho_{i,j}$. The customer visits different products in this fashion until she visits a product that is available for purchase and purchases it. Naturally, we assume that $\sum_{i \in N} \lambda_i = 1$ and $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$. We note that a customer may have the option of leaving the system without a purchase in certain settings. To capture such a setting, we can assume that the option of leaving the system without a purchase corresponds to one of the products in N. This product is always available for purchase and if a customer visits this product, then she leaves the system without a purchase. Blanchet et al. (2013) and Feldman and Topaloglu (2014) show that we can solve a linear system of equations to compute the probability that a customer purchases a certain product when we offer the subset $S \subset N$ of products. In particular, we let $\Theta_i(S)$ be the expected number of times that a customer visits product i during the course of her choice process given that we offer the subset S of products. We can compute $\mathbf{\Theta}(S) = (\Theta_1(S), \ldots, \Theta_n(S))$ by solving

$$\Theta_i(S) = \lambda_i + \sum_{j \in N \setminus S} \rho_{j,i}\, \Theta_j(S) \qquad \forall\, i \in N. \tag{1}$$

We interpret (1) as follows. On the left side, $\Theta_i(S)$ is the expected number of times that a customer visits product i. The expected number of times that a customer visits product i when she arrives into the system is λ_i, yielding λ_i on the right side. The expected number of times that a customer visits some product $j \in N \setminus S$ is $\Theta_j(S)$. In each one of these visits, she transitions from product j to product i with probability $\rho_{j,i}$, yielding $\sum_{j \in N \setminus S} \rho_{j,i}\, \Theta_j(S)$ on the right side. If product i is available for purchase so that $i \in S$, then a customer can visit this product at most once, since she purchases this product whenever she visits it. So, the expected number of times that a customer visits product $i \in S$ is the same as the probability that a customer visits this product, which is, in turn, the same as the probability that a customer purchases this product. Thus, if $\mathbf{\Theta}(S) = (\Theta_1(S), \ldots, \Theta_n(S))$ is the solution to (1), then the probability that a customer purchases product $i \in S$ is $\Theta_i(S)$.

Using the vector $\bar{\boldsymbol{\lambda}} = \{\lambda_i : i \in N \setminus S\}$ and the matrix $\bar{\boldsymbol{\rho}} = \{\rho_{i,j} : i, j \in N \setminus S\}$, by (1), the vector $\bar{\mathbf{\Theta}}(S) = \{\Theta_i(S) : i \in N \setminus S\}$ satisfies $\bar{\mathbf{\Theta}}(S) = \bar{\boldsymbol{\lambda}} + \bar{\boldsymbol{\rho}}^\top \bar{\mathbf{\Theta}}(S)$. Thus, we have $(I - \bar{\boldsymbol{\rho}})^\top \bar{\Theta}(S) = \bar{\boldsymbol{\lambda}}$, where I is the identity matrix with the appropriate dimension. By Corollary C.4 in Puterman (1994), if we have $\sum_{j \in N \setminus S} \rho_{i,j} < 1$ for all $i \in N \setminus S$, then $(I - \bar{\boldsymbol{\rho}})^{-1}$ exists and has non-negative entries. So, since $((I - \bar{\boldsymbol{\rho}})^\top)^{-1} = ((I - \bar{\boldsymbol{\rho}})^{-1})^\top$, we obtain $\bar{\mathbf{\Theta}}(S) = ((I - \bar{\boldsymbol{\rho}})^{-1})^\top \bar{\boldsymbol{\lambda}}$. Once we compute $\bar{\mathbf{\Theta}}(S) = \{\Theta_i(S) : i \in N \setminus S\}$ by using the last equality, (1) implies that we can compute $\{\Theta_i(S) : i \in S\}$ as $\Theta_i(S) = \lambda_i + \sum_{j \in N \setminus S} \rho_{j,i}\, \Theta_j(S)$ for all $i \in S$. This discussion implies that if we have $\sum_{j \in N \setminus S} \rho_{i,j} < 1$ for all $i \in N \setminus S$, then there exists a unique value of $\mathbf{\Theta}(S)$ that satisfies (1) and this value of $\mathbf{\Theta}(S)$ has non-negative entries.

3 Incomplete and Complete Likelihood Functions

Our goal is to estimate the parameters $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ and $\boldsymbol{\rho} = \{\rho_{i,j} : i, j \in N\}$ of the Markov chain choice model by using the data on the subsets of products offered to the customers and the purchase decisions of the customers. We use the maximum likelihood method to estimate the parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. To capture the product that a customer purchases, we define the random variable $Z_i(S) \in \{0, 1\}$ such that $Z_i(S) = 1$ if and only if a customer purchases product i when we offer the subset S of products. Therefore, using $\boldsymbol{e}_i \in \Re^n$ to denote the unit vector with a one in the i-th component, we have $\boldsymbol{Z}(S) = (Z_1(S), \ldots, Z_n(S)) = \boldsymbol{e}_i$ with probability $\Theta_i(S)$. In the data that we have available to estimate the parameters of the Markov chain choice model, there are τ customers indexed by $T = \{1, \ldots, \tau\}$. We use $\hat{S}^t \subset N$ to denote the subset of products offered to customer t. To capture the product purchased by this customer, we define $\hat{Z}_i^t \in \{0, 1\}$ such that $\hat{Z}_i^t = 1$ if and only if customer t purchased product i, which implies that $\hat{\boldsymbol{Z}}^t = (\hat{Z}_1^t, \ldots, \hat{Z}_n^t)$ is a sample of the random variable $\boldsymbol{Z}(\hat{S}^t) = (Z_1(\hat{S}^t), \ldots, Z_n(\hat{S}^t))$. Thus, the data that is available to estimate the parameters of the Markov chain choice model is $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$. The probability that customer t purchases product i is $\Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$, where we explicitly show that the solution $\boldsymbol{\Theta}(S \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}) = (\Theta_1(S \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}), \ldots, \Theta_n(S \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}))$ to (1) depends on the parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ of the Markov chain choice model. In this case, the likelihood of the purchase decision of customer t is given by $\prod_{i \in N} \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})^{\hat{Z}_i^t}$, where we follow the convention that $0^0 = 1$. The log-likelihood of this purchase decision is $\sum_{i \in N} \hat{Z}_i^t \log \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$. Assuming that the purchase decisions of the different customers are independent of each other, the log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ is

$$L_I(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{Z}_i^t \log \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}). \tag{2}$$

To estimate the parameters of the Markov chain choice model, we can maximize $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ subject to the constraint that $\sum_{i \in N} \lambda_i = 1$, $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$, $\boldsymbol{\lambda} \in \Re_+^n$ and $\boldsymbol{\rho} \in \Re_+^{n \times n}$. The difficulty with this approach is that there is no closed-form expression for $\Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$ in the definition of $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$, which is the main motivation for our expectation-maximization algorithm.

In our expectation-maximization algorithm, we use a likelihood function constructed under the assumption that we have access to additional data for each customer. In particular, we define the random variable $F_i \in \{0, 1\}$ such that $F_i = 1$ if and only if a customer arriving into the system is interested in purchasing product i. Thus, we have $\boldsymbol{F} = (F_1, \ldots, F_n) = \boldsymbol{e}_i$ with probability λ_i. Also, we use the random variable $X_{i,j}(S)$ to denote the number of times that a customer transitions from product i to product j during the course of her choice process when we offer the subset S of products. We do not give the probability law for the random variable $X_{i,j}(S)$ explicitly, but we compute certain expectations involving this random variable in the next section. For each customer t, we assume that we have access to additional data so that we know the product that this customer was interested in purchasing when she arrived into the system, as well as the number of times that she transitioned from each product i to each product j during the course of her choice process. (This assumption is temporary to facilitate our analysis and our expectation-maximization algorithm will not require having access to the additional data.) In particular, we define $\hat{F}_i^t \in \{0,1\}$ such that $\hat{F}_i^t = 1$ if and only if customer t was interested in purchasing product i when she arrived into the system. We use $\hat{X}_{i,j}^t$ to denote the number of times that customer t transitioned from product i to product j during the course of her choice process. So, $\hat{\boldsymbol{F}}^t = (\hat{F}_1^t, \ldots, \hat{F}_n^t)$ and $\hat{\boldsymbol{X}}^t = \{\hat{X}_{i,j}^t : i,j \in N\}$ are respectively samples of the random variables $\boldsymbol{F} = (F_1, \ldots, F_n)$ and $\boldsymbol{X}(\hat{S}^t) = \{X_{i,j}(\hat{S}^t) : i,j \in N\}$. We construct a likelihood function under the assumption that the data that is available to estimate the parameters of the Markov chain choice model is $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t, \hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$. The probability that customer t is interested in purchasing product i when she arrives into the system is λ_i. Also, given that customer t is interested in purchasing product i up on arrival, the probability that she visits products $i, i_1, i_2, \ldots, i_{k-1}, i_k, j$ to purchase product j is given by $\rho_{i,i_1}\, \rho_{i_1,i_2} \ldots \rho_{i_{k-1},i_k}\, \rho_{i_k,j}$. In this case, the likelihood of the purchase decision of customer t is $\prod_{i \in N} \lambda_i^{\hat{F}_i^t} \prod_{i,j \in N} \rho_{i,j}^{\hat{X}_{i,j}^t}$. The log-likelihood of this purchase decision is $\sum_{i \in N} \hat{F}_i^t \log \lambda_i + \sum_{i,j \in N} \hat{X}_{i,j}^t \log \rho_{i,j}$. Thus, the log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t, \hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$ is

$$L_C(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^t \log \rho_{i,j}. \tag{3}$$

Note that once we know $\{(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$, $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ does not play a role in (3). To estimate the parameters of the Markov chain choice model, knowing $\{(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$ is equivalent to knowing the path that a customer follows in the Markov chain, since the second term on the right side of (3) does not depend on the order in which the transitions take place.

The likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (3) is constructed under the assumption that we have access to additional data for each customer. This likelihood function is known as the complete likelihood function and the subscript C in $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ stands for complete. In contrast, the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (2) is constructed under the assumption that we do not have access to the additional data for each customer. This likelihood function is known as the incomplete likelihood function and the subscript I in $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ stands for incomplete. Noting that $\log x$ is concave in x, the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is concave in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ and it has a closed-form expression. However, this likelihood function is not immediately useful when estimating the parameters of the Markov chain choice model, since we do not have access to the data $\{(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$ in practice. In the next section, we give an expectation-maximization algorithm that uses the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ to estimate the parameters of the Markov chain choice model, while making sure that we do not need to have access to the data $\{(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$.

4 Expectation-Maximization Algorithm

In this section, we describe our expectation-maximization algorithm. We show how to execute the expectation and maximization steps of this algorithm in detail. Lastly, we discuss the convergence properties of the iterates of our expectation-maximization algorithm.

4.1 Overview of the Algorithm

Our expectation-maximization algorithm estimates the parameters of the Markov chain choice model by using the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (3). Although this algorithm works with the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (3), it requires having access to the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$, but not to the data $\{(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$. In our expectation-maximization algorithm, we start with some estimate of the parameters $(\boldsymbol{\lambda}^1, \boldsymbol{\rho}^1)$ at the first iteration. At iteration ℓ, we estimate $\hat{F}_i^t$ as the expectation of the random variable F_i conditional on the fact that customer t chooses according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and her purchase decision is given by $\hat{\boldsymbol{Z}}^t = (\hat{Z}_1^t, \ldots, \hat{Z}_n^t)$. Similarly, we estimate $\hat{X}_{i,j}^t$ as the expectation of the random variable $X_{i,j}(\hat{S}^t)$ conditional on the fact that customer t chooses according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and her purchase decision is given by $\hat{\boldsymbol{Z}}^t = (\hat{Z}_1^t, \ldots, \hat{Z}_n^t)$. Computing these conditional expectations to estimate $\hat{F}_i^t$ and $\hat{X}_{i,j}^t$ for all $i, j \in N$, $t \in T$ is known as the expectation step. Next, we plug these estimates into (3) to construct the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ and maximize this likelihood function subject to the constraint that $\sum_{i \in N} \lambda_i = 1$, $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$, $\boldsymbol{\lambda} \in \Re_+^n$ and $\boldsymbol{\rho} \in \Re_+^{n \times n}$. The optimal solution to this problem yields parameters $(\boldsymbol{\lambda}^{\ell+1}, \boldsymbol{\rho}^{\ell+1})$ that we use at iteration $\ell + 1$. Maximizing the likelihood function $L_C(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in this fashion is known as the maximization step. Using the parameters $(\boldsymbol{\lambda}^{\ell+1}, \boldsymbol{\rho}^{\ell+1})$, we can go back to the expectation step to estimate $\hat{F}_i^t$ and $\hat{X}_{i,j}^t$ for all $i,j \in N$, $t \in T$. The expectation-maximization algorithm iteratively carries out the expectation and maximization steps to generate a sequence of parameters $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$. We state the expectation-maximization algorithm below.

Step 1. Choose the initial estimates $(\boldsymbol{\lambda}^1, \boldsymbol{\rho}^1)$ of the parameters of the Markov chain choice model arbitrarily, as long as they satisfy $\sum_{i \in N} \lambda_i^1 = 1$, $\sum_{j \in N} \rho_{i,j}^1 = 1$ for all $i \in N$, $\boldsymbol{\lambda}^1 \in \Re_+^n$ and $\boldsymbol{\rho}^1 \in \Re_+^{n \times n}$. Initialize the iteration counter by setting $\ell = 1$.

Step 2. (Expectation) Assuming that the customers choose according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$, set $\hat{F}_i^{t,\ell} = \mathbb{E}\{F_i \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$ and $\hat{X}_{i,j}^{t,\ell} = \mathbb{E}\{X_{i,j}(\hat{S}_t) \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$ for all $i,j \in N$, $t \in T$.

Step 3. (Maximization) Let $(\boldsymbol{\lambda}^{\ell+1}, \boldsymbol{\rho}^{\ell+1})$ be the maximizer of $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j}$ subject to the constraint that $\sum_{i \in N} \lambda_i = 1$, $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$, $\boldsymbol{\lambda} \in \Re_+^n$ and $\boldsymbol{\rho} \in \Re_+^{n \times n}$. Increase ℓ by one and go to Step 2.

In the expectation step, we compute non-trivial conditional expectations. In Section 4.2, we show that we can compute these conditional expectations by solving linear systems of equations. In the maximization step, we maximize the function $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho})$ subject to linear constraints. In Section 4.3, we show that this optimization problem has a closed-form solution. To estimate the parameters of the Markov chain choice model, we need to maximize the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$. In Section 4.4, we consider our expectation-maximization algorithm under the assumption that the parameters that we are trying to estimate are bounded away from zero. We show that the value of the likelihood function $\{L_I(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ at the successive parameter estimates $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ generated by our algorithm monotonically increases and converges to the value of the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ at a local maximizer. In that section, we precisely define what we mean by a local maximizer. We also show that we can still solve the optimization problem in the maximization step in polynomial time when we have a strictly positive lower bound on the parameters.

4.2 Expectation Step

In the expectation step of the expectation-maximization algorithm, customer t is offered the subset $\hat{S}^t$ of products. She chooses among these products according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. We know that the purchase decision of this customer is given by the vector $\hat{\boldsymbol{Z}}^t$, which is to say that if we know that customer t purchased product k, then $\hat{\boldsymbol{Z}}^t = \boldsymbol{e}_k$. We need to compute the conditional expectations $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$ and $\mathbb{E}\{X_{i,j}(\hat{S}_t) \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$. In this section, we show that we can compute these conditional expectations by solving linear systems of equations. For notational brevity, we omit the superscript t indexing the customer and the superscript ℓ indexing the iteration counter. In particular, we consider the case where a customer is offered the subset S of products. She chooses among these products according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. We know that the customer purchased some product k. In other words, noting that the purchase decision of a customer is captured by the random variable $\boldsymbol{Z}(S) = (Z_1(S), \ldots, Z_n(S))$, we know that $\boldsymbol{Z}(S) = \boldsymbol{e}_k$. We want to compute the conditional expectations $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$.

Computation of $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$. The expectation $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ is conditional on having $\boldsymbol{Z}(S) = \boldsymbol{e}_k$. Thus, we know that a customer purchased product k out of the subset S of products, which implies that we must have $k \in S$. Therefore, we assume that $k \in S$ in our discussion. Using the Bayes rule, we have

$$\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} = \mathbb{P}\{F_i = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} = \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}\, \mathbb{P}\{F_i = 1\}}{\mathbb{P}\{Z_k(S) = 1\}}. \tag{4}$$

On the right side of (4), $\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}$ is the probability that a customer purchases product k out of the subset S of products given that she is interested in purchasing product i when she arrives into the system. This probability is simple to compute when $i \in S$. In particular, if product i is offered and a customer is interested in purchasing product i when she arrives into the system, then this customer definitely purchases product i. Thus, letting $\mathbf{1}(\cdot)$ be the indicator function, we have $\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\} = \mathbf{1}(i = k)$ for all $i \in S$. We focus on computing $\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}$ for all $i \in N \setminus S$. Letting $\Psi_k(i, S) = \mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}$ for all $i \in N \setminus S$ for notational brevity, we can compute $\{\Psi_k(i, S) : i \in N \setminus S\}$ by solving the linear system of equations

$$\Psi_k(i, S) = \rho_{i,k} + \sum_{j \in N \setminus S} \rho_{i,j}\, \Psi_k(j, S) \qquad \forall\, i \in N \setminus S. \tag{5}$$

We interpret (5) as follows. On the left side, $\Psi_k(i, S)$ is the probability that a customer purchases product k out of the subset S of products given that she is interested in purchasing product i up on arrival. For this customer to purchase product k, she may transition from product i to product k, yielding $\rho_{i,k}$ on the right side. Alternatively, the customer may transition from product i to some product $j \in N \setminus S$, at which point, she is identical to a customer interested in purchasing product j up on arrival and this customer purchases product k with probability $\Psi_k(j, S)$. This reasoning yields $\sum_{j \in N \setminus S} \rho_{i,j}\, \Psi_k(j, S)$ on the right side. By the same discussion at the end of Section 2, if $\sum_{j \in N \setminus S} \rho_{i,j} < 1$ for all $i \in N \setminus S$, then there exists a unique solution to the system of equations in (5). Thus, if $\{\Psi_k(i, S) : i \in N \setminus S\}$ solve (5), then we have $\Psi_k(i, S) = \mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}$ for all $i \in N \setminus S$. For notational uniformity, noting the discussion right before (5), we let $\Psi_k(i, S) = \mathbf{1}(i = k)$ for all $i \in S$. In this case, we have $\Psi_k(i, S) = \mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}$ for all $i \in N$.

The other probabilities on the right side of (4) are simple to compute. Noting that $\mathbb{P}\{F_i = 1\}$ is the probability that a customer arriving into the system is interested in purchasing product i, we have $\mathbb{P}\{F_i = 1\} = \lambda_i$. Similarly, since $\mathbb{P}\{Z_k(S) = 1\}$ is the probability that a customer purchases product k out of the subset S of products, we have $\mathbb{P}\{Z_k(S) = 1\} = \Theta_k(S)$, where $(\Theta_1(S), \ldots, \Theta_n(S))$ solve the system of equations in (1). Putting the discussion so far together, we compute $\{\Psi_k(i, S) : i \in N \setminus S\}$ by solving the system of equations in (5). Also, letting $\Psi_k(i, S) = \mathbf{1}(i = k)$ for all $i \in S$, by (4), for all $i \in N$, we have

$$\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} = \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}\, \mathbb{P}\{F_i = 1\}}{\mathbb{P}\{Z_k(S) = 1\}} = \frac{\Psi_k(i, S)\, \lambda_i}{\Theta_k(S)}. \tag{6}$$

When we offer the subset S of products, $\Psi_k(i, S)$ is the purchase probability of product k conditional on the fact that the customer is interested in purchasing product i when she arrives into the system, which can be computed by solving (5), whereas $\Theta_k(S)$ is the unconditional purchase probability of product k, which can be computed by solving (1). The systems of equations in (1) and (5) are similar to each other. Blanchet et al. (2013) and Feldman and Topaloglu (2014) use (1) to compute unconditional purchase probabilities, but it is interesting that a slight variation of (1) allows computing conditional purchase probabilities.

Computation of $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$. Similar to our earlier argument, since the expectation $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ is conditional on having $\boldsymbol{Z}(S) = \boldsymbol{e}_k$, we know that a customer purchased product k out of the subset S of products, which implies that $k \in S$. Therefore, we assume that $k \in S$ in our discussion. The random variable $X_{i,j}(S)$ captures the number of times that a customer transitions from product i to product j during the course of her choice process when we offer the subset S of products. If we have $i \in S$, then a customer cannot transition from product i to another product, since the customer purchases product i whenever she visits it. So,

$X_{i,j}(S) = 0$ with probability one for all $i \in S$, which implies that $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} = 0$. We focus on computing the expectation $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ for all $i \in N \setminus S$. We define the random variable $Y_i^m(S) \in \{0, 1\}$ such that $Y_i^m(S) = 1$ if and only if the m-th product that a customer visits during the course of her choice process is product i when we offer the subset S of products. Since $X_{i,j}(S)$ is the number of times a customer transitions from product i to product j when we offer the subset S of products, we have $X_{i,j}(S) = \sum_{m=1}^{\infty} \mathbf{1}(Y_i^m(S) = 1, Y_j^{m+1}(S) = 1)$. So, we have

$$\begin{aligned}\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} &= \sum_{m=1}^{\infty} \mathbb{P}\{Y_i^m(S) = 1, Y_j^{m+1}(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}. \\ &= \sum_{m=1}^{\infty} \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} \, \mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}, \quad (7)\end{aligned}$$

where the second equality is by the Bayes rule. We focus on each one of the probabilities $\mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ on the right side of (7) separately. Considering the probability $\mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$, from the perspective of the final purchase decision, a customer that visits product i as the m-th product is indistinguishable from a customer that visits product i as the first product. Thus, we have $\mathbb{P}\{\boldsymbol{Z}(S) = \boldsymbol{e}_k \,|\, Y_i^m(S) = 1\} = \mathbb{P}\{\boldsymbol{Z}(S) = \boldsymbol{e}_k \,|\, F_i = 1\}$. In this case, by using the Bayes rule once more, we have

$$\begin{aligned}\mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} &= \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, Y_i^m(S) = 1\} \, \mathbb{P}\{Y_i^m(S) = 1\}}{\mathbb{P}\{Z_k(S) = 1\}} \\ &= \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\} \, \mathbb{P}\{Y_i^m(S) = 1\}}{\mathbb{P}\{Z_k(S) = 1\}} = \frac{\Psi_k(i, S) \, \mathbb{P}\{Y_i^m(S) = 1\}}{\Theta_k(S)}, \quad (8)\end{aligned}$$

where we compute $\{\Psi_k(i, S) : i \in N \setminus S\}$ by solving (5). On the other hand, considering the probability $\mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$, by the Bayes rule, we also have

$$\begin{aligned}&\mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k, Y_i^m(S) = 1\} \\ &\quad = \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, Y_j^{m+1}(S) = 1, Y_i^m(S) = 1\} \, \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1\}}{\mathbb{P}\{Z_k(S) = 1 \,|\, Y_i^m(S) = 1\}} \\ &\quad = \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, Y_j^{m+1}(S) = 1\} \, \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1\}}{\mathbb{P}\{Z_k(S) = 1 \,|\, Y_i^m(S) = 1\}} \\ &\quad = \frac{\mathbb{P}\{Z_k(S) = 1 \,|\, F_j = 1\} \, \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1\}}{\mathbb{P}\{Z_k(S) = 1 \,|\, F_i = 1\}} = \frac{\Psi_k(j, S) \, \rho_{i,j}}{\Psi_k(i, S)}. \quad (9)\end{aligned}$$

In the chain of equalities above, the second equality uses the fact that if we know the $(m+1)$-st product that a customer visits, then the distribution of the product that she purchases does not depend on the m-th product that this customer visits. The third equality uses the fact that a customer that visits product j as the $(m+1)$-st product is indistinguishable from a customer that visits product j as the first product from the perspective of the final purchase decision. The fourth equality is by the fact that given that a customer visits product i as the m-th product, the probability that she visits product j next is given by the transition probability $\rho_{i,j}$. To compute the conditional expectation $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$, we use (8) and (9) in (7) to get

$$\begin{aligned}
\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} &= \sum_{m=1}^{\infty} \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} \, \mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\} \\
&= \sum_{m=1}^{\infty} \frac{\Psi_k(j,S)\,\rho_{i,j}}{\Psi_k(i,S)} \, \frac{\Psi_k(i,S)\,\mathbb{P}\{Y_i^m(S) = 1\}}{\Theta_k(S)} \\
&= \frac{\Psi_k(j,S)\,\rho_{i,j}}{\Theta_k(S)} \sum_{m=1}^{\infty} \mathbb{P}\{Y_i^m(S) = 1\} = \frac{\Psi_k(j,S)\,\rho_{i,j}}{\Theta_k(S)} \, \Theta_i(S).
\end{aligned} \tag{10}$$

In the last equality above, we use the fact that $\sum_{m=1}^{\infty} \mathbb{P}\{Y_i^m(S) = 1\}$ corresponds to the expected number of times that a customer visits product i given that we offer the subset S of products, in which case, by the discussion in Section 2, this quantity is given by $\Theta_i(S)$. The discussion in this section shows how we can compute the conditional expectations $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$. The main bulk of the work involves solving the systems of equations in (1) and (5) to obtain $(\Theta_1(S), \ldots, \Theta_n(S))$ and $\{\Psi_k(i,S) : i \in N \setminus S\}$.

Using (6) and (10), we can give explicit expressions to execute the expectation step of our expectation-maximization algorithm. We replace $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ in (1) and solve this system of equations. We use $(\Theta_1^\ell(\hat{S}^t), \ldots, \Theta_n^\ell(\hat{S}^t))$ to denote the solution. Similarly, we replace $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ in (5) and solve this system of equations. We use $\{\Psi_k^\ell(i, \hat{S}^t) : i \in N \setminus S\}$ to denote the solution. Also, we let $\Psi_k^\ell(i, \hat{S}^t) = \mathbf{1}(i = k)$ for all $i \in S$. In this case, by (6), for all $i \in N$ and $t \in T$, we have $\hat{F}_i^{t,\ell} = \mathbb{E}\{F_i \,|\, \boldsymbol{Z}(\hat{S}^t) = \boldsymbol{e}_k\} = \Psi_k^\ell(i, \hat{S}^t)\,\lambda_i^\ell \,/\, \Theta_k^\ell(\hat{S}^t)$. Also, by (10), for all $i \in N \setminus S$, $j \in N$ and $t \in T$, we have $\hat{X}_{i,j}^{t,\ell} = \mathbb{E}\{X_{i,j}(\hat{S}^t) \,|\, \boldsymbol{Z}(\hat{S}^t) = \boldsymbol{e}_k\} = \Psi_k^\ell(j, \hat{S}^t)\,\rho_{i,j}^\ell\,\Theta_i^\ell(\hat{S}^t) \,/\, \Theta_k^\ell(\hat{S}^t)$. Finally, we set $\hat{X}_{i,j}^{t,\ell} = 0$ for all $i \in S$, $j \in N$ and $t \in T$.

4.3 Maximization Step

In the maximization step of the expectation-maximization algorithm we need to maximize the function $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j}$ subject to the constraint that $\sum_{i \in N} \lambda_i = 1$, $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$, $\boldsymbol{\lambda} \in \Re_+^n$ and $\boldsymbol{\rho} \in \Re_+^{n \times n}$. This optimization problem decomposes into $1 + n$ problems given by

$$\max \left\{ \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i : \sum_{i \in N} \lambda_i = 1,\ \lambda_i \geq 0\ \forall\, i \in N \right\} \qquad \text{and} \tag{11}$$

$$\max \left\{ \sum_{t \in T} \sum_{j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j} : \sum_{j \in N} \rho_{i,j} = 1,\ \rho_{i,j} \geq 0\ \forall\, j \in N \right\} \quad \forall\, i \in N. \tag{12}$$

Problem (11) corresponds to the problem of computing the maximum likelihood estimators of the parameters $(\lambda_1, \ldots, \lambda_n)$ of the multinomial distribution, where λ_i is the probability of observing outcome i in each trial, we have a total of $\sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell}$ trials and we observe outcome i in $\sum_{t \in T} \hat{F}_i^{t,\ell}$ trials. In this case, the maximum likelihood estimator of λ_i is known to be $\sum_{t \in T} \hat{F}_i^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{F}_j^{t,\ell}$; see Section 2.2 in Bishop (2006). Therefore, the optimal solution to problem (11) is obtained by setting $\lambda_i = \sum_{t \in T} \hat{F}_i^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{F}_j^{t,\ell}$ for all $i \in N$. In the next section, we focus on our expectation-maximization algorithm under the assumption that the parameters of the Markov chain choice model are known to be bounded away from zero by some $\epsilon > 0$. In this case, the maximization step requires solving the first problem above with a lower bound of ϵ on the decision variables $(\lambda_1, \ldots, \lambda_n)$. In Online Appendix A, we discuss how to solve this optimization problem. Repeating this discussion with $\epsilon = 0$ also shows that we can obtain the optimal solution to problem (11) by setting $\lambda_i = \sum_{t \in T} \hat{F}_i^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{F}_j^{t,\ell}$ for all $i \in N$. Each one of the n problems in (12) has the same structure as problem (11). Following the same argument used to find the optimal value of λ_i, the optimal solution to each one of the n problems in (12) is obtained by setting $\rho_{i,j} = \sum_{t \in T} \hat{X}_{i,j}^{t,\ell} / \sum_{t \in T} \sum_{k \in N} \hat{X}_{i,k}^{t,\ell}$ for all $j \in N$.

Thus, to execute the maximization step of our expectation-maximization algorithm, we simply set $\lambda_i^{\ell+1} = \sum_{t \in T} \hat{F}_i^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{F}_j^{t,\ell}$ and $\rho_{i,j}^{\ell+1} = \sum_{t \in T} \hat{X}_{i,j}^{t,\ell} / \sum_{t \in T} \sum_{k \in N} \hat{X}_{i,k}^{t,\ell}$ for all $i, j \in N$.

4.4 Convergence of the Algorithm

We can give a convergence result for our expectation-maximization algorithm under the assumption that the parameters of the Markov chain choice model that we are trying to estimate are known to be bounded away from zero by some $\epsilon > 0$. Under this assumption, we execute the maximization step of the expectation-maximization algorithm slightly differently. In particular, we let $(\boldsymbol{\lambda}^{\ell+1}, \boldsymbol{\rho}^{\ell+1})$ be the maximizer of $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j}$ subject to the constraint that $\sum_{i\in N}\lambda_i = 1$, $\sum_{j\in N}\rho_{i,j} = 1$ for all $i \in N$, $\lambda_i \geq \epsilon$ for all $i \in N$ and $\rho_{i,j} \geq \epsilon$ for all $i, j \in N$. In other words, we impose a lower bound of ϵ on the decision variables. In Online Appendix A, we show that we can still solve the last optimization problem in polynomial time. The assumption that the parameters of the Markov chain choice model are known to be bounded away from zero by some $\epsilon > 0$ allows us to satisfy certain regularity conditions when we study the convergence of our expectation-maximization algorithm. This assumption is arguably mild, since we can put a small lower bound of $\epsilon > 0$ on the parameters with negligible effect on the choice probabilities. To give a convergence result for our expectation-maximization algorithm, we let $\Omega = \{(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Re^n_+ \times \Re^{n\times n}_+ : \sum_{i\in N}\lambda_i = 1, \sum_{j\in N}\rho_{i,j} = 1\ \forall\, i \in N,\ \lambda_i \geq \epsilon\ \forall\, i \in N,\ \rho_{i,j} \geq \epsilon\ \forall\, i, j \in N\}$, capturing the set of possible parameter values when we have a lower bound of ϵ on the parameters. Also, we define the set of parameters

$$\Phi = \left\{ (\boldsymbol{\lambda}_0, \boldsymbol{\rho}_0) \in \Omega : \frac{dL_I((1-\gamma)\,(\boldsymbol{\lambda}_0, \boldsymbol{\rho}_0) + \gamma\,(\boldsymbol{\lambda}, \boldsymbol{\rho}))}{d\gamma}\Big|_{\gamma=0} \leq 0 \ \ \forall\,(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Omega \right\}.$$

Roughly speaking, having $(\boldsymbol{\lambda}_0, \boldsymbol{\rho}_0) \in \Phi$ implies that if we start from the point $(\boldsymbol{\lambda}_0, \boldsymbol{\rho}_0)$ and move towards any point $(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Omega$ for an infinitesimal step size, then the value of the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ does not improve. In the next theorem, we give a convergence result for our expectation-maximization algorithm when we know that the parameters that we are trying to estimate are bounded away from zero by some $\epsilon > 0$. The proof is in Online Appendix B.

Theorem 1 *Assume that the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ is generated by our expectation-maximization algorithm when we impose a lower bound of $\epsilon > 0$ on the parameters of the Markov chain choice model. Then, we have $L_I(\boldsymbol{\lambda}^{\ell+1}, \boldsymbol{\rho}^{\ell+1}) \geq L_I(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ for all $\ell = 1, 2, \ldots$. Furthermore, all limit points of the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ are in Φ and the sequence $\{L_I(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ converges to $L_I(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\rho}})$ for some $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\rho}}) \in \Phi$.*

Note that $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is the function that we need to maximize to estimate the parameters. By the theorem above, the sequence of parameters generated by our algorithm monotonically improves $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ and we have convergence to a some form of local maximum of $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$. Since $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is not necessarily concave, we are not guaranteed to get to the global maximum. Nettleton (1999) gives regularity conditions to ensure convergence of the expectation-maximization algorithm. The proof of Theorem 1 follows by verifying these regularity conditions. In Online Appendix C, we give an example to show that the regularity conditions in Nettleton (1999) may not hold without a lower bound of $\epsilon > 0$ on the parameters. Wu (1983) gives other regularity conditions but he assumes that the parameters generated by the algorithm are in the interior of the set of possible parameter values, which is difficult to satisfy for the Markov chain choice model. Also, Theorem 1 does not rule out the possibility of multiple limit points for the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$, but all limit points are in Φ. Lastly, we are not able to give a convergence result without a lower bound of $\epsilon > 0$ on the parameters. In Online Appendix D, however, we show that as long as the initial parameter estimates are strictly positive, even if we do not impose a lower bound on the parameters, our algorithm always generates a sequence of parameters $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ such that there exist unique solutions to the systems of equations in (1) and (5) when we solve these systems of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ for any subset in the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$. So, we do not encounter parameters that render the systems of equations in (1) or (5) unsolvable.

5 Computational Experiments

We test the performance of our expectation-maximization algorithm on randomly generated data, as well as on a data set coming from a hotel revenue management application.

5.1 Benchmark Strategies

In our first benchmark, referred to as EM, we estimate the parameters of the Markov chain choice model by using our expectation-maximization algorithm. In our second benchmark, referred to as DM, we continue using the Markov chain choice model to capture the customer choices, but we estimate the parameters by directly maximizing the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (2) through continuous optimization software. In our third benchmark, referred to as ML, we use the multinomial logit model to capture the customer choices and estimate its parameters by using maximum likelihood. We briefly describe the multinomial logit model. In the multinomial logit model, the mean utility of product i is η_i. If we offer the subset S of products, then a customer purchases product i with probability $e^{\eta_i} / \sum_{j \in S} e^{\eta_j}$. As mentioned in Section 2, we represent the no-purchase option as a product always available for purchase. We denote this product by $\phi \in N$. In other words, a customer purchasing product ϕ corresponds to a customer leaving the system without a purchase. If we add the same constant to the mean utilities of all products, then the choice probability $e^{\eta_i} / \sum_{j \in S} e^{\eta_j}$ of each product i does not change. So, we normalize the mean utility of the no-purchase option to zero. In this case, the parameters of the multinomial logit model are $\boldsymbol{\eta} = \{\eta_i : i \in N \setminus \{\phi\}\}$. Assume that we offer the subset $\hat{S}^t$ of products to customer t and the purchase decision of this customer is given by $\hat{\boldsymbol{Z}}^t = (\hat{Z}_1^t, \ldots, \hat{Z}_n^t)$, where $\hat{Z}_i^t = 1$ if and only if the customer purchases product i. The likelihood of the purchase decision of customer t is $\prod_{i \in N} (e^{\eta_i} / \sum_{j \in \hat{S}_t} e^{\eta_j})^{\hat{Z}_i^t}$. Noting that $\sum_{i \in N} \hat{Z}_i^t = 1$, the log-likelihood of this purchase decision is $\sum_{i \in N} \hat{Z}_i^t \, \eta_i - \sum_{i \in N} \hat{Z}_i^t \, \log(\sum_{j \in \hat{S}_t} e^{\eta_j}) = \sum_{i \in N} \hat{Z}_i^t \, \eta_i - \log(\sum_{i \in \hat{S}_t} e^{\eta_i})$; see Section II in McFadden (1974). Thus, the log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ is given by

$$L(\boldsymbol{\eta}) = \sum_{t \in T} \sum_{i \in N} \hat{Z}_i^t \, \eta_i - \sum_{t \in T} \log \Big(\sum_{i \in \hat{S}_t} e^{\eta_j} \Big). \tag{13}$$

In ML, we estimate the parameters of the multinomial logit model by maximizing the log-likelihood function in (13) through the `Matlab` routine `fmincon`. Section 3.1.5 in Boyd and Vandenberghe (2005) shows that $\log(\sum_{i=1}^n e^{x_i})$ is convex in $(x_1, \ldots, x_n) \in \Re^n$. Thus, the log-likelihood function $L(\boldsymbol{\eta})$ in (13) is concave in $\boldsymbol{\eta}$. Letting $w_i = e^{\eta_i}$, we can also express the choice probability of product i out of the subset S as $w_i / \sum_{j \in S} w_j$, but the log-likelihood function $L(\boldsymbol{w}) = \sum_{t \in T} \sum_{i \in N} \hat{Z}_i^t \, \log w_i - \sum_{t \in T} \log(\sum_{i \in \hat{S}_t} w_i)$ is not concave in $\boldsymbol{w} = \{w_i : i \in N \setminus \{\phi\}\}$.

In DM, we directly maximize the log-likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ in (2) also by using the `Matlab` routine `fmincon`. Since the function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is not necessarily concave in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$, the parameters estimated by DM may depend on the initial solution, but exploratory trials indicated that the performance of DM is rather insensitive to the initial solution. We use the initial solution $\lambda_i = 1/n$ for all $i \in N$, $\rho_{i,j} = 1/n$ for all $i \in N \setminus \{\phi\}$, $j \in N$ and $\rho_{\phi,\phi} = 1$. In EM, we use this initial solution as well. In our expectation-maximization algorithm, we do not impose a strictly positive lower bound on the parameters and stop when the incomplete log-likelihood increases by less than 0.01% in two successive iterations. We give the pseudo-code for our algorithm in Online Appendix E. We also used the so-called independent demand model as a benchmark. This model performed consistently worse than EM, DM and ML and we will comment on its performance only briefly.

5.2 Known Ground Choice Model

We provide computational experiments on randomly generated data where we have access to the exact ground choice model that governs the customer choice process.

Experimental Setup. We assume that the ground choice model that governs the customer choice process is the ranking-based choice model. In this choice model, each arriving customer has a ranked list of products in mind and she purchases the most preferred available product in her ranked list. The ranking-based choice model is used in Mahajan and van Ryzin (2001), van Ryzin and Vulcano (2008), Smith et al. (2009), Honhon et al. (2010), Honhon et al. (2012), Farias et al. (2013), van Ryzin and Vulcano (2015), Jagabathula and Vulcano (2016), Jagabathula and Rusmevichientong (2016) and van Ryzin and Vulcano (2016). We have m possible ranked lists indexed by $M = \{1, \ldots, m\}$. We denote the possible ranked lists that an arriving customer can have in mind by using $\{(\sigma_1^g, \ldots, \sigma_n^g) : g \in M\}$, where σ_i^g is the preference order of product i in the ranked list $\boldsymbol{\sigma}^g = (\sigma_1^g, \ldots, \sigma_n^g)$. For example, if $n = 3$ and $\sigma_1^g = 2$, $\sigma_2^g = 3$, $\sigma_3^g = 1$, then product 3 is the most preferred product, product 1 is the second most preferred product and product 2 is the third most preferred product. The probability that an arriving customer has the ranked list $\boldsymbol{\sigma}^g$ in mind is β^g. If we offer the subset S of products, then an arriving customer purchases product i with probability $\sum_{g \in M} \beta^g \, \mathbf{1}(i = \arg\min_{j \in S} \sigma_j^g)$, which is the probability that an arriving customer has product i as her most preferred available product. In our computational experiments, to come up with the possible ranked lists that a customer can have in mind, we generate m random permutations of the products. To come up with the probability β^g that an arriving customer has the ranked list $\boldsymbol{\sigma}^g$ in mind, following van Ryzin and Vulcano (2016), we generate γ^g from the uniform distribution over $[0, 1]$ and set $\beta^g = \gamma^g / \sum_{h \in M} \gamma^h$. We note that $(\beta^1, \ldots, \beta^m)$ generated in this fashion is not uniformly distributed over the $(m-1)$-dimensional simplex.

Once we generate the ground choice model that governs the customer choice process, we generate the past purchase histories $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ of the customers from this ground choice model. To come up with the subset $\hat{S}^t$ of products offered to customer t, we assume that the no-purchase option is always available in the offered subset of products. Each of the other products are included in the subset $\hat{S}^t$ with probability 1/2. Once we come up with the subset of products offered to customer t, we generate the choice of a random customer out of this subset according to the ground choice model and set $\hat{Z}_i^t = 1$ if and only if customer t purchases product i. Using this approach, we generate the purchase history of 50,000 customers to use as the training data and a separate purchase history of 10,000 customers to use as the hold-out data. We use different portions of the generated training data with 2,500, 5,000, 10,000 and 50,000 customers to fit our choice models. In our test problems, the number of products is $n = 11$ or $n = 21$. One of these products corresponds to the no-purchase option. The number of possible ranked lists is $m = 10 + n$, $m = 20 + n$ or $m = 40 + n$. For each product $i \in N$, there is one ranked list where the most preferred product in the ranked list is

(n,m)	Trn. data	Out-of-sample log-likelihood EM	DM	ML	EM-DM	EM-ML	(n,m)	Trn. data	Out-of-sample log-likelihood EM	DM	ML	EM-DM	EM-ML
$(11,21)$	2,500	-16,391	-16,407	-16,677	0.10%	1.72%	$(21,31)$	2,500	-22,839	-23,053	-23,139	0.93%	1.30%
$(11,21)$	5,000	-16,387	-16,390	-16,665	0.02%	1.67%	$(21,31)$	5,000	-22,737	-22,740	-23,116	0.02%	1.64%
$(11,21)$	10,000	-16,366	-16,349	-16,663	-0.10%	1.78%	$(21,31)$	10,000	-22,673	-22,648	-23,100	-0.11%	1.85%
$(11,21)$	50,000	-16,361	-16,327	-16,663	-0.21%	1.81%	$(21,31)$	50,000	-22,627	-22,583	-23,088	-0.19%	2.00%
$(11,31)$	2,500	-16,199	-16,222	-16,298	0.14%	0.61%	$(21,41)$	2,500	-22,581	-22,783	-22,807	0.88%	0.99%
$(11,31)$	5,000	-16,154	-16,161	-16,286	0.04%	0.81%	$(21,41)$	5,000	-22,483	-22,580	-22,781	0.43%	1.31%
$(11,31)$	10,000	-16,141	-16,138	-16,282	-0.02%	0.87%	$(21,41)$	10,000	-22,422	-22,408	-22,773	-0.06%	1.54%
$(11,31)$	50,000	-16,132	-16,120	-16,283	-0.07%	0.93%	$(21,41)$	50,000	-22,378	-22,331	-22,763	-0.21%	1.69%
$(11,51)$	2,500	-17,029	-17,059	-17,116	0.17%	0.51%	$(21,61)$	2,500	-23,406	-23,562	-23,439	0.66%	0.14%
$(11,51)$	5,000	-17,017	-17,032	-17,111	0.09%	0.55%	$(21,61)$	5,000	-23,301	-23,321	-23,426	0.09%	0.53%
$(11,51)$	10,000	-16,992	-16,993	-17,107	0.00%	0.67%	$(21,61)$	10,000	-23,267	-23,271	-23,406	0.02%	0.59%
$(11,51)$	50,000	-16,986	-16,966	-17,104	-0.12%	0.69%	$(21,61)$	50,000	-23,226	-23,183	-23,399	-0.19%	0.74%

Table 1: Out-of-sample log-likelihoods obtained by EM, DM and ML.

product i. In this way, we fix the most preferred product in n of the ranked lists. The preference order of the other products in these n ranked lists are randomly generated. The remaining ranked lists other than these n ranked lists are fully random permutations of the products.

RESULTS. In Table 1, we show the out-of-sample log-likelihoods obtained by EM, DM and ML. The first column in this table shows the parameter combination (n, m) in the ground choice model. The second column shows the number of customers in the training data that we use to fit our choice models. The third column shows the out-of-sample log-likelihood obtained by EM, which is the value of the log-likelihood function in (2) after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ by the parameters estimated by EM and using the hold-out data as the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ in this log-likelihood. The fourth column shows the out-of-sample log-likelihood obtained by DM, which is the value of the same log-likelihood function in the second column, but the parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ correspond to the those estimated by DM. The fifth column shows the out-of-sample log-likelihood obtained by ML, which is the value of the log-likelihood function in (13) after replacing $\boldsymbol{\eta}$ by the parameters estimated by ML and using the hold-out data as the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ in this log-likelihood. The last two columns compare the log-likelihood obtained by EM with those obtained by DM and ML by giving the percent gaps between the corresponding pairs of log-likelihoods. The results indicate that EM provides better out-of-sample log-likelihoods than ML in our test problems. The out-of-sample log-likelihoods obtained by EM and DM are quite close.

When we have 2,500 customers in the training data and 11 products, the average computation times for EM, DM and ML are respectively 12.93, 441.04 and 0.54 seconds on a 2.2 GHz Intel Core i7 CPU with 16 GB RAM. With 50,000 customers and 21 products, the average computation times for EM, DM and ML are respectively 3,456.19, 218,012.22 and 16.99 seconds. EM terminates in 31 to 52 iterations. The computation times for EM and ML are reasonable since we do not solve the estimation problem in real time, but DM is computationally demanding. The computation time for EM is mostly spent on solving the systems of equations in (1) and (5) for each subset in the training data. Thus, EM is drastically faster when the customers in the training data are offered a few different subsets, which is likely to happen in practice. For example, when we have 50,000 customers in the training data, if these customers are offered one of 10 different subsets, then EM takes about 20 seconds. In Online Appendix F, we give the detailed computation times.

EM improves the log-likelihoods obtained by ML, but ML may provide advantages in certain cases. EM estimates $O(n^2)$ parameters given by $\{\lambda_i : i \in N\}$ and $\{\rho_{i,j} : i, j \in N\}$, whereas ML estimates $O(n)$ parameters given by $\{\eta_i : i \in N \setminus \{\phi\}\}$. The Markov chain choice model is more flexible due to its larger number of parameters. However, since the Markov chain choice model has a large number of parameters, EM may over-fit this choice model to the training data, especially when we have too few customers in the training data and too many products so that we need to estimate too many parameters from too little data. In this case, the out-of-sample performance of EM may be inferior. For example, if we have 1,000 customers in the training data and 21 products, so that EM estimates about 400 parameters from 1,000 data points, then the average percent gap between the out-of-sample log-likelihoods obtained by EM and ML is -0.45%, favoring ML, where the average is computed over the test problems with $m \in \{10 + n, 20 + n, 40 + n\}$. Clearly, it is difficult to estimate 400 parameters from 1,000 data points! If we have 1,000 customers in the training data and 11 products, so that EM estimates about 100 parameters instead of 400, then the same average percent gap is 0.57%, favoring EM back again. Thus, we should be cautious about using the Markov chain choice model when we have too little data and too many products.

To form a base line, we also check the out-of-sample log-likelihoods when we fit a ranking-based choice model, which is the ground choice model that actually drives the choice process of the customers in the training and hold-out data. The papers by van Ryzin and Vulcano (2015) and van Ryzin and Vulcano (2016) give algorithms for estimating the parameters of the ranking-based choice model. We fit two versions of the ground choice model. In the first version, we estimate both the ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$ in the ground choice model and the corresponding probabilities $(\beta^1, \ldots, \beta^m)$. In the second version, we assume that we know the ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$ and we estimate only the probabilities $(\beta^1, \ldots, \beta^m)$. We refer to the first and second versions of the ground choice model as UR and KR, standing for unknown ranked lists and known ranked lists. In reality, it is highly unlikely to know the ranked lists of products that drive the choice process of the customers. Therefore, KR represents a situation that is not achievable in practice. For economy of space, we briefly summarize our results and defer the details to Online Appendix G. Over all of our problem instances, the out-of-sample log-likelihoods from EM are better than those from UR by 1.60% on average. Not surprisingly, since KR has access to the ranked lists that drive the customer choice behavior, it provides noticeably better log-likelihoods than EM and UR.

In Table 1, we compared the out-of-sample log-likelihoods obtained by EM, DM and ML. In Online Appendix H, we also compare the three benchmarks from the perspective of the out-of-sample root-mean-square error between the fitted and actual choice probabilities. To give an overview of our results, in our test problems, the root-mean-square errors obtained by EM are significantly smaller than those obtained by ML, but one should still be cautious about using the Markov chain choice model when we have too little data and too many parameters to estimate. In 10 out of 24 test problems, the root-mean-square errors of EM are smaller than those of DM.

Checking the out-of-sample log-likelihoods or root-mean-square errors are two possible methods to control for over-fitting. If a choice model with a large number of parameters over-fits to the training data, then its generalization performance on the hold-out data becomes poor. Another possible method to control for over-fitting is to check the Akaike information criterion (AIC). For any of the fitted choice models, AIC is given by $2p - 2L^*$, where p is the number of parameters of the choice model and L^* is the log-likelihood of the training data; see Akaike (1973). So, AIC penalizes a fitted choice model with a large number of parameters and a small log-likelihood. In Table 2, we show AIC for EM, DM and ML. When we have 10,000 or more customers in the training data, AIC for EM is generally better than AIC for ML. When we have 5,000 or fewer customers in the training data, AIC for ML is generally better than AIC for EM. Overall, the out-of-sample log-likelihoods and root-mean-square errors indicate that EM provides better prediction ability than ML when we have ample data. However, the situation is not as clear when we have little data. While the out-of-sample log-likelihoods and root-mean-square errors indicate that EM can still provide better prediction ability than ML, AIC points out that ML has an edge.

In Online Appendix I, we compare EM and ML with the independent demand model. Both EM and ML perform significantly better than the independent demand model.

(n,m)	Trn. data	Akaike information criterion EM	DM	ML	EM-DM	EM-ML	(n,m)	Trn. data	Akaike information criterion EM	DM	ML	EM-DM	EM-ML
$(11,21)$	2,500	8,386	8,369	8,378	-0.21%	-0.09%	$(21,31)$	2,500	12,046	12,018	11,541	-0.23%	-4.37%
$(11,21)$	5,000	16,498	16,462	16,670	-0.22%	1.03%	$(21,31)$	5,000	23,350	23,298	23,051	-0.23%	-1.30%
$(11,21)$	10,000	32,856	32,781	33,333	-0.23%	1.43%	$(21,31)$	10,000	45,974	45,870	46,134	-0.23%	0.35%
$(11,21)$	50,000	163,969	163,645	167,015	-0.20%	1.82%	$(21,31)$	50,000	227,174	226,669	231,152	-0.22%	1.72%
$(11,31)$	2,500	8,304	8,285	8,212	-0.22%	-1.12%	$(21,41)$	2,500	11,909	11,883	11,392	-0.22%	-4.54%
$(11,31)$	5,000	16,372	16,341	16,378	-0.19%	0.04%	$(21,41)$	5,000	23,149	23,095	22,821	-0.23%	-1.44%
$(11,31)$	10,000	32,687	32,625	32,838	-0.19%	0.46%	$(21,41)$	10,000	45,548	45,446	45,601	-0.23%	0.11%
$(11,31)$	50,000	161,721	161,426	163,148	-0.18%	0.87%	$(21,41)$	50,000	224,711	224,175	227,649	-0.24%	1.29%
$(11,51)$	2,500	8,616	8,597	8,528	-0.22%	-1.03%	$(21,61)$	2,500	12,275	12,246	11,687	-0.23%	-5.03%
$(11,51)$	5,000	17,023	16,991	17,014	-0.19%	-0.05%	$(21,61)$	5,000	23,908	23,848	23,380	-0.25%	-2.26%
$(11,51)$	10,000	33,887	33,825	34,081	-0.18%	0.57%	$(21,61)$	10,000	47,080	46,965	46,779	-0.25%	-0.64%
$(11,51)$	50,000	169,231	168,923	170,698	-0.18%	0.86%	$(21,61)$	50,000	232,667	232,116	233,751	-0.24%	0.46%

Table 2: AIC obtained by EM, DM and ML.

5.3 Hotel Data

We provide computational experiments on the data set based on Bodea et al. (2009), which comes from a hotel revenue management application.

Experimental Setup. In the data, we have booking records for customers with check-in dates between March 12, 2007 and April 15, 2007 for five different hotels. Each booking record gives the rate and room type availability at the time of the booking and the room type that was booked. The paper by van Ryzin and Vulcano (2015) uses this data and we closely follow this paper in our experimental setup. We define each product as a room type. Room types take values such as king non-smoking and queen smoking. We consider the substitution behavior between different products offered by the same hotel. As done in van Ryzin and Vulcano (2015), we preprocess the data to reduce sparsity. We discard the products that have fewer than 10 purchases in the data. Also, we eliminate the booking records where the observed bookings are not among the products that are available at the time of the booking. Finally, the data provides only the bookings and does not provide any information on the customers who did not book. For each booking record, we create four no-purchase records. These four booking records has the same product availability as the original booking record, but no product is purchased in these four booking records. The left side of Table 3 shows the number of products after eliminating those with fewer than 10 purchases and the number of booking records after we add the no-purchase records.

Results. We use five-fold cross-validation to check the out-of-sample performance of the fitted choice models. We randomly partition the data for each hotel into five equal segments. We use

Problem instance characteristics

	Hotel 1	Hotel 2	Hotel 3	Hotel 4	Hotel 5
Prds.	10	6	7	4	6
Custs.	6,575	930	5,730	1,330	1,185

Out-of-sample log-likelihood

	Hotel 1	Hotel 2	Hotel 3	Hotel 4	Hotel 5
EM	-1,049.92	-135.77	-821.44	-185.61	-181.92
DM	-1,049.97	-137.79	-821.18	-185.50	-181.79
ML	-1,095.92	-136.35	-842.03	-187.03	-184.34
EM-DM	0.01%	1.47%	-0.03%	-0.06%	-0.07%
EM-ML	4.20%	0.43%	2.45%	0.76%	1.32%

Table 3: Comparison of EM, DM and ML on the hotel data.

one partition as the hold-out data and the remaining portion as the training data. For EM, DM and ML, we use the training data to estimate the parameters of the corresponding choice model and check the out-of-sample performance of the fitted choice model using the hold-out data. We repeat this process five times, each of the five partitions being used once as the hold-out data. We average the results from the five so-called folds to report our findings for each hotel.

We summarize our computational results on the right side of Table 3, where we show the out-of-sample log-likelihoods obtained by EM, DM and ML averaged over the five folds. The last two rows on the right side of Table 3 compare the log-likelihood obtained by EM with those obtained by DM and ML by giving the percent gaps between the corresponding pairs of log-likelihoods. EM provides larger log-likelihoods than ML. The general trend is that the hotels for which we obtain the largest gaps between the log-likelihoods of EM and ML correspond to those for which we have the largest number of bookings. This observation is consistent with the discussion in the previous section. That is, the Markov chain choice model provides significant improvements over the multinomial logit model particularly when we have relatively ample data to estimate the parameters. In Online Appendix J, we compare the three benchmarks from the perspective of the out-of-sample root-mean-square error between the fitted and empirically observed choice probabilities. EM provides consistent improvements over ML.

6 Conclusions and Research Directions

In this paper, we developed an expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. One issue that often comes up in revenue management applications is demand censorship. Demand censorship refers to the fact that if there is no purchase for a product during a certain duration of time, then the data does not make it clear whether there was no customer arrival during this duration of time or the customers that arrived during this duration of time decided to purchase nothing. Vulcano et al. (2012) show that the expectation-maximization framework is useful to deal with demand censorship, since we can treat whether there is a customer arrival or not during a certain duration of time as a missing piece of the data. In Online Appendix K, we show how we can follow the work of Vulcano et al. (2012) to modify our expectation-maximization algorithm to deal with demand censorship. There are other research directions to pursue to improve the practical appeal of the Markov chain choice model. In the multinomial logit model, we can parameterize the mean utility of a product as a function of attributes such as price, quality and richness of features. In this case, rather than estimating the mean utility of each product separately, we estimate how much weight a customer puts on each attribute. Thus, the number of parameters in the multinomial logit model increases linearly with the number of attributes, rather than the number of products. In many practical applications, the number of attributes is much smaller than the number of products. We do not have an approach to parameterize the Markov chain choice model by using a small number of attributes. Thus, the number of parameters of the Markov chain choice model increases quadratically with the number of products. An interesting research direction is to parameterize the initial arrival and transition probabilities in the Markov chain choice model as a function of a small number of attributes. Such a parameterization currently appears to be non-trivial. After parameterizing the initial arrival and transition probabilities as a function of a small number of attributes, the next step is to come up with methods to estimate the parameters in the choice model.

Acknowledgments. We thank the associate editor and two anonymous referees whose comments improved the paper in many ways.

References

Akaike, H. (1973), Information theory and an extension of the maximum likelihood principle, *in* B. N. Petrov and F. Csaki, eds, '*Second International Symposium on Information Theory*', Akademia Kiado, Budapest, Hungary, pp. 267–281.

Bishop, C. M. (2006), *Pattern Recognition and Machine Learning*, Springer, New York, NY.

Blanchet, J., Gallego, G. and Goyal, V. (2013), A Markov chain approximation to choice modeling, Technical report, Columbia University, Department of Industrial Engineering and Operations Research.

Bodea, T., Ferguson, M. and Garrow, L. (2009), 'Data set–choice-based revenue management: Data from a major hotel chain', *Manufacturing and Service Operations Management* **11**(2), 356–361.

Boyd, S. and Vandenberghe, L. (2005), *Convex Optimization*, Cambridge University Press, Cambridge, UK.

Campo, K., Gijsbrechts, E. and Nisol, P. (2003), 'The impact of retailer stockouts on whether, how much, and what to buy', *International Journal of Research in Marketing* **30**(3), 273–286.

Chong, J.-K., Ho, T.-H. and Tang, C. S. (2001), 'A modeling framework for category assortment planning', *Manufacturing & Service Operations Management* **3**(3), 191–210.

Dai, J., Ding, W., Kleywegt, A. J., Wang, X. and Zhang, Y. (2014), Choice based revenue management for parallel flights, Technical report, Georgia Tech, Atlanta, GA.

Dempster, A., Laird, N. and Rubin, D. (1977), 'An expectation-maximization algorithm to estimate a general class of non-parametric choice models', *Journal of the Royal Statistical Society B* **39**(1), 1–38.

Desir, A., Goyal, V., Segev, D. and Ye, C. (2015), Capacity constrained assortment optimization under the Markov chain based choice model, Technical report, Columbia University, New York, NY.

Farias, V., Jagabathula, S. and Shah, D. (2013), 'A non-parametric approach to modeling choice with limited data', *Management Science* **59**(2), 305–322.

Feldman, J. B. and Topaloglu, H. (2014), Revenue management under the Markov chain choice model, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at `http://legacy.orie.cornell.edu/~huseyin/publications/publications.html`.

Guadagni, P. M. and Little, J. D. C. (2008), 'A logit model of brand choice calibrated on scanner data', *Marketing Science* **27**(1), 29–48.

Honhon, D., Gaur, V. and Seshadri, S. (2010), 'Assortment planning and inventory decisions under stock-out based substitution', *Operations Research* **58**(5), 1364–1379.

Honhon, D., Jonnalagedda, S. and Pan, X. A. (2012), 'Optimal algorithms for assortment selection under ranking-based consumer choice models', *Manufacturing & Service Operations Management* **14**(2), 279–289.

Jagabathula, S. and Rusmevichientong, P. (2016), 'A nonparametric joint assortment and price choice model', *Management Science* (to appear).

Jagabathula, S. and Vulcano, G. (2016), A partial-order-based model to estimate individual preferences using panel data, Technical report, New York University, New York, NY.

Kok, A. and Fisher, M. (2007), 'Demand estimation and assortment optimization under susbstitution', *Operations Research* **55**, 1001–1021.

Mahajan, S. and van Ryzin, G. (2001), 'Stocking retail assortments under dynamic customer substitution', *Operations Research* **49**(3), 334–351.

McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., '*Frontiers in Economics*', Academic Press, New York, NY, pp. 105–142.

Misra, K. (2008), Understanding retail assortments in competitive markets, Technical report, Northwestern University, Evanston, IL.

Nettleton, D. (1999), 'Convergence properties of the EM algorithm in constrained parameter spaces', *The Canadian Journal of Statistics* **27**(3), 639–648.

Puterman, M. L. (1994), *Markov Decision Processes*, John Wiley and Sons, Inc., New York.

Sloot, L. M., Verhoef, P. C. and Franses, P. H. (2005), 'The impact of brand equity and the hedonic level of products on consumer stock-out reactions', *Journal of Retailing* **81**(1), 15–34.

Smith, J. C., Lim, C. and Alptekinoglu, A. (2009), 'New product introduction against a predator: A bilevel mixed-integer programming approach', *Naval Research Logistics* **59**(8), 714–729.

van Ryzin, G. and Vulcano, G. (2008), 'Computing virtual nesting controls for network revenue management under customer choice behavior', *Manufacturing & Service Operations Management* **10**(3), 448–467.

van Ryzin, G. and Vulcano, G. (2015), 'A market discovery algorithm to estimate a general class of nonparametric choice models', *Management Science* **61**(2), 281–300.

van Ryzin, G. and Vulcano, G. (2016), 'An expectation-maximization method to estimate a rank-based choice model of demand', *Operations Research* (to appear).

Vulcano, G., van Ryzin, G. and Chaar, W. (2010), 'Choice-based revenue management: An empirical study of estimation and optimization', *Manufacturing and Service Operations Management* **12**(3), 371–392.

Vulcano, G., van Ryzin, G. J. and Ratliff, R. (2012), 'Estimating primary demand for substitutable products from sales transaction data', *Operations Research* **60**(2), 313–334.

Wu, C. F. J. (1983), 'On the convergence of the EM algorithm', *The Annals of Statistics* **11**(1), 95–103.

Zinn, W. and Liu, P. C. (2001), 'Consumer response to retail stockouts', *Journal of Business Logistics* **22**(1), 49–71.

A Online Appendix: Maximization Step with Parameter Lower Bounds

We consider the maximization step of our expectation-maximization algorithm when the parameters that we are trying to estimate are known to be bounded away from zero by some $\epsilon > 0$. Noting problem (11), when we have a lower bound of $\epsilon > 0$ on the parameters, to obtain estimates of the parameters $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ at iteration ℓ, we solve the problem

$$\max \left\{ \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i \; : \; \sum_{i \in N} \lambda_i = 1, \; \lambda_i \geq \epsilon \; \forall \, i \in N \right\}. \tag{14}$$

Associating the Lagrange multiplier γ with the first constraint, the Lagrangian is $\mathcal{L}(\boldsymbol{\lambda} \,|\, \gamma) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \gamma \, (1 - \sum_{i \in N} \lambda_i)$. For a fixed value of the Lagrange multiplier γ, we maximize $\mathcal{L}(\boldsymbol{\lambda} \,|\, \gamma)$ subject to the constraint that $\lambda_i \geq \epsilon$ for all $i \in N$. Since $\mathcal{L}(\boldsymbol{\lambda} \,|\, \gamma)$ is separable by $(\lambda_1, \ldots, \lambda_n)$, for all $i \in N$, we can equivalently solve

$$\max \left\{ \sum_{t \in T} \hat{F}_i^{t,\ell} \log \lambda_i - \gamma \, \lambda_i \; : \; \lambda_i \geq \epsilon_i \right\}. \tag{15}$$

Setting the derivative of the objective function above with respect to λ_i to zero, the unconstrained optimal value of λ_i in the problem above is given by $\sum_{t \in T} \hat{F}_i^{t,\ell} / \gamma$. Thus, noting that the objective function of problem (15) is concave in λ_i, the optimal value of λ_i in this problem is given by $\lambda_i^*(\gamma) = \max\{\sum_{t \in T} \hat{F}_i^{t,\ell} / \gamma, \epsilon\}$, where we make the dependence of the optimal solution on the Lagrange multiplier γ explicit. By the discussion in Section 5.5.5 in Boyd and Vandenberghe (2005), since the objective function of problem (14) is strictly concave and its feasible space is convex, the optimal solution to problem (14) is unique and it is of the form $(\lambda_1^*(\gamma^*), \ldots, \lambda_n^*(\gamma^*))$ for the optimal value γ^* of the Lagrange multiplier. Since the optimal solution to problem (14) satisfies the first constraint in this problem, we must have $\sum_{i \in N} \lambda_i^*(\gamma^*) = 1$, which implies that we must have $\sum_{i \in N} \max\{\sum_{t \in T} \hat{F}_i^{t,\ell} / \gamma^*, \epsilon\} = 1$. Thus, we can find the optimal value of the Lagrange multiplier by solving $\sum_{i \in N} \max\{\sum_{t \in T} \hat{F}_i^{t,\ell} / \gamma, \epsilon\} = 1$ for γ. Once we have the optimal value γ^* of the Lagrange multiplier, the optimal solution to problem (14) is $(\lambda_1^*(\gamma^*), \ldots, \lambda_n^*(\gamma^*))$.

Next, we consider the problem of solving $\sum_{i \in N} \max\{\sum_{t \in T} \hat{F}_i^{t,\ell} / \gamma, \epsilon\} = 1$ for γ. For notational brevity, we let $\Delta_i = \sum_{t \in T} \hat{F}_i^{t,\ell}$ so that we consider the problem of solving $\sum_{i \in N} \max\{\Delta_i / \gamma, \epsilon\} = 1$ for γ. Without loss of generality, we can assume that $\Delta_i > 0$ for all $i \in N$. In particular, if we have $\Delta_i = 0$, then the optimal value of λ_i in problem (14) is zero and we can drop this decision variable from problem (14). Also, without loss of generality, we can assume that $\epsilon < 1/n$. In particular, noting the constraints $\lambda_i \geq \epsilon$ for all $i \in N$ in problem (14), if we have $\epsilon > 1/n$, then $\sum_{i \in N} \lambda_i \geq \sum_{i \in N} \epsilon = n\epsilon > 1$, which implies that problem (14) does not have a feasible solution. If, on the other hand, $\epsilon = 1/n$, then the only feasible solution to problem (14) is obtained by setting $\lambda_i = 1/n$ for all $i \in N$. Since we can assume that $\epsilon < 1/n$, we have $\lim_{\gamma \to \infty} \sum_{i \in N} \max\{\Delta_i/\gamma, \epsilon\} = n\,\epsilon < 1$. Also, we have $\lim_{\gamma \to 0} \sum_{i \in N} \max\{\Delta_i/\gamma, \epsilon\} = \infty > 1$. Noting also that $\sum_{i \in N} \max\{\Delta_i/\gamma, \epsilon\}$ is continuous and strictly decreasing in γ until it reaches a value that is strictly less than one, there exists a unique value of γ that satisfies $\sum_{i \in N} \max\{\Delta_i/\gamma, \epsilon\} = 1$. To find this value of γ, we partition the interval $[0, \infty)$ into at most $n+1$ intervals such that the largest one of Δ_i/γ and ϵ does not change for any $i \in N$ as γ takes values in one of these intervals. In particular, indexing the products such that $\Delta_1 \leq \Delta_2 \leq \ldots \leq \Delta_n$, we define the interval I_i as $I_i = [\Delta_i/\epsilon, \Delta_{i+1}/\epsilon)$, with the convention that $I_0 = (0, \Delta_1/\epsilon)$ and $I_n = [\Delta_n/\epsilon, \infty)$. If we have $\gamma \in I_i$, then $\Delta_1/\epsilon \leq \ldots \leq \Delta_i/\epsilon \leq \gamma < \Delta_{i+1}/\epsilon \leq \ldots \leq \Delta_n/\epsilon$, which implies that $\max\{\Delta_1/\gamma, \epsilon\} = \epsilon, \ldots, \max\{\Delta_i/\gamma, \epsilon\} = \epsilon$ and $\max\{\Delta_{i+1}/\gamma, \epsilon\} = \Delta_{i+1}/\gamma, \ldots, \max\{\Delta_n/\gamma, \epsilon\} = \Delta_n/\gamma$. In this case, if we have $\gamma \in I_i$, then $\sum_{k \in N} \max\{\Delta_k/\gamma, \epsilon\} = i\,\epsilon + \sum_{k=i+1}^n \Delta_k/\gamma$. Therefore, for all $i \in N \cup \{0\}$, we can solve $i\,\epsilon + \sum_{k=i+1}^n \Delta_k/\gamma = 1$ for γ. Using $\gamma_i^* = \sum_{k=i+1}^n \Delta_k/(1 - i\,\epsilon)$ to denote the solution, if $\gamma_i^* \in I_i$, then γ_i^* is the value of γ that satisfies $\sum_{k \in N} \max\{\Delta_k/\gamma, \epsilon\} = 1$. Since we know that there exists a unique value of such γ, exactly one of $\{\gamma_i^* : i \in N \cup \{0\}\}$ satisfies $\gamma_i^* \in I_i$.

B Online Appendix: Proof of Theorem 1

The proof of Theorem 1 is based on checking the regularity conditions in Nettleton (1999). In Lemma 2 below, we show that the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is a continuous and differentiable function of $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ over Ω. In Lemma 3 below, we show that the set $\Omega_\alpha = \{(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Omega : L_I(\boldsymbol{\lambda}, \boldsymbol{\rho}) \geq \alpha\}$ is compact for any $\alpha \in \Re$. Finally, we observe that $\hat{F}_i^{t,\ell}$ and $\hat{X}_{i,j}^{t,\ell}$ in the expectation step of the expectation-maximization algorithm depend on the current parameter estimates $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. In Lemma 4 below, we show that the likelihood function $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j}$ in the maximization step of the expectation-maximization algorithm is a continuous function of $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ over $\Omega \times \Omega$. These three lemmas collectively show that our expectation-maximization algorithm satisfies the regularity conditions in Nettleton (1999), in which case, Theorem 1 in our paper follows from Theorem 2 in Nettleton (1999).

In the next lemma, we focus on the first regularity condition.

Lemma 2 *The function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is a continuous and differentiable function of $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ over Ω.*

Proof. Consider any non-empty subset of products $S \subset N$. For any $(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Omega$, we note that $\sum_{j \in N \setminus S} \rho_{i,j} \leq \sum_{j \in N} \rho_{i,j} - |S|\,\epsilon = 1 - |S|\,\epsilon < 1$ for all $i \in N \setminus S$. Therefore, by the discussion at the end of Section 2, using the vector $\bar{\boldsymbol{\lambda}} = \{\lambda_i : i \in N \setminus S\}$ and the matrix $\bar{\boldsymbol{\rho}} = \{\rho_{i,j} : i, j \in N \setminus S\}$, the vector $\bar{\boldsymbol{\Theta}}(S) = \{\Theta_i(S) : i \in N \setminus S\}$ is given by $\bar{\boldsymbol{\Theta}}(S) = ((I - \bar{\boldsymbol{\rho}})^{-1})^\top \bar{\boldsymbol{\lambda}}$. Each entry in the inverse of a matrix, when it exists, is continuous and differentiable in the entries of the matrix involved in the inversion operation. Similarly, each entry in the product of two matrices is continuous and differentiable in the entries of the matrices that are involved in the multiplication operation. Since we have $\bar{\boldsymbol{\Theta}}(S) = ((I - \bar{\boldsymbol{\rho}})^{-1})^\top \bar{\boldsymbol{\lambda}}$, it follows that $\boldsymbol{\Theta}_i(S)$ is continuous and differentiable in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ for all $i \in N \setminus S$. In this case, noting that we also have $\Theta_i(S) = \lambda_i + \sum_{j \in N \setminus S} \rho_{j,i}\, \Theta_j(S)$ for all $i \in S$, it also follows that $\Theta_i(S)$ is differentiable and continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ for all $i \in S$. Therefore, noting (2), since $\log x$ is continuous and differentiable in x, $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous and differentiable in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$, which is the desired result. □

In the next lemma, we focus on the second regularity condition.

Lemma 3 *For any $\alpha \in \Re$, the set $\Omega_\alpha = \{(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Omega : L_I(\boldsymbol{\lambda}, \boldsymbol{\rho}) \geq \alpha\}$ is compact.*

Proof. Noting the definition of Ω, we observe that Ω is closed and bounded. Thus, Ω_α is bounded as well, since $\Omega_\alpha \subset \Omega$. We show that Ω_α is closed. Consider any sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ in Ω_α with a limit point $(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*)$. Since $\Omega_\alpha \subset \Omega$, the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ is also in Ω. Noting that Ω is closed, any limit point of the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ is in Ω, which implies that $(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \in \Omega$. We claim that $L_I(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \geq \alpha$. To get a contradiction to the claim, assume that $L_I(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) < \alpha$. Thus, there exists some $\delta > 0$ such that $L_I(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \leq \alpha - \delta$. Since the sequence $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ is in Ω_α, we have $L_I(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) \geq \alpha$ for all $\ell = 1, 2, \ldots$. By Lemma 2, $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. Having $L_I(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \leq \alpha - \delta$ for some $\delta > 0$ and $L_I(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) \geq \alpha$ for all $\ell = 1, 2, \ldots$ contradicts the fact that $(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*)$ is a limit point of the sequence of $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ and $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$, in which case, the claim follows. Therefore, we have $(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \in \Omega$ and $L_I(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \geq \alpha$, which implies that $(\boldsymbol{\lambda}^*, \boldsymbol{\rho}^*) \in \Omega_\alpha$. Thus, a limit point of any sequence in Ω_α is also in Ω_α, which implies that Ω_α is closed. □

In the next lemma, we focus on the third regularity condition.

Lemma 4 *Noting that* $\hat{F}_i^{t,\ell}$ *and* $\hat{X}_{i,j}^{t,\ell}$ *in the expectation step of the expectation-maximization algorithm depend on the current parameter estimates* $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$*, the likelihood function* $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^{t,\ell} \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^{t,\ell} \log \rho_{i,j}$ *in the maximization step is a continuous function of* $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ *and* $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ *over* $\Omega \times \Omega$*.*

Proof. Since $\log x$ is continuous in x, $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. Next, we show that $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous in $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. By the argument in the proof of Lemma 2, $\Theta_i(S)$ computed through (1) is continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. By the same argument, $\Psi_k(i, S)$ computed through (5) is continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. Thus, noting (6) and (10), $\mathbb{E}\{F_i \mid \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{E}\{X_{i,j}(S) \mid \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ are continuous in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$, although we do not make the dependence of these quantities on $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ explicit. In the expectation step, $\hat{F}_i^{t,\ell}$ and $\hat{X}_{i,j}^{t,\ell}$ are given by the last two conditional expectations when the parameters of the Markov chain choice model are $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and we replace S with $\hat{S}^t$. Therefore, $\hat{F}_i^{t,\ell}$ and $\hat{X}_{i,j}^{t,\ell}$ are continuous in $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. In the expression for $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho})$, only $\{\hat{F}_i^{t,\ell} : i \in N,\ t \in T\}$ and $\{\hat{X}_{i,j}^{t,\ell} : i, j \in N,\ t \in T\}$ depend on $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. So, $L_C^\ell(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is continuous in $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. $\square$

C Online Appendix: Regularity Conditions without Parameter Lower Bounds

In this section, we give an example to demonstrate that Lemmas 2, 3 and 4 may not hold when we do not have a lower bound of $\epsilon > 0$ on the parameters of the Markov chain choice model, which implies that the regularity conditions that allow us to show Theorem 1 may not hold without a lower bound of $\epsilon > 0$ on the parameters. We focus on Lemmas 2 and 4, but we can also provide an example to demonstrate that Lemma 3 may not hold. When we do not have a lower bound of $\epsilon > 0$ on the parameters, the set of possible parameter values is given by $\bar{\Omega} = \{(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \Re_+^n \times \Re_+^{n \times n} : \sum_{i \in N} \lambda_i = 1,\ \sum_{j \in N} \rho_{i,j} = 1\ \forall\, i \in N\}$. It is simple to see that there exist values of $(\boldsymbol{\lambda}, \boldsymbol{\rho}) \in \bar{\Omega}$ and subsets $S \subset N$ such that the system of equations in (1) does not have a solution when we solve this system of equations with such values of $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ and subsets S. Consider the case $n = 2$. We fix $\lambda_1 > 0$, $\lambda_2 > 0$, $\rho_{11} > 0$ and $\rho_{12} = 1 - \rho_{11} > 0$ arbitrarily. If we have $\rho_{21} = 0$ and $\rho_{22} = 1$, then the system of equations in (1) with $S = \{1\}$ is given by $\Theta_1(S) = \lambda_1 + \rho_{21}\, \Theta_2(S)$ and $\Theta_2(S) = \lambda_2 + \Theta_2(S)$. Noting the second equality and the fact that $\lambda_2 > 0$, the last system of equations does not have a solution. Thus, the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$ is not well-defined at any parameter value $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ that satisfies $\rho_{21} = 0$ and $\rho_{22} = 1$ when the subset $S = \{1\}$ is offered to the customers in the data. In this case, Lemma 2 does not hold since the likelihood function $L_I(\boldsymbol{\lambda}, \boldsymbol{\rho})$

is not a continuous function of $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ at this parameter value when we offer the subset $S = \{1\}$ to the customers in the data. Similarly, Lemma 4 does not hold either since $(\Theta_1(S), \Theta_2(S))$ with $S = \{1\}$ is not a continuous function of $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ at this parameter value.

D Online Appendix: Parameter Estimates without Parameter Lower Bounds

In Section 4.4, we give a convergence result for our expectation-maximization algorithm when we impose a lower bound of $\epsilon > 0$ on the parameters. In the expectation step of our algorithm, we need to solve the systems of equations in (1) and (5) to compute $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$ and $\mathbb{E}\{X_{i,j}(\hat{S}_t) \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$. A natural question is whether our expectation-maximization algorithm successfully avoids parameter values that render these systems of equations unsolvable, even when we do not impose a lower bound of $\epsilon > 0$ on the parameters. Answering this question affirmatively ensures that our algorithm is well-posed. In this section, we show that as long as the initial parameter estimates are strictly positive, even if we do not impose a lower bound of $\epsilon > 0$ on the parameters, our algorithm always generates a sequence of parameters $\{(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell) : \ell = 1, 2, \ldots\}$ such that there exist unique solutions to the systems of equations in (1) and (5) when we solve these systems of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ for any subset in the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$. In the next lemma, we show that certain parameter estimates remain strictly positive over the successive iterations of the expectation-maximization algorithm.

Lemma 5 *Fix an arbitrary customer $t \in T$. Letting product $k \in \hat{S}^t$ be such that $\hat{\boldsymbol{Z}}^t = \boldsymbol{e}_k$, assume that $\lambda_k^\ell > 0$, $\lambda_i^\ell > 0$ for all $i \in N \setminus \hat{S}^t$ and $\rho_{i,k}^\ell > 0$ for all $i \in N \setminus \hat{S}^t$. Then, we have $\lambda_k^{\ell+1} > 0$, $\lambda_i^{\ell+1} > 0$ for all $i \in N \setminus \hat{S}^t$ and $\rho_{i,k}^{\ell+1} > 0$ for all $i \in N \setminus \hat{S}^t$.*

Proof. We use $(\Theta_1^\ell(\hat{S}^t), \ldots, \Theta_n^\ell(\hat{S}^t))$ to denote the solution to (1) when we solve this system of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. For all $i \in N \setminus \hat{S}^t$, we observe that $\sum_{j \in N \setminus \hat{S}^t} \rho_{i,j}^\ell < \sum_{j \in N \setminus \hat{S}^t} \rho_{i,j} + \rho_{i,k}^\ell \leq \sum_{j \in N} \rho_{i,j}^\ell = 1$, where the first inequality uses the fact that $\rho_{i,k}^\ell > 0$ for all $i \in N \setminus \hat{S}^t$ and the second inequality uses the fact that $k \in \hat{S}^t$. In this case, since $\sum_{j \in N \setminus \hat{S}^t} \rho_{i,j}^\ell < 1$, by the argument at the end of Section 2, there exists a unique solution to (1) when we solve this system of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. Furthermore, this solution has non-negative entries, so that $\Theta_i^\ell(\hat{S}^t) \geq 0$ for all $i \in N$. By the definition of $(\Theta_1^\ell(\hat{S}^t), \ldots, \Theta_n^\ell(\hat{S}^t))$, we have $\Theta_i^\ell(\hat{S}^t) = \lambda_i^\ell + \sum_{j \in N \setminus \hat{S}^t} \rho_{j,i}^\ell \, \Theta_j^\ell(\hat{S}^t)$ for all $i \in N$. Since $\lambda_i^\ell > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$ and $\Theta_i^\ell(\hat{S}^t) \geq 0$ for all $i \in N$, the last equality implies that $\Theta_i^\ell(\hat{S}^t) > 0$ for all

$i \in \{k\} \cup (N \setminus \hat{S}^t)$. We use $\{\Psi_k^\ell(i, \hat{S}^t) : i \in N \setminus \hat{S}^t\}$ to denote the solution to (5) when we solve this system of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. Since $\sum_{j \in N \setminus \hat{S}^t} \rho_{i,j}^\ell < 1$ for all $i \in N \setminus \hat{S}^t$, an argument similar to the one at the end of Section 2 implies that there exists a unique solution to (5) when we solve this system of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. Furthermore, this solution has non-negative entries, so that $\Psi_k^\ell(i, \hat{S}^t) \geq 0$ for all $i \in N$. In this case, since $\rho_{i,k}^\ell > 0$ for all $i \in N \setminus \hat{S}^t$, by (5), it follows that $\Psi_k(i, \hat{S}^t) > 0$ for all $i \in N \setminus \hat{S}^t$. Also, we have $\Psi_k^\ell(k, \hat{S}^t) = 1$ by definition. Thus, we have $\Psi_k(i, \hat{S}^t) > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$. So, by (6), $\{\hat{F}_i^{t,\ell} : i \in \{k\} \cup (N \setminus \hat{S}^t)\}$ computed in the expectation step satisfy $\hat{F}_i^{t,\ell} = \Psi_k^\ell(i, \hat{S}^t)\, \lambda_i^\ell / \Theta_k^\ell(\hat{S}^t) > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$, where we use the fact that $\lambda_i^\ell > 0$ and $\Psi_k^\ell(i, \hat{S}^t) > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$ and $\Theta_k^\ell(\hat{S}^t) > 0$. Noting the explicit description of the maximization step at the end of Section 4.3, we get $\lambda_i^{\ell+1} = \sum_{t \in T} \hat{F}_i^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{F}_j^{t,\ell} > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$. Similarly, noting that $\Psi_k^\ell(k, \hat{S}^t) = 1$, by (10), $\{\hat{X}_{i,k}^{t,\ell} : i \in N \setminus \hat{S}^t\}$ computed in the expectation step of the expectation-maximization algorithm satisfy $\hat{X}_{i,k}^{t,\ell} = \rho_{i,k}^\ell\, \Theta_i^\ell(\hat{S}^t) / \Theta_k^\ell(\hat{S}^t) > 0$ for all $i \in N \setminus \hat{S}^t$, where we use the fact that $\rho_{i,k}^\ell > 0$ for all $i \in N \setminus \hat{S}^t$ and $\Theta_i^\ell(\hat{S}^t) > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$. Noting the explicit description of the maximization step at the end of Section 4.3, we get $\rho_{i,k}^{\ell+1} = \sum_{t \in T} \hat{X}_{i,k}^{t,\ell} / \sum_{t \in T} \sum_{j \in N} \hat{X}_{i,j}^{t,\ell} > 0$ for all $i \in N \setminus \hat{S}^t$. $\square$

Fix an arbitrary customer $t \in T$ and let product $k \in \hat{S}^t$ be such that $\hat{\boldsymbol{Z}}^t = \boldsymbol{e}_k$. Consider solving the systems of equations in (1) and (5) after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. By Lemma 5, if the initial parameter estimates in our expectation-maximization algorithm satisfy $\lambda_i^1 > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$ and $\rho_{i,k}^1 > 0$ for all $i \in N \setminus \hat{S}^t$, then we have $\lambda_i^\ell > 0$ for all $i \in \{k\} \cup (N \setminus \hat{S}^t)$ and $\rho_{i,k}^\ell > 0$ for all $i \in N \setminus \hat{S}^t$ and $\ell = 1, 2, \ldots$. In this case, by the same argument at the beginning of the proof of Lemma 5, it follows that there exist unique solutions to the systems of equations in (1) and (5) when we solve these systems of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$. By this discussion, if the initial parameter estimates in our expectation-maximization algorithm satisfy $\lambda_i^1 > 0$ for all $i \in N$ and $\rho_{i,j}^1 > 0$ for all $i, j \in N$, then there exist unique solutions to the systems of equations in (1) and (5) when we solve these systems of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ for any customer $t \in T$.

E Online Appendix: Pseudo-Code for the Expectation-Maximization Algorithm

In this section, we give a pseudo-code for our expectation-maximization algorithm. The inputs to our expectation-maximization algorithm are the initial parameter estimates $(\boldsymbol{\lambda}^1, \boldsymbol{\rho}^1)$, which satisfy

$\sum_{i\in N}\lambda_i^1 = 1$, $\sum_{j\in N}\rho_{i,j}^1 = 1$ for all $i \in N$, $\lambda_i^1 \geq 0$ for all $i \in N$ and $\rho_{i,j}^1 \geq 0$ for all $i, j \in N$. By the discussion in Online Appendix D, if the initial parameter estimates satisfy $\lambda_i^1 > 0$ for all $i \in N$ and $\rho_{i,j}^1 > 0$ for all $i, j \in N$, then there exist unique solutions to the systems of equations in (1) and (5) when we solve these systems of equations after replacing $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ with $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$ and S with $\hat{S}^t$ for any customer $t \in T$ and for any iteration $\ell = 1, 2, \ldots$. Thus, we choose initial parameter estimates that satisfy this condition. Below is a pseudo-code for our expectation-maximization algorithm.

Step 1. Choose the initial parameter estimates $(\boldsymbol{\lambda}^1, \boldsymbol{\rho}^1)$ such that $\sum_{i\in N}\lambda_i^1 = 1$, $\sum_{j\in N}\rho_{i,j}^1 = 1$ for all $i \in N$, $\lambda_i^1 > 0$ for all $i \in N$ and $\rho_{i,j}^1 > 0$ for all $i, j \in N$. Set the iteration counter $\ell = 1$.

Step 2. (Expectation) For each customer $t \in T$, let $k_t \in \hat{S}^t$ be the product such that $\hat{\boldsymbol{Z}}^t = \boldsymbol{e}_{k_t}$ and execute Steps 2.a and 2.b.

Step 2.a. Compute $\{\Theta_i^\ell(\hat{S}^t) : i \in N\}$ and $\{\Psi_{k_t}^\ell(i, \hat{S}_t) : i \in N\}$ as follows. To compute $\{\Theta_i^\ell(\hat{S}^t) : i \in N \setminus \hat{S}^t\}$, solve the system of equations

$$\Theta_i^\ell(\hat{S}^t) = \lambda_i^\ell + \sum_{j\in N\setminus\hat{S}^t}\rho_{j,i}^\ell\,\Theta_j^\ell(\hat{S}^t) \qquad \forall\, i \in N \setminus \hat{S}^t.$$

To compute $\{\Theta_i^\ell(\hat{S}^t) : i \in \hat{S}^t\}$, after solving the system above, set $\Theta_i^\ell(\hat{S}^t) = \lambda_i + \sum_{j\in N\setminus\hat{S}^t}\rho_{j,i}^\ell\,\Theta_j^\ell(\hat{S}^t)$ for all $i \in \hat{S}^t$. To compute $\{\Psi_{k_t}^\ell(i, \hat{S}_t) : i \in N \setminus \hat{S}^t\}$, solve the system of equations

$$\Psi_{k_t}^\ell(i, \hat{S}^t) = \rho_{i,k_t}^\ell + \sum_{j\in N\setminus\hat{S}^t}\rho_{i,j}^\ell\,\Psi_{k_t}^\ell(j, \hat{S}^t) \qquad \forall\, i \in N \setminus \hat{S}^t.$$

To compute $\{\Psi_{k_t}^\ell(i, \hat{S}_t) : i \in \hat{S}^t\}$, set $\Psi_{k_t}^\ell(i, \hat{S}_t) = \mathbf{1}(i = k_t)$ for all $i \in \hat{S}^t$, where we use $\mathbf{1}(\cdot)$ to denote the indicator function.

Step 2.b. Compute $\{\hat{F}_i^{t,\ell} : i \in N\}$ and $\{\hat{X}_{i,j}^{t,\ell} : i, j \in N\}$ as follows. For all $i \in N$, set $\hat{F}_i^{t,\ell} = \Psi_{k_t}^\ell(i, \hat{S}^t)\,\lambda_i^\ell\,/\,\Theta_{k_t}^\ell(\hat{S}^t)$. For all $i, j \in N$, set $\hat{X}_{i,j}^{t,\ell} = \mathbf{1}(i \notin \hat{S}^t)\,\Psi_{k_t}^\ell(j, \hat{S}^t)\,\rho_{i,j}^\ell\,\Theta_i^\ell(\hat{S}^t)\,/\,\Theta_{k_t}^\ell(\hat{S}^t)$.

Step 3. (Maximization) For all $i \in N$, set $\lambda_i^{\ell+1} = \sum_{t\in T}\hat{F}_i^{t,\ell} / \sum_{t\in T}\sum_{j\in N}\hat{F}_j^{t,\ell}$. For all $i, j \in N$, set $\rho_{i,j}^{\ell+1} = \sum_{t\in T}\hat{X}_{i,j}^{t,\ell} / \sum_{t\in T}\sum_{k\in N}\hat{X}_{i,k}^{t,\ell}$. Increase the iteration counter ℓ by one and go to Step 2.

Note that Step 2.a uses (1) and (5). Step 2.b uses (6) and (10). Finally, Step 3 builds on the discussion in Section 4.3.

(n,m)	Trn. data	EM CPU secs.	EM No. itns.	DM CPU secs.	ML CPU secs.
$(11,21)$	2,500	13.60	42	404.11	1.17
$(11,21)$	5,000	14.76	40	1,054.66	0.46
$(11,21)$	10,000	17.32	37	2,147.39	0.31
$(11,21)$	50,000	44.84	35	19,382.17	4.41
$(11,31)$	2,500	12.78	40	449.86	0.22
$(11,31)$	5,000	12.81	35	1,029.46	0.33
$(11,31)$	10,000	15.27	33	2,827.09	0.48
$(11,31)$	50,000	39.43	31	19,984.20	2.83
$(11,51)$	2,500	12.41	39	469.17	0.22
$(11,51)$	5,000	13.19	36	1,110.74	0.34
$(11,51)$	10,000	16.46	35	2,657.42	0.36
$(11,51)$	50,000	42.75	33	20,280.00	4.81

(n,m)	Trn. data	EM CPU secs.	EM No. itns.	DM CPU secs.	ML CPU secs.
$(21,31)$	2,500	212.78	52	5,530.82	0.56
$(21,31)$	5,000	412.27	50	10,917.31	0.62
$(21,31)$	10,000	779.42	48	34,005.74	1.57
$(21,31)$	50,000	3818.24	47	197,743.57	29.74
$(21,41)$	2,500	210.51	52	5,077.26	0.47
$(21,41)$	5,000	386.43	47	12,339.03	0.53
$(21,41)$	10,000	740.44	46	44,566.01	2.88
$(21,41)$	50,000	3421.86	43	226,993.18	8.35
$(21,61)$	2,500	207.07	50	5,961.53	0.37
$(21,61)$	5,000	365.23	44	12,308.80	0.58
$(21,61)$	10,000	699.87	43	33,039.75	1.50
$(21,61)$	50,000	3128.47	39	229,299.92	12.86

Table 4: Computation times for EM, DM and ML, along with number of iterations for EM.

F Online Appendix: Computation Times for Estimating the Parameters

We provide the details of the computation times for EM, DM and ML. We give our results in Table 4. The first column in this table shows the parameter combination (n, m) in the ground choice model. The second column shows the number of customers in the training data. The third column shows the CPU seconds for EM. We stop our expectation-maximization algorithm when the incomplete log-likelihood function increases by less than 0.01% in two successive iterations. The fourth column shows the number of iterations for our algorithm until termination. The fifth and sixth columns respectively show the CPU seconds for DM and ML. Our results indicate that EM takes somewhere between about 10 seconds to about an hour. Since we do not solve parameter estimation problems in real time, such computation times are reasonable, but if we have a large number of products and a large number of customers in the training data, such as $n = 50$ and $\tau = 50,000$, then the computation times for EM will be impractical. DM is computationally demanding even with $n = 20$. In particular, the computation times for DM can exceed 60 hours. The computation times for ML are consistently below 30 seconds.

Note that EM needs to solve the system of equations in (1) for the subset of products offered to each customer in the training data. If there are n products, then there are at most 2^n different subsets in the training data. For the test problems with $n = 11$ products, we have at most $2,048$ different subsets that can potentially be offered to the customers in the training data and we can a priori solve the system of equations in (1) for all possible subsets that can potentially be offered to the customers and store the solutions. A priori solving the system of equations in (1) for all possible subsets that can potentially be offered to the customers and storing the solutions make the implementation of our expectation-maximization algorithm particularly efficient. For the test problems with $n = 21$ products, we have more than two million different subsets that can be offered to the customers. In this case, it is more efficient to solve the system of equations in (1) for $|T|$ subsets $\{\hat{S}^t : t \in T\}$ that are in the training data. Also, EM needs to solve the system of equations in (5) for the subset of products offered to each customer and the product purchased by each customer in the training data. Similar observations apply when solving this system of equations a priori and storing the solutions. This discussion indicates that EM can be particularly efficient when the customers in the training data are offered one of a few possible subsets of products. For example, if the customers in the training data are offered one of 10 possible subsets of products, then EM takes about 20 seconds, even when we have 50,000 customers in the training data. It is worthwhile to emphasize that customers in practice are often offered one of a few possible subsets of products, which correspond to the subset of products offered on a store shelf, possibly adjusted by stock-outs as we run out of some products.

G Online Appendix: Estimating the Parameters of the Ground Choice Model

To understand the best achievable performance, we also check the out-of-sample log-likelihoods when we fit a ranking-based choice model to the training data, which is the ground choice model that drives the choices of the customers in the training and hold-out data. We fit two versions of the ground choice model. In the first version of the ground choice model, we estimate both the ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$ and the corresponding probabilities $(\beta^1, \ldots, \beta^m)$. We refer to this version of the fitted ground choice model as UR, standing for unknown ranked lists. The paper by van Ryzin and Vulcano (2015) gives an algorithm to estimate both the ranked lists and the probabilities from the training data. We use a cross-validation approach to choose the appropriate number of ranked lists in the ground choice model. In our cross-validation approach, we split the training data into two parts, where we use three-fourths of the training data to fit the ground choice model and one-fourth of the training data to choose the appropriate number of ranked lists. In particular, we gradually increase the number of ranked lists m one by one. For each value of m, we use the so-called market discovery algorithm proposed by van Ryzin and Vulcano (2015) on the three-fourths of the training data to come up with m ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$ and their corresponding probabilities $(\beta^1, \ldots, \beta^m)$. Once we fit the ground choice model with a particular value of m, we use the fitted

(n, m)	Trn. data	Out-of-sample log-likelihood EM	UR	KR	EM-UR
$(11, 21)$	2,500	-16,391	-16,717	-16,196	1.96%
$(11, 21)$	5,000	-16,387	-16,531	-16,183	0.87%
$(11, 21)$	10,000	-16,366	-16,366	-16,173	0.00%
$(11, 31)$	2,500	-16,199	-16,251	-16,129	0.32%
$(11, 31)$	5,000	-16,154	-16,226	-16,081	0.44%
$(11, 31)$	10,000	-16,141	-16,136	-16,066	-0.03%
$(11, 51)$	2,500	-17,029	-17,199	-16,958	0.99%
$(11, 51)$	5,000	-17,017	-17,105	-16,943	0.51%
$(11, 51)$	10,000	-16,992	-16,993	-16,930	0.01%

(n, m)	Trn. data	Out-of-sample log-likelihood EM	UR	KR	EM-UR
$(21, 31)$	2,500	-22,839	-24,001	-22,387	4.84%
$(21, 31)$	5,000	-22,737	-23,422	-22,341	2.93%
$(21, 31)$	10,000	-22,673	-22,822	-22,329	0.65%
$(21, 41)$	2,500	-22,581	-24,581	-22,179	8.14%
$(21, 41)$	5,000	-22,483	-23,006	-22,139	2.27%
$(21, 41)$	10,000	-22,422	-22,767	-22,115	1.51%
$(21, 61)$	2,500	-23,406	-23,867	-23,141	1.93%
$(21, 61)$	5,000	-23,301	-23,469	-23,103	0.71%
$(21, 61)$	10,000	-23,267	-23,425	-23,060	0.67%

Table 5: Out-of-sample log-likelihoods obtained by EM, UR and KR.

ground choice model with the particular value of m to compute the log-likelihood of the remaining one-fourth of the training data. We choose the value of m that provides the largest log-likelihood on one-fourth of the training data. The algorithm proposed by van Ryzin and Vulcano (2015) solves a sequence of integer programs to come up with the ranked lists, so it becomes computationally prohibitive when there are too many customers in the training data. Therefore, we are not able to present comparisons for 50,000 customers in the training data. In the second version of the ground choice model, we assume that we know the ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$ and we estimate only the probabilities $(\beta^1, \ldots, \beta^m)$. The paper by van Ryzin and Vulcano (2016) gives an algorithm to estimate the probabilities from the training data when we know the ranked lists $\{\boldsymbol{\sigma}^g : g \in M\}$. We refer to this version of the fitted ground choice model as KR, standing for known ranked lists. In reality, it is highly unlikely to know the ranked lists of products that drive the choice process of the customers. Thus, KR represents a situation that is not achievable in practice.

In Table 5, we show the out-of-sample log-likelihoods obtained by EM, UR and KR, which are computed by using the testing data. The out-of-sample log-likelihoods obtained by EM are generally better than those obtained by UR. The percent gaps between the two log-likelihoods diminish as we have more training data, but the percent gaps between the log-likelihoods obtained EM and UR can still be as large 1.51% even when we have 10,000 customers. Furthermore, UR becomes computationally prohibitive when the number of customers in the training data exceeds 10,000. Also, as the number of products gets larger, the number of possible ranked lists gets larger and the percent gaps between the log-likelihoods from EM and UR get larger. Not surprisingly, since KR has access to the ranked lists of products in the ground choice model, the log-likelihoods obtained by KR are noticeably better than those obtained by EM and UR, but as discussed in the previous paragraph, the log-likelihoods obtained by KR are unlikely to be achievable in practice.

(n,m)	Trn. data	Root-mean-square error EM	DM	ML	EM-DM	EM-ML	(n,m)	Trn. data	Root-mean-square error EM	DM	ML	EM-DM	EM-ML
$(11,21)$	2,500	0.0334	0.0330	0.0523	-1.21%	36.14%	$(21,31)$	2,500	0.0275	0.0287	0.0344	4.18%	20.06%
$(11,21)$	5,000	0.0326	0.0318	0.0518	-2.52%	37.07%	$(21,31)$	5,000	0.0247	0.0243	0.0341	-1.65%	27.57%
$(11,21)$	10,000	0.0309	0.0286	0.0513	-8.04%	39.77%	$(21,31)$	10,000	0.0232	0.0222	0.0339	-4.50%	31.56%
$(11,21)$	50,000	0.0297	0.0265	0.0510	-12.08%	41.76%	$(21,31)$	50,000	0.0218	0.0202	0.0336	-7.92%	35.12%
$(11,31)$	2,500	0.0274	0.0295	0.0368	7.12%	25.54%	$(21,41)$	2,500	0.0257	0.0267	0.0317	3.75%	18.93%
$(11,31)$	5,000	0.0233	0.0235	0.0363	0.85%	35.81%	$(21,41)$	5,000	0.0235	0.0236	0.0311	0.42%	24.44%
$(11,31)$	10,000	0.0218	0.0202	0.0361	-7.92%	39.61%	$(21,41)$	10,000	0.0215	0.0209	0.0309	-2.87%	30.42%
$(11,31)$	50,000	0.0207	0.0170	0.0358	-21.76%	42.18%	$(21,41)$	50,000	0.0200	0.0179	0.0307	-11.73%	34.85%
$(11,51)$	2,500	0.0274	0.0302	0.0366	9.27%	25.14%	$(21,61)$	2,500	0.0239	0.0264	0.0248	9.47%	3.63%
$(11,51)$	5,000	0.0239	0.0246	0.0354	2.85%	32.49%	$(21,61)$	5,000	0.0197	0.0208	0.0240	5.29%	17.92%
$(11,51)$	10,000	0.0223	0.0218	0.0352	-2.29%	36.65%	$(21,61)$	10,000	0.0179	0.0185	0.0232	3.24%	22.84%
$(11,51)$	50,000	0.0207	0.0180	0.0347	-15.00%	40.35%	$(21,61)$	50,000	0.0162	0.0144	0.0230	-12.50%	29.57%

Table 6: Root-mean-square errors obtained by EM, DM and ML.

H Online Appendix: Comparison of Root-Mean-Square Errors

We compare EM, DM and ML from the perspective of the root-mean-square error between the fitted choice probabilities and the actual choice probabilities. Given that the customers choose under a choice model Mod, we use $P_i^{\mathsf{Mod}}(S)$ to denote the purchase probability of product i when we offer the subset S. The root-mean-square error for EM is given by $(\sum_{t\in T}\sum_{i\in\hat{S}^t}(P_i^{\mathsf{GR}}(\hat{S}^t)-P_i^{\mathsf{MC}}(\hat{S}^t))^2/\sum_{t\in T}|\hat{S}^t|)^{1/2}$, where GR is the actual ground choice model, MC is the Markov chain choice model whose parameters are estimated by EM and $\{\hat{S}^t : t\in T\}$ are the subsets in the hold-out data. The root-mean-square errors for DM and ML are defined similarly. Thus, the root-mean-square error indicates how closely a fitted choice model tracks the ground choice model. We show our results in Table 6. The layout of this table is identical to that of Table 1, except that we compare the root-mean-square errors in Table 6, whereas we compare the out-of-sample log-likelihoods in Table 1. The results in Table 6 indicate that EM performs noticeably better than ML. Nevertheless, if we have too few customers in the training data and too many products so that we estimate too many parameters from too little data, then the root-mean-square errors obtained by ML may be better than those obtained by EM. For example, if we have 1,000 customers in the training data and 21 products, so that EM estimates about 400 parameters from 1,000 data points, then the average percent gap between the root-mean-square errors obtained by EM and ML is -3.38%, favoring ML. It is difficult for EM to estimate 400 parameters from 1,000 data points. From a statistical viewpoint, expecting EM to be able to effectively estimate about 400 parameters from 1,000 data points is indeed unrealistic.

I Online Appendix: Comparison with the Independent Demand Model

In the independent demand model, each customer arrives into the system to purchase a particular product. If this product is available for purchase, then the customer purchases it. Otherwise, the customer leaves the system without a purchase. We use ID to refer to the benchmark that uses the independent demand model to capture the customer choices. We proceed to giving a brief formulation of the independent demand model. A customer arrives into the system to purchase product i with probability p_i. Naturally, we have $\sum_{i\in N} p_i = 1$. The parameters of the independent demand model are $\boldsymbol{p} = (p_1, \ldots, p_n)$. Recalling that we use $\phi \in N$ to denote the product that corresponds to the no-purchase option, if we offer the subset S of products, then the probability that a customer purchases product $i \in S \setminus \{\phi\}$ is p_i, whereas the probability that a customer purchases product ϕ is $p_\phi + \sum_{i\in N\setminus S} p_i$. In the last expression, we use the fact that if a customer arrives into the system to purchase one of the products in $N \setminus S$, then she leaves the system without a purchase. We construct the likelihood function to estimate the parameters of the independent demand model. Assume that we offer the subset $\hat{S}^t$ of products to customer t and the purchase decision of this customer is given by $\hat{\boldsymbol{Z}}^t = (\hat{Z}^t_1, \ldots, \hat{Z}^t_n)$, where $\hat{Z}^t_i = 1$ if and only if the customer purchases product i. The likelihood of the purchase decision of customer t is $\prod_{i\in N\setminus\{\phi\}} p_i^{\hat{Z}^t_i} \times (p_\phi + \sum_{i\in N\setminus S} p_i)^{\hat{Z}^t_\phi}$. So, the log-likelihood of this purchase decision is $\sum_{i\in N\setminus\{\phi\}} \hat{Z}^t_i \log p_i + \hat{Z}^t_\phi \log(p_\phi + \sum_{i\in N\setminus S} p_i)$; see Section 2.2 in Bishop (2006). The log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ is given by

$$L(\boldsymbol{p}) = \sum_{t\in T}\sum_{i\in N} \hat{Z}^t_i \log p_i + \sum_{t\in T} \hat{Z}^t_\phi \log\Big(p_\phi + \sum_{i\in N\setminus S} p_i\Big), \tag{16}$$

which is concave in $\boldsymbol{p}$. In ID, we use the `Matlab` routine `fmincon` to maximize the log-likelihood function above subject to the constraint that $\sum_{i\in N} p_i = 1$ and $\boldsymbol{p} \in \Re^n_+$.

In Tables 7 and 8, we respectively compare the out-of-sample log-likelihoods and AIC obtained by EM, ML and ID. The layouts of these tables are identical to those of Tables 1 and 2. Our results indicate that both EM and ML perform significantly better than ID.

J Online Appendix: Comparison of Root-Mean-Square Errors for Hotel Data

Using the same notation in Appendix H and following van Ryzin and Vulcano (2016), we define the root-mean-square error for EM as $(\sum_{t\in T}\sum_{i\in\hat{S}_t} (P^{\mathsf{MC}}_i(\hat{S}_t) - \mathbf{1}(\hat{Z}^t_i = 1))^2 / \sum_{t\in T} |\hat{S}^t|)^{1/2}$, where

(n,m)	Trn. data	Out-of-sample log-likelihood			EM-ID	ML-ID	(n,m)	Trn. data	Out-of-sample log-likelihood			EM-ID	ML-ID
		EM	ML	ID					EM	ML	ID		
$(11,21)$	2,500	-16,391	-16,677	-18,697	12.34%	10.80%	$(21,31)$	2,500	-22,839	-23,139	-27,507	16.97%	15.88%
$(11,21)$	5,000	-16,387	-16,665	-18,686	12.30%	10.81%	$(21,31)$	5,000	-22,737	-23,116	-27,489	17.29%	15.91%
$(11,21)$	10,000	-16,366	-16,663	-18,683	12.40%	10.81%	$(21,31)$	10,000	-22,673	-23,100	-27,470	17.46%	15.91%
$(11,21)$	50,000	-16,361	-16,663	-18,684	12.43%	10.82%	$(21,31)$	50,000	-22,627	-23,088	-27,460	17.60%	15.92%
$(11,31)$	2,500	-16,199	-16,298	-17,412	6.97%	6.40%	$(21,41)$	2,500	-22,581	-22,807	-27,564	18.08%	17.26%
$(11,31)$	5,000	-16,154	-16,286	-17,400	7.16%	6.40%	$(21,41)$	5,000	-22,483	-22,781	-27,534	18.34%	17.26%
$(11,31)$	10,000	-16,141	-16,282	-17,396	7.22%	6.40%	$(21,41)$	10,000	-22,422	-22,773	-27,526	18.54%	17.27%
$(11,31)$	50,000	-16,132	-16,283	-17,397	7.27%	6.40%	$(21,41)$	50,000	-22,378	-22,763	-27,519	18.68%	17.28%
$(11,51)$	2,500	-17,029	-17,116	-19,697	13.54%	13.10%	$(21,61)$	2,500	-23,406	-23,439	-27,817	15.86%	15.74%
$(11,51)$	5,000	-17,017	-17,111	-19,702	13.63%	13.15%	$(21,61)$	5,000	-23,301	-23,426	-27,808	16.21%	15.76%
$(11,51)$	10,000	-16,992	-17,107	-19,667	13.60%	13.02%	$(21,61)$	10,000	-23,267	-23,406	-27,793	16.28%	15.79%
$(11,51)$	50,000	-16,986	-17,104	-19,667	13.63%	13.03%	$(21,61)$	50,000	-23,226	-23,399	-27,787	16.41%	15.79%

Table 7: Out-of-sample log-likelihoods obtained by EM, ML and ID.

(n,m)	Trn. data	Akaike information criterion			EM-ID	ML-ID	(n,m)	Trn. data	Akaike information criterion			EM-ID	ML-ID
		EM	ML	ID					EM	ML	ID		
$(11,21)$	2,500	8,386	8,378	9,446	11.23%	11.31%	$(21,31)$	2,500	12,046	11,541	13,802	12.73%	16.38%
$(11,21)$	5,000	16,498	16,670	18,771	12.11%	11.20%	$(21,31)$	5,000	23,350	23,051	27,461	14.97%	16.06%
$(11,21)$	10,000	32,856	33,333	37,511	12.41%	11.14%	$(21,31)$	10,000	45,974	46,134	54,879	16.23%	15.93%
$(11,21)$	50,000	163,969	167,015	187,024	12.33%	10.70%	$(21,31)$	50,000	227,174	231,152	274,858	17.35%	15.90%
$(11,31)$	2,500	8,304	8,212	8,768	5.30%	6.34%	$(21,41)$	2,500	11,909	11,392	13,772	13.53%	17.29%
$(11,31)$	5,000	16,372	16,378	17,517	6.54%	6.50%	$(21,41)$	5,000	23,149	22,821	27,516	15.87%	17.06%
$(11,31)$	10,000	32,687	32,838	35,257	7.29%	6.86%	$(21,41)$	10,000	45,548	45,601	55,068	17.29%	17.19%
$(11,31)$	50,000	161,721	163,148	174,518	7.33%	6.52%	$(21,41)$	50,000	224,711	227,649	275,840	18.54%	17.47%
$(11,51)$	2,500	8,616	8,528	9,716	11.32%	12.23%	$(21,61)$	2,500	12,275	11,687	13,885	11.60%	15.83%
$(11,51)$	5,000	17,023	17,014	19,596	13.13%	13.18%	$(21,61)$	5,000	23,908	23,380	27,711	13.72%	15.63%
$(11,51)$	10,000	33,887	34,081	39,365	13.92%	13.42%	$(21,61)$	10,000	47,080	46,779	55,602	15.33%	15.87%
$(11,51)$	50,000	169,231	170,698	196,963	14.08%	13.33%	$(21,61)$	50,000	232,667	233,751	278,461	16.45%	16.06%

Table 8: AIC obtained by EM, ML and ID.

$\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ is the hold-out-data and MC is the Markov chain choice model whose parameters are estimated by EM. The root-mean-square errors for DM and ML are defined similarly. In Table 9, we show the root-mean-square errors for EM, DM and ML, where the results correspond to the root-mean-square errors averaged over the five folds. DM provides slightly better root-mean-square errors than EM. EM provides better root-mean-square errors than ML.

K Online Appendix: Extension to Censored Demands

Demand censorship refers to the fact that if there is no purchase for a product during a certain duration of time, then the data does not make it clear whether there was no customer arrival during this duration of time or the customers arriving during this duration of time decided to purchase nothing. As pointed out by Vulcano et al. (2012) and van Ryzin and Vulcano (2016),

Root-mean-square error					
	Hotel 1	Hotel 2	Hotel 3	Hotel 4	Hotel 5
EM	0.2175	0.2551	0.2191	0.2802	0.2375
DM	0.2174	0.2550	0.2190	0.2801	0.2374
ML	0.2195	0.2567	0.2202	0.2812	0.2382
EM-DM	-0.03%	-0.02%	-0.01%	-0.02%	-0.02%
EM-ML	0.95%	0.63%	0.53%	0.38%	0.30%

Table 9: Root-mean-square errors obtained by EM, DM and ML on the hotel data.

demand censorship creates another missing piece in the data, since we do not know whether the absence of purchase is due to the fact that there was no customer arrival or the arriving customers did not purchase anything. Our goal in this section is to show how we can extend our expectation-maximization algorithm to deal with demand censorship.

We use one of the products $\phi \in N$ to denote the no-purchase option, which is always available to the customers. If a customer visits product ϕ during the course of her choice process, then she leaves the system without a purchase. Under censored demands, we make purchase observations at multiple time periods. These time periods correspond to small enough intervals of time that there is at most one customer arrival at each time period. At each time period, a customer arrives into the system with probability α. An arriving customer is interested in purchasing product i with probability λ_i. For notational uniformity, we do not restrict λ_ϕ to take value zero so that an arriving customer may be interested in the no-purchase option, in which case, she immediately leaves without a purchase. A customer visiting product i and finding product i unavailable transitions from product i to product j with probability $\rho_{i,j}$. Note that $\rho_{i,\phi}$ is the probability of transitioning from product i to the no-purchase option. Using $(\Theta_1(S), \ldots, \Theta_n(S))$ to denote the solution to (1), if we offer the subset S of products, then an arriving customer does not purchase anything with probability $\Theta_\phi(S)$. Our goal is to estimate the parameters $(\alpha, \boldsymbol{\lambda}, \boldsymbol{\rho})$.

We define the random variable $Z_i(S)$ such that $Z_i(S) = 1$ if and only if there is a purchase for product i at a time period when we offer the subset S of products. For all $i \in S \setminus \{\phi\}$, we have $Z_i(S) = 1$ with probability $\alpha\,\Theta_i(S)$, where we use the fact that we have a purchase for product $i \in S \setminus \{\phi\}$ at a time period if there is a customer arrival and the customer purchases product i. Note that $Z_\phi(S) = 1$ if and only if there is no purchase at a time period when we offer the subset S of products. Therefore, we have $Z_\phi(S) = 1$ with probability $1 - \alpha + \alpha\,\Theta_\phi(S)$, where we use the fact that we have no purchase at a time period if there is no customer arrival or there is a customer arrival and the arriving customer does not purchase anything.

K.1 Incomplete and Complete Likelihood Functions

In the data that we have available to estimate the parameters of the Markov chain choice model, there are τ time periods indexed by $T = \{1, \ldots, \tau\}$. We use $\hat{S}^t \subset N$ to denote the subset of products offered at time period t. We define $\hat{Z}_i^t \in \{0, 1\}$ such that $\hat{Z}_i^t = 1$ if and only if there is a purchase for product i at time period t. Note that $\hat{Z}_\phi^t = 1$ if and only if there is no purchase at time period t. If we have $\hat{Z}_\phi^t = 1$, then we either had no customer arrival at time period t or the customer arriving at time period t did not purchase anything, but we do not know which one of these events occurred. The data that is available to estimate the parameters of the Markov chain choice model is given by $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$. We emphasize that the set T indexes the time periods in the censored demand case, whereas the set T indexes the customers in the uncensored demand case. Therefore, our interpretation of the data is slightly different when we deal with demand censorship. Our setup for the censored demand case closely follows the one in van Ryzin and Vulcano (2016).

Considering the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ that is available to estimate the parameters of the Markov chain choice model, the probability that there is a purchase for product $i \in \hat{S}^t \setminus \{\phi\}$ at time period t is given by $\alpha\, \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$, where we explicitly show that the solution $(\Theta_1(S \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}), \ldots, \Theta_n(S \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}))$ to (1) depends on the parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ of the Markov chain choice model. On the other hand, the probability that there is no purchase at time period t is given by $1 - \alpha + \alpha\, \Theta_\phi(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$, where we use the fact that there is no purchase at time period t if there is no customer arrival or there is a customer arrival and the arriving customer does not purchase anything. Therefore, the likelihood of the observation at time period t is given by $\prod_{i \in N \setminus \{\phi\}} (\alpha\, \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}))^{\hat{Z}_i^t} \times (1 - \alpha + \alpha\, \Theta_\phi(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}))^{\hat{Z}_\phi^t}$. So, the log-likelihood of this observation is $\sum_{i \in N \setminus \{\phi\}} \hat{Z}_i^t \log \alpha + \sum_{i \in N \setminus \{\phi\}} \hat{Z}_i^t \log \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}) + \hat{Z}_\phi^t \log(1 - \alpha + \alpha\, \Theta_\phi(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}))$. In this case, the log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t) : t \in T\}$ is given by

$$L_I(\alpha, \boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N \setminus \{\phi\}} \hat{Z}_i^t \log \alpha + \sum_{t \in T} \sum_{i \in N \setminus \{\phi\}} \hat{Z}_i^t \log \Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho}) + \sum_{t \in T} \hat{Z}_\phi^t \log(1 - \alpha + \alpha\, \Theta_\phi(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})).$$

Similar to the uncensored demand case, maximizing the log-likelihood function above directly is difficult since we do not have a closed-form expression for $\Theta_i(\hat{S}^t \,|\, \boldsymbol{\lambda}, \boldsymbol{\rho})$. Therefore, we rely on the expectation-maximization framework for the censored demand case as well. We define F_i, $X_{i,j}(S)$,

$\hat{F}_i^t$ and $\hat{X}_{i,j}^t$ as in Section 3. In particular, we define the random variable $F_i \in \{0,1\}$ such that $F_i = 1$ if and only if a customer arrives into the system at a time period with an interest in purchasing product i. Given that there is a customer arrival at a time period, we define the random variable $X_{i,j}(S)$ as the number of times the customer transitions from product i to product j during the course of her choice process when we offer the subset S of products. The definitions of $\hat{F}_i^t$ and $\hat{X}_{i,j}^t$ are analogous. In the censored demand case, using $\mathbf{0} \in \Re_+^n$ to denote the vector of all zeros, we may have $\boldsymbol{F} = (F_1, \ldots, F_n) = \mathbf{0}$, since there may be no customers arriving at a time period. In particular, we have $\boldsymbol{F} = \mathbf{0}$ with probability $1 - \alpha$. Also, we have $\boldsymbol{F} = \boldsymbol{e}_i$ with probability $\alpha \, \lambda_i$, since a customer arrives into the system at a time period with an interest in purchasing product i if there is a customer arrival at a time period and the arriving customer is interested in purchasing product i. If we have access to the data $(\hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t)$, then the likelihood of the observation at time period t is given by $\prod_{i \in N} (\alpha \lambda_i)^{\hat{F}_i^t} \prod_{i,j \in N} \rho_{i,j}^{\hat{X}_{i,j}^t} \times (1 - \alpha)^{1 - \sum_{i \in N} \hat{F}_i^t}$. So, the log-likelihood of this observation is $\sum_{i \in N} \hat{F}_i^t \log \alpha + \sum_{i \in N} \hat{F}_i^t \log \lambda_i + \sum_{i,j \in N} \hat{X}_{i,j}^t \log \rho_{i,j} + (1 - \sum_{i \in N} \hat{F}_i^t) \log(1 - \alpha)$. Therefore, the log-likelihood of the data $\{(\hat{S}^t, \hat{\boldsymbol{Z}}^t, \hat{\boldsymbol{F}}^t, \hat{\boldsymbol{X}}^t) : t \in T\}$ is given by

$$L_C(\alpha, \boldsymbol{\lambda}, \boldsymbol{\rho}) = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \log \alpha + \sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \log \lambda_i + \sum_{t \in T} \sum_{i,j \in N} \hat{X}_{i,j}^t \log \rho_{i,j} + \sum_{t \in T} \left(1 - \sum_{i \in N} \hat{F}_i^t \right) \log(1 - \alpha). \quad (17)$$

The log-likelihood function above has a closed-form expression and it is concave in $(\alpha, \boldsymbol{\lambda}, \boldsymbol{\rho})$. We note that another approach for dealing with demand censorship is to interpret λ_i as the probability that there is a customer arrival at a time period and the arriving customer is interested in purchasing product i. In this case, we need to assume that $\sum_{i \in N} \lambda_i \leq 1$ and there is no customer arrival at a time period with probability $1 - \sum_{i \in N} \lambda_i$. Therefore, we can replace $\alpha \, \lambda_i$ simply with λ_i and $1 - \alpha$ simply with $1 - \sum_{i \in N} \lambda_i$ in the likelihood function, eliminating the need for the parameter α. We keep the parameter α so that we explicitly estimate the arrival probability of a customer, in which case, our development becomes consistent with van Ryzin and Vulcano (2016).

K.2 Expectation Step

At iteration ℓ, in the expectation step of our expectation-maximization algorithm, we offer the subset $\hat{S}^t$ of products at time period t. A customer arrives into the system with probability α^ℓ

and chooses among these products according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}^\ell, \boldsymbol{\rho}^\ell)$. The purchase observation at this time period is given by the vector $\hat{\boldsymbol{Z}}^t$. Thus, for all $i \in \hat{S}^t \setminus \{\phi\}$, we have $\hat{Z}_i^t = 1$ if and only if there is a purchase for product i at time period t. We have $\hat{Z}_\phi^t = 1$ if and only if there is no purchase at time period t. We need to compute the conditional expectations $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$ and $\mathbb{E}\{X_{i,j}(\hat{S}_t) \,|\, \boldsymbol{Z}(\hat{S}^t) = \hat{\boldsymbol{Z}}^t\}$. For notational brevity, we omit the superscript t indexing the time period and the superscript ℓ indexing the iteration counter. In particular, we consider the case where we offer the subset S of products at a time period. A customer arrives into the system with probability α. She chooses among these products according to the Markov chain choice model with parameters $(\boldsymbol{\lambda}, \boldsymbol{\rho})$. We know the purchase observation at this time period. That is, noting that the purchase decision of a customer is captured by the random variable $\boldsymbol{Z}(S) = (Z_1(S), \ldots, Z_n(S))$, we know that $\boldsymbol{Z}(S) = \boldsymbol{e}_k$ for some $k \in N$. We want to compute the conditional expectations $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$.

The difference between the uncensored and censored demand cases arises only when $k = \phi$, which corresponds to the case where we do not observe a purchase in the data. If $k \in S \setminus \{\phi\}$, then we can compute $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ and $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ by using (6) and (10) without any modifications. In the rest of this section, we focus on the case where $k = \phi$. Considering the conditional expectation $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$, we use the Bayes rule to obtain

$$\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\} = \mathbb{P}\{F_i = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\} = \frac{\mathbb{P}\{Z_\phi(S) = 1 \,|\, F_i = 1\} \, \mathbb{P}\{F_i = 1\}}{\mathbb{P}\{Z_\phi(S) = 1\}}.$$

On the right side above, $\mathbb{P}\{Z_\phi(S) = 1 \,|\, F_i = 1\}$ corresponds to the probability that we do not have a purchase when we offer the subset S of products conditional on the fact that a customer arrives with an interest in purchasing product i. This conditional probability is still given by $\Psi_\phi(i, S)$, where $\{\Psi_\phi(i, S) : i \in N \setminus S\}$ is the solution to (5) after replacing k with ϕ. We recall that we also have $\Psi_\phi(i, S) = 1(i = \phi)$ for all $i \in S$. Also, $\mathbb{P}\{F_i = 1\}$ corresponds to the probability that a customer arrives into the system with an interest in purchasing product i. So, we have $\mathbb{P}\{F_i = 1\} = \alpha \, \lambda_i$. Finally, $\mathbb{P}\{Z_\phi(S) = 1\}$ corresponds to the probability that there is no purchase. Thus, we have $\mathbb{P}\{Z_\phi(S) = 1\} = 1 - \alpha + \alpha \, \Theta_\phi(S)$. In this case, using the equality above, we obtain $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\} = \Psi_\phi(i, S) \, \alpha \, \lambda_i / (1 - \alpha + \alpha \, \Theta_\phi(S))$, which shows how to compute the conditional expectation $\mathbb{E}\{F_i \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ under censored demands. Next, we consider the conditional expectation $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\}$. Similar to our earlier notation, we define the random variable $Y_i^m(S) \in \{0, 1\}$ such that $Y_i^m(S) = 1$ if and only if the m-th product that a customer visits during the course of her choice process is product i when we offer the subset S of products. Note that $Y_\phi^m(S)$ is well-defined. The chain of equalities in (7) holds without any modifications after replacing k with ϕ. We modify the chain of equalities in (8) as

$$\mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\} = \frac{\mathbb{P}\{Z_\phi(S) = 1 \,|\, Y_i^m(S) = 1\}\,\mathbb{P}\{Y_i^m(S) = 1\}}{\mathbb{P}\{Z_\phi(S) = 1\}}$$
$$= \frac{\mathbb{P}\{Z_\phi(S) = 1 \,|\, F_i = 1\}\,\mathbb{P}\{Y_i^m(S) = 1\}}{\mathbb{P}\{Z_\phi(S) = 1\}} = \frac{\Psi_\phi(i,S)\,\mathbb{P}\{Y_i^m(S) = 1\}}{1 - \alpha + \alpha\,\Theta_\phi(S)}.$$

The chain of equalities in (9) holds without any modifications after replacing k with ϕ so that we have $\mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi, Y_i^m(S) = 1\} = \Psi_\phi(j,S)\,\rho_{i,j}/\Psi_\phi(i,S)$. Using this equality and the chain of equalities above in (7), after replacing k in (7) with ϕ, we obtain

$$\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\} = \sum_{m=1}^{\infty} \mathbb{P}\{Y_j^{m+1}(S) = 1 \,|\, Y_i^m(S) = 1, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\}\,\mathbb{P}\{Y_i^m(S) = 1 \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_\phi\}$$
$$= \sum_{m=1}^{\infty} \frac{\Psi_\phi(j,S)\,\rho_{i,j}}{\Psi_\phi(i,S)}\,\frac{\Psi_\phi(i,S)\,\mathbb{P}\{Y_i^m(S) = 1\}}{1 - \alpha + \alpha\,\Theta_\phi(S)}$$
$$= \frac{\Psi_\phi(j,S)\,\rho_{i,j}}{1 - \alpha + \alpha\,\Theta_\phi(S)} \sum_{m=1}^{\infty} \mathbb{P}\{Y_i^m(S) = 1\} = \frac{\Psi_\phi(j,S)\,\rho_{i,j}}{1 - \alpha + \alpha\,\Theta_\phi(S)}\,\alpha\,\Theta_i(S).$$

In the last equality above, $\sum_{m=1}^{\infty} \mathbb{P}\{Y_i^m(S) = 1\}$ is the expected number of times a customer visits product i given that we offer the subset S of products. Conditional on a customer arrival, $\Theta_i(S)$ is the expected number of times a customer visits product i. So, the expected number of times a customer visits product i is $\alpha\,\Theta_i(S)$. The last chain of equalities above shows how we can compute the conditional expectation $\mathbb{E}\{X_{i,j}(S) \,|\, \boldsymbol{Z}(S) = \boldsymbol{e}_k\}$ under censored demands. Thus, the discussion in this section shows how to execute the expectation step under censored demands.

K.3 Maximization Step

In the maximization step of our expectation-maximization algorithm, we need to maximize the log-likelihood function $L_C(\alpha, \boldsymbol{\lambda}, \boldsymbol{\rho})$ in (17) subject to the constraint that $\alpha \leq 1$, $\sum_{i \in N} \lambda_i = 1$, $\sum_{j \in N} \rho_{i,j} = 1$ for all $i \in N$, $\alpha \in \Re_+$, $\boldsymbol{\lambda} \in \Re_+^n$ and $\boldsymbol{\rho} \in \Re_+^{n \times n}$. This optimization problem decomposes into three optimization problems, where each one of the three problems involves the decision variables α, $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$. The problems that involve the decision variables $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$ are identical to problems (11) and (12). Therefore, we can use the discussion in Section 4.3 to compute the optimal values of the decision variables $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$. On the other hand, the problem that involves the decision variable α is of the form

$$\max \left\{ \sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \log \alpha + \sum_{t \in T} \left(1 - \sum_{i \in N} \hat{F}_i^t \right) \log (1 - \alpha) \,:\, 0 \leq \alpha \leq 1 \right\}.$$

The problem above corresponds to the problem of computing the maximum likelihood estimator of the success probability of a Bernoulli random variable, where we have a total of $|T|$ trials and we observe success in $\sum_{t \in T} \sum_{i \in N} \hat{F}_i^t$ trials. In this case, it is known that the maximum likelihood estimator of the success probability is given by $\sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \,/\, |T|$; see Section 2.1 in Bishop (2006). It is also possible to find the optimal solution to the problem above directly. The objective function above is concave in α. We write the objective function as $\sum_{t \in T} \sum_{i \in N} \hat{F}_i^t \, (\log \alpha - \log (1 - \alpha)) + |T| \, \log (1 - \alpha)$. Ignoring the constraint $0 \leq \alpha \leq 1$ for the moment, the first-order condition for the unconstrained problem shows that the optimal solution is obtained by setting $\alpha = \sum_{t \in T} \sum_{i \in N} \hat{F}_i^t / |T|$. Since $\sum_{i \in N} \hat{F}_i^t \leq 1$, we have $\sum_{t \in T} \sum_{i \in N} \hat{F}_i^t / |T| \leq 1$, which implies that the solution to the unconstrained problem is optimal to the problem above as well. Thus, the discussion in this section shows how to execute the maximization step under censored demands.