



Management Science

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

A Partial-Order-Based Model to Estimate Individual Preferences Using Panel Data

Srikanth Jagabathula, Gustavo Vulcano

To cite this article:

Srikanth Jagabathula, Gustavo Vulcano (2017) A Partial-Order-Based Model to Estimate Individual Preferences Using Panel Data. Management Science

Published online in Articles in Advance 24 Apr 2017

<http://dx.doi.org/10.1287/mnsc.2016.2683>

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

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

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

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



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

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

A Partial-Order-Based Model to Estimate Individual Preferences Using Panel Data

Srikanth Jagabathula,^a Gustavo Vulcano^{a,b,c}

^aLeonard N. Stern School of Business, New York University, New York, New York 10012; ^bSchool of Business, Universidad Torcuato di Tella, Buenos Aires C1428BCW, Argentina; ^cConsejo Nacional de Investigaciones Científicas y Técnicas, Argentina

Contact: sjagabat@stern.nyu.edu (SJ); gvulcano@stern.nyu.edu (GV)

Received: February 5, 2015

Accepted: August 23, 2016

Published Online in Articles in Advance:
April 24, 2017

<https://doi.org/10.1287/mnsc.2016.2683>

Copyright: © 2017 INFORMS

Abstract. In retail operations, customer choices may be affected by stockout and promotion events. Given panel data with the transaction history of customers, and product availability and promotion data, our goal is to predict future individual purchases. We use a general nonparametric framework in which we represent customers by partial orders of preferences. In each store visit, each customer samples a full preference list of the products consistent with her partial order, forms a consideration set, and then chooses to purchase the most preferred product among the considered ones. Our approach involves: (a) defining behavioral models to build consideration sets as subsets of the products on offer, (b) proposing a clustering algorithm for determining customer segments, and (c) deriving marginal distributions for partial preferences under the multinomial logit model. Numerical experiments on real-world panel data show that our approach allows more accurate, fine-grained predictions for individual purchase behavior compared to state-of-the-art alternative methods.

History: Accepted by Vishal Gaur, operations management.

Funding: This research was supported in part by the National Science Foundation [Grant CMMI 1454310].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2016.2683>.

Keywords: nonparametric choice models • inertia in choice • brand loyalty • panel data • partial-orders • customized promotions • personalized predictions

1. Introduction

Demand estimates are key inputs for inventory control and price optimization models used in retail operations and revenue management (RM).¹ In the last decade, there has been a trend of switching from independent demand models to choice-based models of demand in both academia and industry practice. For simplicity, the traditional approach assumed that each product has its own independent stream of demand. However, if products are substitutes and their availabilities vary over time, then the demand for each product will be a function of the set of alternatives available to consumers when they make their purchase decisions, so that ignoring stockout and substitution effects that occur over time can introduce biases in the estimation process.

The building block for estimating customer choice is the specification of a choice model, either parametric or nonparametric. Most of the proposals in the operations-related literature have been on the former for both estimation and assortment optimization (e.g., Kök and Fisher 2007, Musalem et al. 2010). By parametric, we mean that the number of parameters that describe the family of underlying distributions is fixed and independent of the training data set volume. The parameterized model structure relates product

attributes to the utility values or choice probabilities of the different options. The greatest advantage of parametric models is the ability to include covariates, such as product features and price, that can help explain consumer preferences for alternatives. This also enables parametric models like the multinomial logit (MNL) or nested logit (NL) with linear-in-parameters utilities to extrapolate choice predictions to new alternatives that have not been observed in historical data and to predict how changes in product attributes such as price affect choice outcomes. Yet, the drawback of any parametric model is that one must make assumptions about the structure of preferences and the relevant covariates that influence it, which requires expert judgment and trial-and-error testing to determine an appropriate specification.

Recently, with the rise of business analytics and data-driven approaches, there has been a growing interest in revenue predictions and demand estimation derived from nonparametric choice models (e.g., Farias et al. 2013, Haensel and Koole 2011, van Ryzin and Vulcano 2014) that provide inputs for various optimization problems (see discussion in Section 1.2). These operations-related, nonparametric proposals specify customer types defined by their rank ordering of all alternatives (along with the no-purchase alternative).²

When faced with a choice from an offer set, a customer is assumed to purchase the available product that ranks highest in her preference list—or to quit without purchasing. The flexibility of this nonparametric choice model comes at a price: The potential number of preference lists (customer types) in this rank-based choice model is factorial in the number of alternatives, which challenges the estimation procedure.

The empirical evidence both in academia and industry practices strongly sustains choice-based demand models over the independent demand assumption (e.g., Ratliff et al. 2008, Newman et al. 2014, Vulcano et al. 2010) when the data source consists of product availability and sales transaction data. Yet, these models overlook two important features: First, they ignore the repeated interactions of a given customer with the firm and treat every transaction as coming from a different individual; and second, they usually assume that the items evaluated by a given customer in any store visit are all of the available ones within a category (or subcategory) of products, which likely overestimates the size of the true *consideration set*. The first limitation can be addressed by keeping track of the repeated interactions between a customer and the firm by tagging transactions with customer id. This information, popularly referred to as *panel data*, can be used subsequently to learn the preferences of individual customers. The canonical example is customers buying groceries on a weekly basis from a grocery retailer, but more broadly the setting includes any application in which customers exhibit loyalty through repeated purchases be it apparel, hotel, airline, etc. These types of data are very common in practical settings because of the proliferation of loyalty cards and other marketing programs, and can later be used to customize the offering (e.g., via personalized assortments and prices (Clifford 2012) or personalized mobile phone coupons (Danaher et al. 2015)). The limitation on the unobserved consideration set of the customer requires building a behavioral model of choice that captures individual bounded rationality.

In this paper, we contribute to the literature by proposing a nonparametric, choice-based demand framework that overcomes the two limitations discussed above, and that incorporates both operations- and marketing-related components. We infer individual customer preferences from panel data in a setting where (i) products are not always available (e.g., due to stockouts or deliberate scarcity introduced by the firm), (ii) preferences may be altered by price or display promotions, and (iii) customers exhibit bounded rationality in the sense that they cannot evaluate all of the products on offer, and their consideration sets are unobservable. For the purpose of estimation, the framework can accommodate both parametric and nonparametric models.

1.1. Summary of Results

We propose a choice-based demand model that is a generalization of the aforementioned rank-based model, which typically associates each customer with a fixed preference list that remains constant over time. In our model, we focus on a fixed set of m customers who visit a firm repeatedly and make purchases from a particular category or subcategory of products (e.g., coffee or shampoo). The full assortment is defined by a set of products, together with the no-purchase option. Each customer belongs to a market segment and is characterized by a set of partial preferences, represented by a directed acyclic graph (DAG) with products as nodes. The DAG captures partial preferences of the form “product a is preferred to product b ” through a directed edge from product a to product b but typically does not provide a full ranking of the products. Upon each store visit, the customer samples a full ranking consistent with her DAG according to a distribution specified for her particular market segment and chooses the product within her *consideration set* that ranks highest in her preference list. Of course, neither the customer’s DAG nor her sampled preference list nor her consideration set are observable. For each customer id, the firm only observes a collection of revealed preferences: the offered products at the moment of purchase and the chosen product. As a consequence, our estimation procedure involves three sequential phases: (1) building the DAG for each customer, (2) clustering the DAGs into a prespecified number of segments, and (3) estimating distributions over preference lists that best explain the purchasing patterns of the customers.

The first phase requires a key element of our work: the modeling of behavioral biases of each individual customer. Existing literature has established several behavioral biases that are common in customer choice behavior. We focus on one such bias that is relevant to the construction of consideration sets; namely, *inertia in choice*. Inertia in choice—also referred to as *short-term brand loyalty* by Jeuland (1979)—claims that when facing frequently purchased consumer goods, customers tend to stick to the same option rather than evaluate all products on offer in each store visit. We capture the effect of such inertia in choice by assuming that customers tend to purchase what they bought previously until there is a “trigger event” that forces them to consider other products on offer. We distinguish two important trigger events: stockout of the previously purchased product (which induces the customer to reevaluate all products on offer) and existence of product promotions (in particular, display and price promotions). Based on these behavioral assumptions, we define a set of *behavioral rules* that allow us to dynamically build the customer’s DAG as we track the sequence of her interactions with the firm and

she reveals her preferences through her purchasing patterns.

Next, to better capture customer heterogeneity, we cluster the m DAGs into K classes, where K is a predetermined small number (e.g., $K = 5$). We formulate this clustering problem as an integer program (IP) to systematically capture the idea that the DAGs of individual customers assigned to the same class are “close” to each other (according to a distance metric).

Finally, in the estimation phase, we calibrate a MNL model to assess probabilistic distributions over preference lists consistent with the DAGs of the first phase and the clustering of the second phase.

The predictive performance of our method is demonstrated through an exhaustive set of numerical experiments using real-world panel data on the purchase of 29 grocery categories across two big U.S. markets in year 2007. We divide our data set in two pieces. On the first part (i.e., the training data), we perform the three phases summarized above. Then, on the second part (i.e., the holdout sample), we predict what the customer would purchase when confronted with the offer sets and products on promotion, and compare with the recorded purchase. We observe that our method with behavioral rules and clustering obtains up to 40% improvement in prediction accuracy on standard metrics over state-of-the-art benchmarks based on variants of the MNL, commonly used in current industry practice.

1.2. Positioning in the Literature

Our work has several connections to the literature in both operations and marketing. While existing work in operations has largely focused on using transaction data aggregated over all customers to estimate choice models, literature in marketing has extensively used panel data to fit choice models. The latter extends back to the seminal work of Guadagni and Little (1983), in which they fit a MNL model to household panel data on the purchases of regular ground coffee, and which has paved the way for choice modeling in marketing using scanner panel data³; see Chandukala et al. (2008) and Wierenga (2008) for a detailed overview of choice modeling using panel data in the area of marketing. Much of this work focuses on understanding how various panel covariates influence individual choice making. For that, a model is specified that relates the panel covariates to a distribution over preference lists (induced, for instance, by the MNL model) that describes the preferences of each customer. The model is then estimated from choice data assuming that different choices of the same customer are characterized by independent draws of preference lists from the same distribution.

Our work tightens this assumption by allowing customers to be “partially consistent” in their preferences

across different choice instances, so that the draws of preference lists are no longer fully independent but are compatible with the customer’s DAG. This additional structure allows our model to capture individual preferences more accurately, especially when observations are sparse.

In the context of the operations-related literature, rank-based choice models of demand have been used as inputs in retail operations proposals, pioneered by the work of Mahajan and van Ryzin (2001), who analyze a single-period, stochastic inventory model in which a sequence of customers described by preference lists substitute among product variants within a retail assortment while inventory is depleted. Other recent retail operations papers dealing with variants of rank-based choice models include Rusmevichientong et al. (2006), Smith et al. (2009), Honhon et al. (2010, 2012), and Jagabathula and Rusmevichientong (2014). Rank-based models also reached the airline-related RM literature (e.g., Zhang and Cooper 2005, Chen and Homem-de Mello 2010, Chaneton and Vulcano 2011, Kunnumkal 2014). All of these proposals assume that data is obtained at the market level and do not account for the repeated interaction of a customer with a firm. Therefore, their applicability to making fine-grained, individual level predictions of purchasing patterns is unclear.

Finally, our work has rich methodological connections to the area that studies distributions over rankings in statistics and machine learning. We discuss these connections in Sections 3 and 5.

2. Model Description

We start this section with a description of the general modeling framework, where we introduce basic notation and formally define the DAGs that represent the customers. We continue with a discussion of the model assumptions, where we position our framework with respect to the traditional approaches toward customer choice. Next, we describe the type of data needed by our framework and finish with a detailed explanation of the DAGs’ construction process.

2.1. General Modeling Framework

We model consumer preferences using a general rank-based choice model of demand. The product universe \mathcal{N} consists of n products $\{a_1, a_2, \dots, a_n\}$. The “no-purchase” or “outside” alternative is denoted by 0. Preferences over the universe of products are captured by an antireflexive, antisymmetric, and transitive relation $>$, which induces a total ordering or ranking over all of the products in the universe, and we write $a > b$ to mean “ a is preferred to b .” Preferences can also be represented through rankings or permutations. Each preference list σ specifies a rank ordering over the $n + 1$ products in the universe $\mathcal{N} \cup \{0\}$, with $\sigma(a)$ denoting

the preference rank of product a . Lower ranks indicate higher preference so that a customer described by σ prefers product a to product b if and only if $\sigma(a) < \sigma(b)$, or equivalently, if $a >_{\sigma} b$.

The population consists of m customers who make purchases over T discrete time periods. We assume that the set of customers and the set of products remain constant over time. Each customer i is described by a general partial order D_i . A general partial order specifies a collection of pairwise preference relations, $D_i \subset \{(a_j, a_{j'}) : 0 \leq j, j' \leq n, j \neq j'\}$, so that product a_j is preferred to product $a_{j'}$ for any $(a_j, a_{j'}) \in D_i$. In addition, the customer belongs to one of K segments in the market, where segment k is characterized by a probability distribution λ_k over preference lists. In every interaction with the firm, the customer will sample a full ranking consistent with D_i according to the distribution λ_k . We say that a preference list σ is consistent with partial order D_i if and only if $\sigma(a_j) < \sigma(a_{j'})$ for each $(a_j, a_{j'}) \in D_i$. The partial order D_i could be empty (i.e., have no arcs), in which case the consistency requirement is vacuous.

To illustrate, suppose $n = 3$ and take a customer of type k with partial order $D = \{(1, 2), (3, 2)\}$. For ease of exposition, we ignore here the no-purchase alternative.⁴ There are two possible rankings consistent with D in this case: $1 > 3 > 2$ and $3 > 1 > 2$. The distribution λ_k specifies the point probabilities for each of the six possible rankings. Now, since $D = \{(1, 2), (3, 2)\}$, the customer consistently prefers both 1 and 3 over 2 in every choice instance. Hence, conditioned on D , she samples the preference list $1 > 3 > 2$ or $3 > 1 > 2$, say with probabilities 0.6 and 0.4, respectively.

The choice process proceeds as follows. In each purchase or choice instance, customer i of type k is offered a subset S of products. She focuses on a consideration set $C \subset S \cup \{0\}$, samples a preference list σ according to distribution λ_k from her collection of rankings consistent with D_i , and purchases from C the most preferred product according to σ —i.e., $\arg \min_{a_j \in C} \sigma(a_j)$. Different choice instances for customer i are independent but always consistent with her own collection of preference lists described by her partial order D_i . Thus, in the previous example, and assuming that the customer considers all products on offer, if the subset $S = \{1, 2\}$ is offered, the customer chooses product 1 for sure. If $S = \{1, 2, 3\}$ is offered, then the customer purchases 1 with probability 0.6 (if $1 > 3 > 2$ is sampled) and 3 with probability 0.4 (if $3 > 1 > 2$ is sampled).

The model described above is quite general and requires further restrictions to be estimable from data. The model is specified by the partial order D_i for each customer i , the number of customer classes K , the class membership of each of the customers, and the distribution λ_k over preference lists for each class k . The partial orders are built dynamically as we keep track of the store visits of each customer recorded in the training

data set, as explained later in Section 2.4. Their construction requires a behavioral model relating the observed choices to underlying preferences. Next, for an input parameter K , we determine the class membership of each of the customers through a clustering technique described in Section 4. Broadly speaking, the clustering technique clusters together customers with similar partial preferences. The distributions λ_k can in principle be induced by any of the commonly used choice models. In this paper, we focus on the most commonly used model: the MNL model, which has been extensively applied in operations, marketing, economics, and transportation science. As shown in Section 3, the MNL model allows for tractable (or approximately tractable) estimation from general partial orders.

2.2. Discussion of Model Assumptions

The inference of the customers' partial orders assumes that both the population of customers and the product universe remain constant along the horizon T . Our approach can be repeated from time to time (e.g., once a year) to collect new data points and update both the customer base and their preferences. In the meantime, new items in the category could be considered within the scope of an existing product, which is our minimal level of aggregation in the analysis to establish preferences.⁵

The potential benefit of the DAGs is the boost in accuracy it could provide for fine-grained individual-level predictions. These gains could be particularly significant for cases in which only a small sample of information is available for each customer, and customers are endowed with consistent preferences, like brand-loyal customers. In such cases, it is impractical to fit a choice model separately to each customer. For this reason, one must make distributional assumptions relating the preferences of the different customers, while retaining sufficient customer-level heterogeneity. A common approach is to assume that customers belong to K classes/segments and all customers in class k are homogeneous in the sense that they sample preferences from the same distribution λ_k , though the distributions are different across different classes. For instance, all customers of class k may sample preferences from the same MNL distribution, but the MNL parameters are different across different classes.

A challenge in implementing this approach is setting the appropriate value for the number of segments K . On the one hand, smaller values of K result in parsimonious models that can be efficiently fit to available data. On the other hand, larger values of K can better capture customer-level heterogeneity. Our model avoids this difficulty by providing an additional lever to capture heterogeneity while retaining parsimony. Specifically, even if customers are not segmented and are assumed to sample preferences from the same distribution (e.g., a single-class MNL), our model captures

customer-level heterogeneity by allowing the partial orders to differ across customers. When computing the probability of customer i choosing a particular option, we use the demand estimates but also condition on the customer being characterized by D_i , which imposes additional structure (and constraints) on the utility drawn for the different products. The DAGs are inferred from small sample data and typically require significantly less computational effort than the estimation of a multiclass demand model. Recall that our model is quite general in the sense that if the DAGs are empty (i.e., the nodes are isolated), no structure is superimposed and we recover a standard, unconstrained, choice-based demand model (e.g., a typical, single-class MNL).

On the other hand, we expect the gains to be subdued or nonexistent for customers who are variety-seeking and inconsistent in their preferences across multiple purchase instances. Such customers are readily identified with only a few observations (we would quickly identify products a and b such that a is purchased over b in one instance, and b is purchased over a in another instance). These customers will be represented with empty DAGs, and traditional techniques may be applied.

It is also worth pointing out that using the transitivity of the pair preferences captured by arcs in the DAGs (e.g., $a > b$ and $b > c$) allows us to infer a richer set of preferences not directly revealed (e.g., $a > c$) and that are not explicitly subsumed in traditional models.

Finally, note that our model for customer choice behavior is consistent with the random utility maximization (RUM) class of models. These models assume that each customer samples utilities for each of the considered products and chooses to purchase the product with the maximum utility. A particular model specification consists of a definition of the distribution from which utilities are drawn. Since each realization of utilities induces a preference list with higher utility products preferred to lower utility products, the RUM model can be equivalently specified through a distribution over preference lists (e.g., see Strauss 1979 and section 2 of Mahajan and van Ryzin 2001 for further discussion). In a similar fashion, our model assumes that in each choice instance, a customer samples a preference list according to a distribution over preference lists (conditioned on her DAG) and chooses the most preferred product. Consequently, our model description of customer choice behavior is consistent with that of a RUM class.

2.3. Data Model

We assume access to panel data with the panel consisting of m customers purchasing over T periods. In particular, the purchases of customer i are tracked over T_i discrete time periods for a given category of products.

To simplify notation, we relabel the periods on a per-customer basis: $t = 1$ corresponds to his first visit to the store, $t = 2$ corresponds to his second visit to the store, and T_i corresponds to the last one in the training data set. For each customer i , the basic data consist of purchases of the customer over time, denoted by tuples (a_{jit}, S_{it}) for $t = 1, 2, \dots, T_i$, where $S_{it} \subset \mathcal{N}$ denotes the subset of products on offer in period t and $a_{jit} \in S_{it}$ denotes the product purchased in period t .

In addition, the data set includes information about the products $P_{it} \subset S_{it}$ that were on promotion in period t . A product may be on display promotion (prominently displayed on the store shelves) or on price promotion (offered at a discounted price). In either case, the product exhibits a distinctive feature that makes it stand out from the others on offer, potentially impacting the purchase behavior of the customer (more on this below). To simplify the analysis, we consider the promotion feature as a binary attribute of a product. However, our model could be extended to distinguish a finite number of price discount depths and display formats.

2.4. Building the Customers' Partial Orders

We dynamically build customers' partial orders from the panel data by keeping track of the sequences of interactions between the customers and the firm. Specifically, for each customer i , we start with the empty DAG D_i (i.e., a DAG with isolated nodes) and sequentially add preference arcs (pairwise comparisons) inferred from the sequence of purchase observations of the customer.

To explain how we infer the preference arcs, consider the store visit corresponding to the purchase observation (transaction) (a_j, S) of a customer, in which product a_j was purchased from the subset $S \cup \{0\}$. Let σ denote the preference list used by the customer for this purchase instance. The classical assumption is that the customer considers all of the offered products and decides to purchase the most preferred one. In other words, when offered subset S , the customer purchases $\arg \min_{a_l \in S \cup \{0\}} \sigma(a_l)$. This classical assumption then implies that since a_j is the observed purchase, the underlying preference list σ is such that $a_j = \arg \min_{a_l \in S \cup \{0\}} \sigma(a_l)$, so that $\sigma(a_j) < \sigma(a_l)$ for all $a_l \in S \cup \{0\}, l \neq j$. But customers are rationally bounded and may not consider all of the products on offer. A typical supermarket carries tens (or even hundreds) of products under each category. Further, the customer may not adopt a complex purchase process for frequently purchased products. Given this, it is unreasonable to expect the customer to evaluate all of the products offered on the shelves. As a result, we assume that customers evaluate only a subset of the offered products, often called the *consideration set*. Existing literature provides evidence for various behavioral heuristics that

customers use to construct consideration sets (e.g., see the recent survey by Hauser 2014).

Now, if customers only evaluate the products in a consideration set, the purchased product a_j would belong to and only be preferred to all other products in the consideration set. Let $C \subset S \cup \{0\}$ denote the consideration set. We must then have that $\sigma(a_j) < \sigma(a_i)$ for all $a_i \in C \setminus \{a_j\}$, and therefore we would add the arcs: $D \leftarrow D \cup \{(a_j, a_i) : a_i \in C \setminus \{a_j\}\}$. Note that ignoring the consideration set results in inferring spurious comparisons of the form $\sigma(a_j) < \sigma(a_i)$, for $a_i \in (S \cup \{0\}) \setminus C$.

Unfortunately, the consideration set in each store visit is unobservable. As a result, we must infer it from the observed transaction in a sequential basis. We face the following challenge: on the one hand, incorrectly assuming a large consideration set (say $C = S \cup \{0\}$) increases both the number of correct and spurious comparisons, but on the other hand, conservatively assuming a small consideration set decreases both types of comparisons. Hence, we must balance the requirement of maximizing the number of correct pairwise comparisons (to obtain a more accurate estimate of the underlying preference distribution) and minimizing the number of spurious pairwise comparisons (which introduce biases in the estimates). To address this challenge, we focus on the three consideration set definitions below. In all of them, during the DAG building process, as soon as the addition of arcs from $a_{j,t}$ to C_{it} implies the creation of a cycle in D_i , then we stop, delete all of the arcs, and keep D_i as the empty DAG.

1. *Standard model*, where the consideration set is the same as the offer set. More precisely, for customer i in period t , we define the consideration set $C_{it} = S_{it} \cup \{0\}$ —i.e., the customer evaluates all of the products on offer, even ignoring the effect of promotions. This assumption is reasonable for product categories in which the number of offered products is relatively small, or for a customer who goes through an exhaustive purchase process before making the purchase decision. Stores with selective offerings and high-priced product categories are good examples.

2. *Inertial model*, designed to capture the *inertia of choice* or *short-term brand loyalty*. The key intuition behind this model is that customers will continue to purchase the same product unless there is a “trigger event” that forces them to consider other products. Starting from $C_{i1} = S_{i1} \cup \{0\}$, for $t = 2, \dots, T_i$, we do: given the previous product purchased by customer i , $a_{j_{i,t-1}}$, and given the set of promoted products in period t , $P_{it} \subset S_{it}$, her consideration set in period t is

$$C_{it} = \begin{cases} \{a_{j_{i,t-1}}\} \cup P_{it} \cup \{0\} & \text{if } a_{j_{i,t-1}} \neq a_{j_{i,t}} \text{ and } a_{j_{i,t}} \notin P_{it}, \\ S_{it} \cup \{0\} & \text{if } a_{j_{i,t-1}} \notin S_{it} \text{ and } a_{j_{i,t}} \notin P_{it}, \\ P_{it} & \text{if } a_{j_{i,t}} \in P_{it}. \end{cases} \quad (1)$$

In words, we are modeling two trigger events: promotions and stockouts. We first assume that if the previous purchase is in stock and the current purchase is not on promotion, then the customer is expected to stick to the previous purchase—i.e., it is expected that $a_{j_{i,t}} = a_{j_{i,t-1}}$. In this case, the customer considers only the previously purchased product in addition to any other products that are on promotion. In the second case, if the previously purchased product is stocked out, then the customer is forced to evaluate all products on offer to determine the “new” best product. Finally, the customer buys a product on promotion—despite the fact that the previous purchase is in stock—which is consistent with existing work that notes that promotions induce trials from customers (e.g., Gupta 1988, Chintagunta 1993).

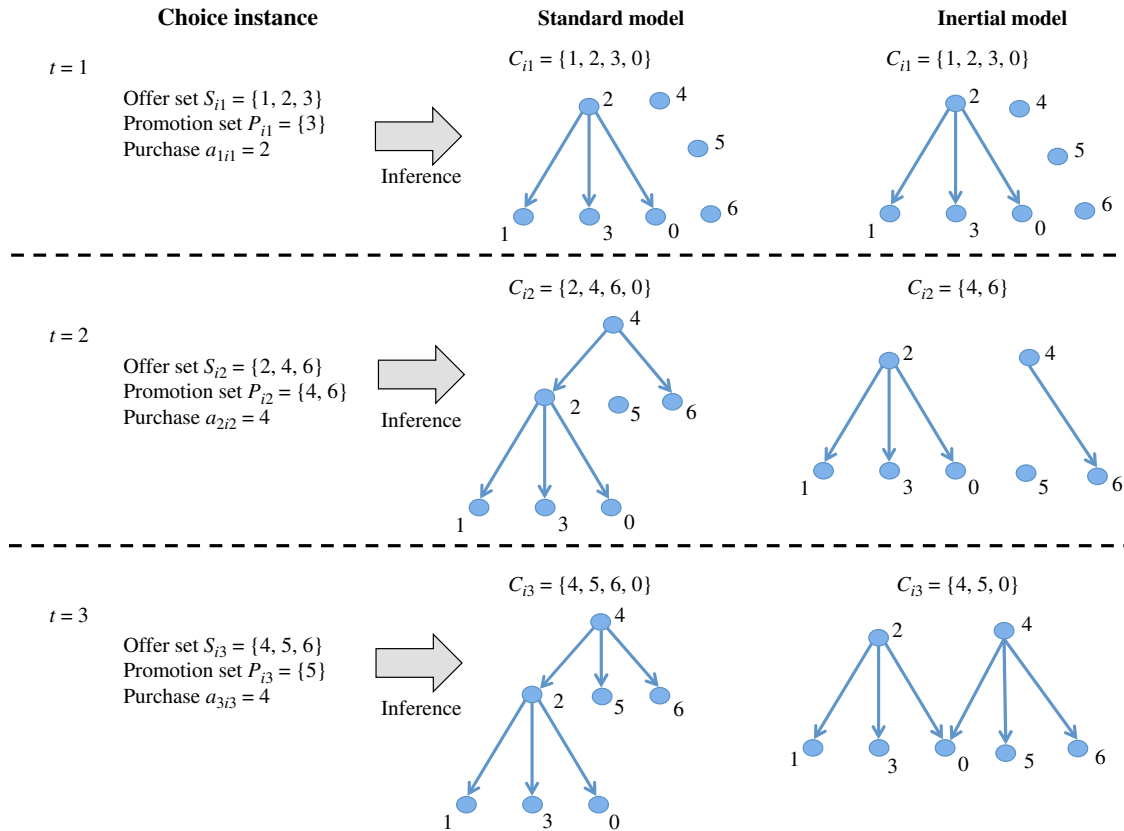
In any case, the inertial model attempts to explain any “switches” (current purchase is different from previous purchase) due to either a stockout or promotion. The customer will set the last purchase (either on promotion or not) as her *sticky product* for the next purchase. This model is reasonable for frequently purchased products in instances in which customers tend to repeat their selection.

Similar to the case of adding a cycle to D_i , under the inertial model it could happen that the observed transaction is not consistent with any of the three cases in (1). In that situation, we stop, delete all of the arcs, and keep D_i as the empty DAG.

3. *Censored model*, which extends the inertial model by allowing customers to randomly deviate from it. Thus, the censored model can capture customers whose “switches” are not explained either by a stockout or a promotion but rather by idiosyncratic reasons—i.e., $a_{j_{it}} \neq a_{j_{i,t-1}}$, but $a_{j_{i,t-1}} \in S_{it}$ and $a_{j_{it}} \notin P_{it}$. In such cases, the consideration set cannot be inferred from the inertial principles, and we set $C_{it} = \{a_{j_{it}}\}$. In other words, the model is a relaxed version of the inertial definition and does not infer any pairwise comparisons in time period t when the choice instance cannot be explained by the consideration set definition in (1).

Figure 1 illustrates the construction of a DAG D_i under the standard and inertial models for the same sequence of transactions of a given customer. The horizon has $T_i = 3$ periods, and there are $n = 6$ products in the category. For the standard model, $C_{it} = S_{it} \cup \{0\}$, for all t . For the inertial model, the consideration set captures the sticky principle which is only altered by stockout or promotion events. We always start from the empty DAG and proceed similarly to the standard model when $t = 1$ (second case of (1)). When $t = 2$, the promotion effect prevails according to the third case in (1). Finally, when $t = 3$, the sticky principle dominates, illustrating the first case in (1). The customer’s preferences are then described by the final DAG in the corresponding sequence (either standard or inertial).

Figure 1. (Color online) Construction of a DAG for a Given Customer Under the Standard and Inertial Models, with $T_i = 3$ and $n = 6$



Notes. The first column describes the offer set S_{it} , the promotion set P_{it} , and the observed purchase $a_{j_{i,t}}$ of each choice instance. The second and third columns show the evolution of the construction of the DAGs under the standard and inertial models, respectively.

Figure 2 illustrates a DAG construction under the censored model in a case where the inertial model cannot explain the purchasing behavior of a given customer. In the second period, the customer is expected to buy either 2 (her *sticky product*) or 1 (which is on promotion). However, the customer chooses 4. This is set as the new sticky product for the customer, and without adding any arc to the DAG, we proceed to the next period.

Note that our procedure to build the DAGs under each of the consideration set definitions is designed so that at the end of the process, for each arc that is present, there is evidence in the data to include it and, at the same time, no evidence to exclude it. Overall, this cautious approach tends to maximize the number of meaningful edges in each DAG.

Finally, we point out that the aforementioned consideration set definitions are only three plausible ones among many others that the modeler can propose to better explain the purchasing pattern of the customers and predict future transactions. For instance, in the extreme case where we only assess that $C_{it} = \{a_{j_{i,t}}\}$, all DAGs will be empty, which brings us back to the typical

unconstrained, choice-based demand model (e.g., the single-class MNL).

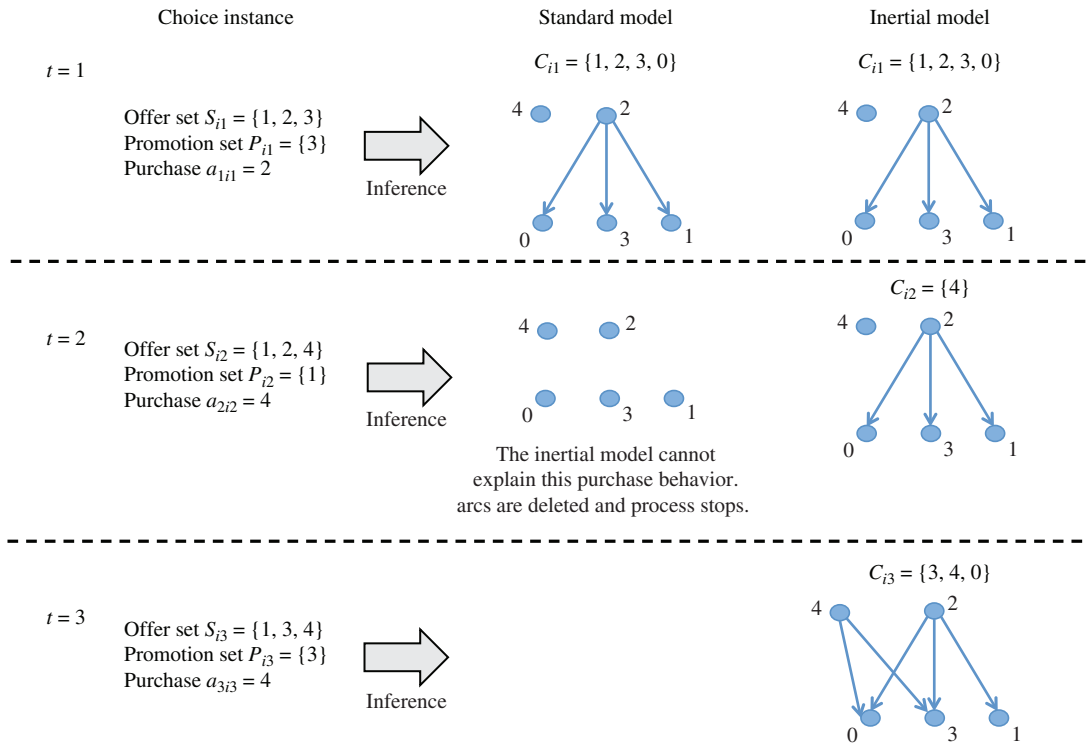
3. Analysis: Partial-Order-Based MNL Likelihoods

In what follows, we derive the probabilities of certain classes of partial orders under the MNL/Plackett–Luce model, which specifies distributions over complete orderings. We will use these results later in Section 5 for the estimation of parameters using the maximum likelihood criterion.

To streamline our discussion, in this section, we ignore the no-purchase option labeled 0 and without loss of generality focus on the standard consideration set definition where $C_{it} = S_{it}$.

Different probability distributions can be defined over the set of rankings (e.g., see Marden 1995). Specifically, we consider probabilities under the MNL model. Let $\lambda(\sigma)$ denote the probability assigned to a full ranking σ . We can then extend the calculation of probabilities to partial orders defined on the same set of products \mathcal{N} . More specifically, we can interpret any partial order D as a *censored* observation of an underlying

Figure 2. (Color online) Construction of a DAG for a Given Customer Under the Inertial and Censored Models, with $T_i = 3$ and $n = 4$



Notes. When $t = 2$, the inertial model cannot explain the purchase behavior, the building process stops, and the customer is represented with the empty DAG.

full ranking σ sampled by a customer according to the distribution over the total orders. Hence, if we define $S_D = \{\sigma: \sigma(a) < \sigma(b) \text{ whenever } (a, b) \in D\}$ as the collection of full rankings consistent with D , then we can compute the probability

$$\lambda(D) = \sum_{\sigma \in S_D} \lambda(\sigma).$$

From this expression, it follows that computing the likelihood of a general partial order D is #P-hard because counting the number of full rankings consistent with D is indeed #P-hard (Brightwell and Winkler 1991). Given this, we will restrict our attention to a particular class of partial orders (specified below) and derive computationally tractable, closed-form expressions for its likelihood.

3.1. MNL/Plackett–Luce Model

The Plackett–Luce model (e.g., see Marden 1995) is a distribution over rankings defined on the universe of items \mathcal{N} with parameter $v_a > 0$ for item a . For a given ranking σ , the likelihood of σ under this model is

$$\lambda(\sigma) = \prod_{r=1}^n \frac{v_{\sigma_r}}{\sum_{j=r}^n v_{\sigma_j}}.$$

For a given indexing of the products, we also use v_j to denote the parameter associated with product a_j .

We highlight here that the Plackett–Luce model is defined by a distribution over rankings that leads to choice probabilities consistent with the MNL choice probabilities. Online Appendix A1 provides a self-contained collection of preliminary results that serve as building block for the next ones. Here, we start from a corollary that states the probability of choosing a particular product from a subset $S \subset \mathcal{N}$.

Corollary 3.1. For a given subset $S \subset \mathcal{N}$, the choice probability for $a_i \in S$ under the Plackett–Luce model is

$$\mathbb{P}(a_i | S) = \frac{v_i}{\sum_{a_j \in S} v_j}.$$

As proved in the online appendix, Corollary 3.1 follows from the results of Propositions A1.2 and A1.3 therein. This result specifies the probability of an important special class of DAG: the star graph. The choice probability corresponds to the probability that a particular product is top-ranked among a subset S of products. The partial preference that prefers product a to all other products in set S can be represented as a star graph with directed edges from node a to all of the other nodes in $S \setminus \{a\}$. The expression for the probability of a star graph under the Plackett–Luce model

is the well-known choice probability expression under the MNL model (see Ben-Akiva and Lerman 1994). While traditionally, the choice probability under the MNL model is derived using the random utility specification of the model, our development here presents an alternate proof starting from the primitive of the distribution of a specific ranking.

We now consider a class of DAGs that can be represented as a forest of *directed trees*⁶ with unique roots. To aid our development, we introduce the concept of *reachability* of an acyclic directed graph. Any DAG D can be equivalently represented by its reachability function Ψ that specifies all of the nodes that can be reached from each node in the graph. More precisely, $\Psi(a) = \{b: \text{there is a directed path from } a \text{ to } b \text{ in } D\}$. We implicitly assume that a node is reachable from itself so that $a \in \Psi(a)$ for all a . Hence, $\Psi(a)$ is never empty and is a singleton set if the out-degree of a node is zero. Further, we assume that for a given DAG D , we have already computed its unique transitive reduction—i.e., a graph with as few edges as possible that has the same reachability relation as D (e.g., see Aho et al. 1972).

Now, we will focus on the class of connected DAGs that are directed trees. In particular, we will constrain to directed trees with exactly one node that has no incoming arc. We call that node the *root* of the tree. Equivalently, a directed tree can be described as a hierarchical tree, where each level k is defined by the nodes that are at distance k from the root. The probability of a given directed tree D under the MNL/Plackett–Luce model is provided in the next proposition.

Proposition 3.1. Consider a directed tree D with a unique root defined over elements in $S \subset \mathcal{N}$. Then, under the Plackett–Luce model, the likelihood of D is given by

$$\lambda(D) = \prod_{a \in \mathcal{N}} \frac{v_a}{\sum_{a' \in \Psi(a)} v_{a'}}.$$

The above result can be extended to a forest of directed trees, each with a unique root, by invoking the result of Corollary A1.1 in the online appendix. Specifically, we have the following new result.

Proposition 3.2. Consider a forest D of directed trees, each with a unique root defined over elements in $S \subset \mathcal{N}$. Then, under the Plackett–Luce model, the likelihood of D is given by

$$\lambda(D) = \prod_{a \in \mathcal{N}} \frac{v_a}{\sum_{a' \in \Psi(a)} v_{a'}}.$$

Finally, we derive the expression for the probability that a customer chooses a specific product from an offer set S conditioned on the partial order that describes the preferences of the customer. We use this expression to make individual-level purchase predictions once we infer the underlying partial orders of

each of the customers. More precisely, we are interested in the probability that product a_j from set S will be chosen given that the sampled preference list is consistent with DAG D . For that, let $C(a_j, S)$ denote the star graph with root a_j . Then, our goal is to compute the probability

$$f(a_j, S, D) \triangleq \Pr(S_{C(a_j, S)} | S_D) = \frac{\Pr(S_D \cap S_{C(a_j, S)})}{\Pr(S_D)},$$

where the second equality above follows from a straightforward application of Bayes rule. The next result computes the conditional choice probability when D and S satisfy some conditions, for which we need to define $h_D(S) \subset S$ as the subset of “heads” (i.e., the set of nodes without parents) in the subgraph of D restricted to S .

Proposition 3.3. Suppose we are given a DAG D that is a forest of directed trees, each with a unique root. Let S be a collection of products, and further assume that all nodes in $h_D(S)$ are also roots in D . Then, under the Plackett–Luce model, the probability of choosing product a_j from offer set S conditioned on the fact that the sampled preference list is consistent with DAG D is given by

$$f(a_j, S, D) = \frac{\Pr(S_{C(a_j, S)} \cap S_D)}{\Pr(S_D)} = \begin{cases} \frac{v_{\Psi(a_j)}}{\sum_{a_i \in h_D(S)} v_{\Psi(a_i)}} & \text{if } a_j \in h_D(S), \\ 0 & \text{otherwise,} \end{cases}$$

where Ψ is the reachability function of D , and $v_A \triangleq \sum_{a \in A} v_a$ for any subset A of products.

The proposition above states that, given a DAG D , the customer will only purchase products in $h_D(S)$ when offered subset S . In principle, the customer may choose any product in $h_D(S)$, and Proposition 3.3 specifies the conditional probability of the customer choosing each of the products there. The expression for the probability has an intuitive form that is consistent with the unconditional choice probability given in Corollary 3.1, with the “weight” v_a of each product replaced by the “weight” $v_{\Psi(a)}$ of the entire subtree “hanging” from node a .

4. Clustering Individuals

We now discuss how to account for heterogeneity in customer preferences. If we had sufficient data for each customer, then we can fit a model separately for each customer. However, in practice, data are sparse, and we typically have only a few observations for each customer. We overcome the data sparsity issue by assuming that customers—represented by their respective DAGs—belong to a small, predetermined number of

classes, where customers belonging to the same class k are “close” to a full preference list σ_k . In this section, we formulate an IP that segments the DAGs into K classes and compute the corresponding central orders σ_k , $k = 1, \dots, K$. In a follow-up step (Section 5), we will further estimate a separate distribution λ_k over preference lists for each cluster.

As a preprocessing step, given a collection of DAGs where customer c is represented by DAG D_c , we augment the arcs in each DAG by taking its transitive closure: we add edge $(a_j, a_{j'})$ whenever there is a directed path from node a_j to $a_{j'}$ in the original graph. The transitive closure of any directed graph with $O(n)$ nodes can be computed with $O(n^3)$ computational complexity using the Floyd–Warshall algorithm. Given this, we let D_c denote the graph obtained after completing the transitive closure.

We measure similarity between DAGs using a distance function based on the level of conflict of each customer assigned to cluster k with respect to the centroid σ_k . For each DAG D_c , the preference of a_j over $a_{j'}$ described by an edge $(a_j, a_{j'})$ is either verified in the total order σ_k or violated. Then, we define the *distance* between customer c and the centroid of the cluster, σ_k , as

$$\text{dist}(c, \sigma_k) = (\text{number of edges in } D_c \text{ in disagreement}) \\ - (\text{number of edges in } D_c \text{ in agreement}).$$

Denoting $|E_c|$ the total number of edges in D_c , and since any edge is either in agreement or disagreement, we can substitute above and get

$$\text{dist}(c, \sigma_k) \\ = 2 \times (\text{number of edges in } D_c \text{ in disagreement}) - |E_c|.$$

This distance measure penalizes the number of disagreements between a customer assigned to cluster k and σ_k and, at the same time, rewards the number of agreements.⁷ Given the level of conflict of a customer, the optimal clustering of the DAGs into K classes minimizes the aggregate level of conflict of all customers.

The IP is formulated as follows. Define binary linear ordering variables δ_{hjk} , which are equal to one if product a_h goes before a_j in the sequence (i.e., preference list) for cluster k and zero otherwise. In addition, define binary variables T_{ck} , which are equal to one if customer c is assigned to cluster k and zero otherwise. Finally, let w_{hjc} be a binary indicator of disagreement for edge $(a_h, a_j) \in D_c$ with respect to the total order that characterizes the cluster to which the customer was assigned. The IP is given by

$$\min \sum_{c=1}^m \sum_{(h,j) \in D_c} (2w_{hjc} - 1) \\ \text{s.t.: } \sum_{k=1}^K T_{ck} = 1, \quad \forall c,$$

$$\begin{aligned} \delta_{hjk} + \delta_{jkh} &= 1, \quad \forall h < j, \forall k, \\ \delta_{hrk} + \delta_{rjk} + \delta_{jkh} &\leq 2, \quad \forall h, r, j, j > h, j > r, h \neq r, \forall k, \\ \delta_{jkh} + T_{ck} - w_{hjc} &\leq 1, \quad \forall (h, j) \in D_c, \forall k, c, \\ \delta_{hjk}, w_{hjc}, T_{ck} &\in \{0, 1\}, \quad \forall h, j, k, c. \end{aligned} \quad (2)$$

Overall, there are $O(n^2(m + K) + mK)$ binary variables and $O(n^3K + n^2mK)$ constraints. The first set of equalities guarantees that any customer c is assigned to exactly one cluster k . The second set of equalities ensures that for any cluster k , either product a_h goes before product a_j in the preference list or product a_j goes before product a_h . The third set of constraints ensure a linear ordering among three products. The last set of constraints counts conflicts: when customer c is assigned to cluster k , and product a_j goes before product a_h in the total order associated with cluster k , then edge $(a_h, a_j) \in D_c$ must be counted as a conflict (i.e., $w_{hjc} = 1$). The objective is to minimize the aggregate level of conflict across all clusters. The final position of product a_j in the total order σ_k , $\sigma_k(a_j)$, can be determined from $\sigma_k(a_j) = \sum_{i:i \neq j} \delta_{ijk} + 1$.

We note that the development above can be formalized using the maximum likelihood estimation (MLE) framework when the underlying preferences of the customers are described by a Mallows model.⁸

Solving the IP to optimality is challenging in general. It has a very poor linear programming relaxation where all variables w_{hjc} take value zero and, in our experience, becomes very hard to solve to optimality. To ease the computational process, we develop a heuristic that provides a nontrivial initial solution. The heuristic is of the greedy-type. It starts by sorting all of the DAGs in decreasing order of number of edges. The DAG with the largest number of edges, $D_{(1)}$, is assigned to cluster 1. Then, it picks the $(K - 1)$ DAGs with the most number of conflicts with $D_{(1)}$, and assigns them to corresponding clusters $k = 2, \dots, K$. Finally, it goes sequentially over all of the remaining DAGs (following the decreasing number of edges) and counts the total number of conflicts between an unassigned DAG and all of the DAGs of each cluster $k = 1, \dots, K$. It assigns the DAG to the cluster with least conflicts.

Clustering and centroid calculation heuristic

Input Given a collection of DAGs $\{D_1, \dots, D_m\}$ and a number of clusters K , do:

Step 1 Sort the customers in decreasing order of number of edges $|E_c|$. Let $(D_{(1)}, \dots, D_{(m)})$ be the new order.

Step 2 Pick $D_{(1)}$ and assign it to cluster 1. Mark $D_{(1)}$ as an *assigned* customer.

Step 3 Count the number of conflicts between each DAG in $\{D_{(2)}, \dots, D_{(m)}\}$ and $D_{(1)}$. Then, pick the $(K - 1)$ DAGs with the most number of conflicts and assign each of them to the empty clusters $k = 2, \dots, K$. Mark the selected $K - 1$ customers as *assigned*.

Step 4 For $c := 1$ to m do

If customer $D_{(c)}$ was not assigned before, do:
Mark customer $D_{(c)}$ as assigned.
For $k := 1$ to K do
 Count the number of conflicted edges
 between $D_{(c)}$ and all of the DAGS
 already assigned to cluster k
Endfor
Assign $D_{(c)}$ to the cluster with minimum
number of conflicts
Endif

Step 5 For $k := 1$ to K do:
 Solve the IP (2) for cluster k (assuming $K = 1$)
 and come up with a ranking σ_k .

Step 6 Stop.

For a single cluster, the problem is the generalization of the *Kemeny optimization problem*—defined for total orders—to partial orders (e.g., see Ali and Meliã 2012).

5. Case Study on the IRI Academic Data Set

In this section, we present the results from a case study using the IRI Academic data set (Bronnenberg et al. 2008), which consists of real-world purchase transactions from grocery and drug stores. Our computations were carried out on a computer with Intel Core i7 (4th Gen), 2 GHz processor, and 8 Gb of RAM.

The purpose of the case study is twofold: (1) to demonstrate the application of our predictive method to a real-world setting and determine the accuracy of the assessments, and (2) to pit our framework in a horse race against three popular benchmarks which are variations of the MNL model. We start with the latent-class MNL (LC-MNL) and the random parameters logit (RPL) models,⁹ followed by the MNL model of brand choice proposed by Guadagni and Little (1983). We tested all of the approaches on the accuracy of two prediction measures on holdout data and find that our methods outperform the benchmarks in most of the categories analyzed.

5.1. Data Analysis

We considered one year (calendar year 2007) of data on the sales of consumer packaged goods (CPG) for chains of grocery stores in the two largest Behavior Scan markets. We focused on a total of 29 categories, listed in Table 1. The data consist of 1.2 M records of weekly sales transactions from 84 K customers spanning 52 weeks.¹⁰ For each purchase transaction, we have the week and the store id of the purchase, the Universal Product Code (UPC) of the purchased item, the panel id of the purchasing customer, the quantity purchased, the price paid, and an indicator of whether the purchased item is on price/display promotion. Because there is no explicit information about the assortment faced by

each consumer upon her visit, the offer set is “approximately built” by aggregating all of the transactions within a given category observed from the same store during a particular week.

We split the transaction data into the training set—consisting of the first 26 weeks of transactions—and the test/hold-out set—consisting of the last 26 weeks of transactions. We focused on customers with at least two purchases over the training period. This resulted in a total of 64 K customers and 1.1 M transactions. To address data sparsity, we aggregated the sales data by vendor as follows. Each purchased item in the data is identified by its collapsed UPC code, which is a 13-digit code. We aggregated all of the items with the same vendor code (comprising digits three through seven) into a single “product.” A detailed summary of the data is provided in Table 1.

We note from Table 1 that we observed nonempty DAGs for 31.2% of the customers under the standard, 38.5% under the inertial, and 71.4% under the censored consideration set definitions. The DAGs for the remaining customers were empty because of the appearance of directed cycles during the construction of their DAGs, or deviations from the assumptions in the case of the inertial model.

These numbers confirm our discussion in Section 2.4: The standard consideration set appears to be restrictive in the sense of imposing a heavy computational burden on the customer side by assuming that a transaction implies the evaluation of all products on offer, potentially adding several spurious arcs to the DAGs (as verified in the “Dens. st.” column), and resulting in directed cycles for the majority of the individuals. Consequently, we infer nonempty DAGs that are denser but for only a minor fraction of the customers. The inertial model reduces the consideration set cardinality (and hence, the customer’s evaluation burden and the associated DAG density) by supporting the “stickiness principle,” being able to define meaningful arcs in more customers. Overall, the inertial model increases the number of customers with nonempty DAGs but decreases the average DAG density by being conservative. The censored model, which relaxes the inertial model by ignoring switches that cannot be explained through stockouts or promotions, increases both the number of customers for whom we infer nonempty DAGs and the associated DAG density significantly.

This preliminary analysis of the data suggests that the censored model provides a good compromise between the number of customers described with nonempty DAGs and the richness of the arcs included (in the sense that the associated DAGs are denser and more meaningful).

Our benchmarking study focuses on the customers with nonempty DAGs to assess the additional benefit

Table 1. Summary of the Data

Category			Customers			Nonempty DAGs					
Shorthand	Expanded name	Vendors	Total	≥ 2 sales	Avg. trans.	# St.	Dens. st.	# Iner.	Dens. iner.	# Cens.	Dens. cens.
beer	Beer	67	1,796	1,154	7.11	420	52.39	475	24.66	978	43.18
blades	Blades	9	703	243	3.07	121	4.62	134	3.11	220	4.29
carbbev	Carbonated beverages	46	4,677	4,387	17.63	618	20.29	643	7.64	1,700	16.29
cigets	Cigarettes	13	452	307	10.39	231	8.27	232	6.81	303	7.55
coffee	Coffee	59	3,101	2,255	5.59	810	26.21	1,071	16.19	1,886	22.27
coldcer	Cold cereal	39	4,438	3,998	10.94	534	22.14	699	12.18	1,948	19.36
deod	Deodorant	32	1,345	653	3.47	275	17.92	321	11.32	581	15.39
diapers	Diapers	4	337	173	5.09	88	3.59	93	2.52	152	3.44
factiss	Facial tissue	10	2,967	2,063	4.96	903	5.47	1,005	3.60	1,843	4.82
fzdinent	Frozen dinners/Entrees	77	3,707	3,288	13.46	700	41.35	927	23.41	1,799	40.05
fzpizza	Frozen pizza	38	3,460	2,946	7.83	866	17.92	1,170	11.67	1,859	17.75
hhclean	Household cleaners	68	2,725	1,699	4.14	250	42.85	346	28.32	1,536	34.77
hotdog	Hot dogs	41	3,318	2,187	3.82	813	21.49	1,110	13.12	2,015	18.45
laundet	Laundry detergent	18	3,196	2,181	4.04	1,008	12.60	1,339	7.23	1,940	11.70
margbutr	Margarine/Butter	16	3,474	2,750	5.65	1,170	11.51	1,385	8.64	2,468	10.80
mayo	Mayonnaise	14	3,761	2,386	3.28	1,662	8.49	1,709	4.26	2,291	6.97
milk	Milk	33	4,851	4,652	14.90	1,540	17.84	1,660	9.55	3,349	15.23
mustketc	Mustard	52	3,728	2,515	3.66	480	22.35	872	12.60	2,234	18.56
paptowl	Paper towels	11	3,072	2,051	5.20	671	7.76	1,002	6.24	1,792	7.78
peanbutr	Peanut butter	19	3,153	1,923	3.89	1,041	9.98	1,181	5.83	1,825	9.07
saltsnck	Salt snacks	95	4,727	4,446	15.09	546	38.05	622	20.67	1,857	33.40
shamp	Shampoo	41	1,466	738	3.73	270	24.86	338	14.61	614	21.42
soup	Soup	90	4,636	4,322	12.02	724	41.93	967	17.85	2,397	37.18
spagsauc	Spaghetti/Italian sauce	52	3,473	2,698	5.46	1,026	22.92	1,352	12.57	2,185	20.04
sugarsub	Sugar substitutes	10	750	308	3.30	246	6.22	250	4.64	303	5.24
toitisu	Toilet tissue	11	3,760	2,817	5.10	1,144	8.54	1,547	6.42	2,460	8.30
toothbr	Toothbrushes	36	1,115	499	3.06	189	18.84	239	12.58	459	17.18
toothpa	Toothpaste	25	2,110	1,186	3.58	610	14.69	686	7.35	1,051	12.53
yogurt	Yogurt	26	3,766	3,491	19.81	1,136	11.39	1,376	6.89	1,903	10.84

Notes. The column “Vendors” reports the number of *products* in training data after aggregating different UPCs by vendor. Focusing on the training data, for customers buying from each category, we report: the original total number of them, the number of customers with at least two transactions, and among the latter, the average number of transactions per customer. Then, for each consideration set model, we report the number of customers with nonempty DAGs in the training data set and the average DAG density (defined as the average number of arcs in the nonempty DAGs).

that can be obtained from the DAG structure. When a customer has an empty DAG, our model reduces to a classical RUM model; therefore, customers with empty DAGs can be analyzed with well-established methods (e.g., MNL and LC-MNL).

5.2. Models Compared

First, we pitted our partial-order MNL (PO-MNL) method against two popular benchmarks based on the MNL model: the LC-MNL and the RPL models. All three models belong to the general RUM model class, which assumes that a customer samples product utilities in each purchase instance and chooses the product giving the highest one. The difference is that in our model, the draws are independent but conditioned on being consistent with the customer partial order across different purchase instances.

Within the RUM class, the single-class MNL model is the most popular member, with several other sophisticated models being extensions of it. The MNL model assumes that customers assign utility $U_j = u_j + \varepsilon_j$ to

product j and choose the product with the maximum utility. The error terms ε_j are Gumbel distributed with location parameter 0 and scale parameter 1. The nominal utility u_j for product j is usually assumed to depend on covariates x_{ji} in a linear-in-parameter form: $u_j = \beta_0 + \sum_i \beta_i x_{ji}$. However, because we train and test on the same universe of products, we directly estimate the nominal utilities u_j from transactions without requiring attribute selection. This approach is common in operations-related applications, where the product universe is fixed and attribute selection is nontrivial.¹¹ Nevertheless, we emphasize that like in the case of the benchmarks, our PO-MNL method allows the incorporation of covariates. Next, we briefly describe how we fit each of the models to data.

5.2.1. Model Fit According to Our Method. Once we infer the DAGs according to the consideration sets models described in Section 2, we have two different treatments: (i) the whole population as a single class of

customers and (ii) the case where we cluster the individuals. To cluster the DAGs, we solve the IP described in (2) with a time limit of five minutes on MATLAB combined with ILOG CPLEX callable library (v12.4). We tried $K = 1, \dots, 5$, classes, and retain the solutions attaining the largest objective function value in (2).

The output of the clustering is used to estimate a K latent-class PO-MNL in which we assume that each customer belongs to one of the K (latent) classes, with $K = 1$ for treatment (i). A customer belonging to class h samples her DAG according to an MNL model with parameters β_h . A priori, a customer has a probability γ_h of belonging to class h , where $\gamma_h \geq 0$ and $\sum_{h=1}^K \gamma_h = 1$. We approximate the likelihood of a DAG under a single-class PO-MNL with the expression described in Proposition 3.2, which becomes exact when the DAG is a forest of directed trees. After marginalizing the likelihood of the DAG over all possible latent classes, the *regularized maximum likelihood estimation problem* that we need to solve can be written as

$$\max_{\beta, \gamma} \left\{ \sum_{i=1}^m \log \left[\sum_{h=1}^K \gamma_h \prod_{j=1}^n \frac{\exp(\beta_{hj})}{\sum_{a_l \in \Psi_i(a_j)} \exp(\beta_{hl})} \right] - \alpha \sum_{h=1}^K \|\beta_h\|_1 \right\},$$

where $\Psi_i(a_j)$ denotes the set of nodes that can be reached from a_j in DAG D_i (recall, always including a_j).¹² Note that the estimation relies only on the DAGs of individual customers and not on their detailed purchase transactions.

The above estimation problem can be shown to be nonconcave for $K > 1$, even for a fixed value of α . To overcome the complexity of directly solving it, we used the EM algorithm. As part of the EM initialization, we used the output from the clustering in Section 4 as our starting point. More precisely, the clustering yields subsets $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K$, which form a partition of the collection of all of the customer DAGs. To get a parameter vector $\beta_h^{(0)}$, we fit a PO-MNL model to each subset of DAGs \mathcal{D}_h by solving the following single-class PO-MNL maximum likelihood estimation problem:

$$\max_{\beta_h} \left\{ \sum_{i \in \mathcal{D}_h} \sum_{j=1}^n \left[\beta_{hj} - \log \left(\sum_{a_l \in \Psi_i(a_j)} \exp(\beta_{hl}) \right) \right] - \alpha \|\beta_h\|_1 \right\}. \quad (3)$$

We tuned the value of α by fivefold cross-validation, as described in Online Appendix A2.1. We also set $\gamma_h^{(0)} = |\mathcal{D}_h| / (\sum_{l=1}^K |\mathcal{D}_l|)$. Then, using $\{\gamma^{(0)}, (\beta_h^{(0)})_{h=1}^K\}$ as the starting point, we carried out EM iterations. Further details are provided in Online Appendix A2.1.2.

Prediction. Given the parameter estimates, we make predictions as follows. For a customer i with DAG D_i , and defining $v_{hj} = \exp(\beta_{hj})$, we estimate the posterior

membership probabilities $\hat{\gamma}_{ih}$, for each h , at the beginning of the holdout sample horizon, and make the prediction:

$$f_i(a_j, S) = \sum_{h=1}^K \hat{\gamma}_{ih} f_h(a_j, S, D_i),$$

$$\text{where } \hat{\gamma}_{ih} = \frac{\gamma_h \prod_{j=1}^n v_{hj} / (1 + \sum_{l \in \Psi_i(a_j)} v_{hl})}{\sum_{d=1}^K \gamma_d \prod_{j=1}^n v_{dj} / (1 + \sum_{l \in \Psi_i(a_j)} v_{dl})},$$

where $f_h(a_j, S, D_i)$ is the approximation from Proposition 3.3 for predicting the choice probability for an individual described by parameters \mathbf{v}_h . Similar to estimation, the prediction relies only on the DAG structure of customer i and not on detailed transaction information.

5.2.2. Benchmark Models Fit According to the Classical Approach.

The two MNL-based benchmarks are briefly described here, with further details in Online Appendix A2.1. The first model is the *k-latent-class MNL* (LC-MNL) model,¹³ where each customer belongs to one of k unobservable classes, $k = 1, \dots, 10$, and remains there during the whole horizon. We keep track of the customer id in the panel data to estimate the nominal utilities, and we treat transactions as independent realizations. This model makes individual-level predictions by averaging the predictions from k single-class models, weighted by the posterior probability of class membership. We fit the model with k classes, and for each performance metric (to be introduced in Section 5.3), we report the best out-of-sample performance from these 10 models.¹⁴ In our case, the estimation requires assessing the value of kn parameters.

The second model is the RPL model, which captures heterogeneity in customer preferences by assuming that each customer samples the β parameters of the utilities according to some distribution and makes choices according to a single-class MNL model with parameter vector β . The key distinction from the LC-MNL model is that the distribution describing the parameter vector β may be continuous and not just discrete with finite support. In our case, we assume that β is sampled according to $N(\mu, \Sigma)$, which denotes the multivariate normal distribution with mean μ and variance-covariance matrix Σ , where Σ is a diagonal matrix with the j th diagonal entry equal to σ_j^2 . RPL requires the estimation of $2n$ parameters within a computationally intensive sample-average approximation approach.

For both benchmarks, and following the common practice, we use the standard consideration set definition. Given the parameter estimates, we make individual-level predictions by updating the class membership probabilities at the beginning of the holdout sample horizon.

5.3. Experimental Setup

We test the predictive power of the models on a *one-step-ahead* prediction experiment in each category under two different metrics: chi-square and miss rates. Here, the periods are labeled over the holdout horizon as $t = 1, 2, \dots, T$. For each customer i , the objective is to predict the product she will purchase in period $t + 1$ when given the last purchase of the customer up to period t and offer set S_{t+1} and promoted products P_{t+1} she faces in period $t + 1$.

For each (category, consideration set definition) combination, there is a specific subset of customers described with nonempty DAGs (see Table 1), and for that specific subset, we calibrate the two benchmarks and our PO-MNL model.

Both metrics considered rely on the definition of the indicator function $y_i(a_j, t)$, taking the value one if customer i makes a purchase in period t and product a_j has the highest predicted choice probability, and zero otherwise—i.e.,

$$y_i(a_j, t) = I\{\text{customer } i \text{ makes a purchase in period } t \text{ and } f_i(a_j, S_t) \geq f_i(a_l, S_t) \forall l \in S_t\}.$$

The first metric computes the following “chi-square” score:

$$\text{X2 score} = \frac{1}{|\mathcal{N}||U|} \sum_{i \in U, a_j \in \mathcal{N}} \frac{(n_{ij} - \hat{n}_{ij})^2}{0.5 + \hat{n}_{ij}}, \quad (4)$$

where $\hat{n}_{ij} = \sum_{t=1}^T y_i(a_j, t)$,

where U is the set of all individuals and n_{ij} is the observed number of times individual i purchased product a_j during the horizon of length T . The term \hat{n}_{ij} denotes the aggregate predicted number of purchases of product a_j by individual i .

To make the prediction, each model combination defined as the pair (choice model, consideration set definition) is given the offer set S_{t+1} ; promoted products P_{t+1} ; the set of individuals U_{t+1} who purchase in period $t + 1$; and all of the purchase transactions of all of the individuals, offer sets, and promoted products up to time period t . Using this information, each model combination provides choice probabilities for each of the offered products.

The score in (4) is similar to the popular chi-square measure of goodness-of-fit of the form $(O - E)^2/E$, where O refers to the observed value and E refers to the expected value. The use of this score in our setting is justified by considering the observed counts n_{ij} as the realization of the sum of independent (but not identically distributed) Bernoulli random variables.¹⁵ We add 0.5 to the denominator to smooth the score and deal with undefined instances.¹⁶ The score measures

the ability of the model combinations to predict aggregate market shares, with lower scores indicating better predictive accuracy.

We also measure the miss-rate of the model combinations, defined as

$$\text{miss rate} = \frac{1}{|U||T|} \sum_{i \in U} \sum_{t=1}^T I\{\text{customer } i \text{ makes a purchase in period } t \text{ and } y_i(a_{jit}, t) = 0\}, \quad (5)$$

where $I[A]$ is the indicator function taking the value one if A is true and zero otherwise. Recall that a_{jit} denotes the product purchased by individual i in period t . Again, lower miss rates are better. Note that the miss-rate measure is very stringent, since it rewards or penalizes the assessment of the right customer purchase on a per-period basis (rather than aggregating over the whole horizon as in the X2 case).

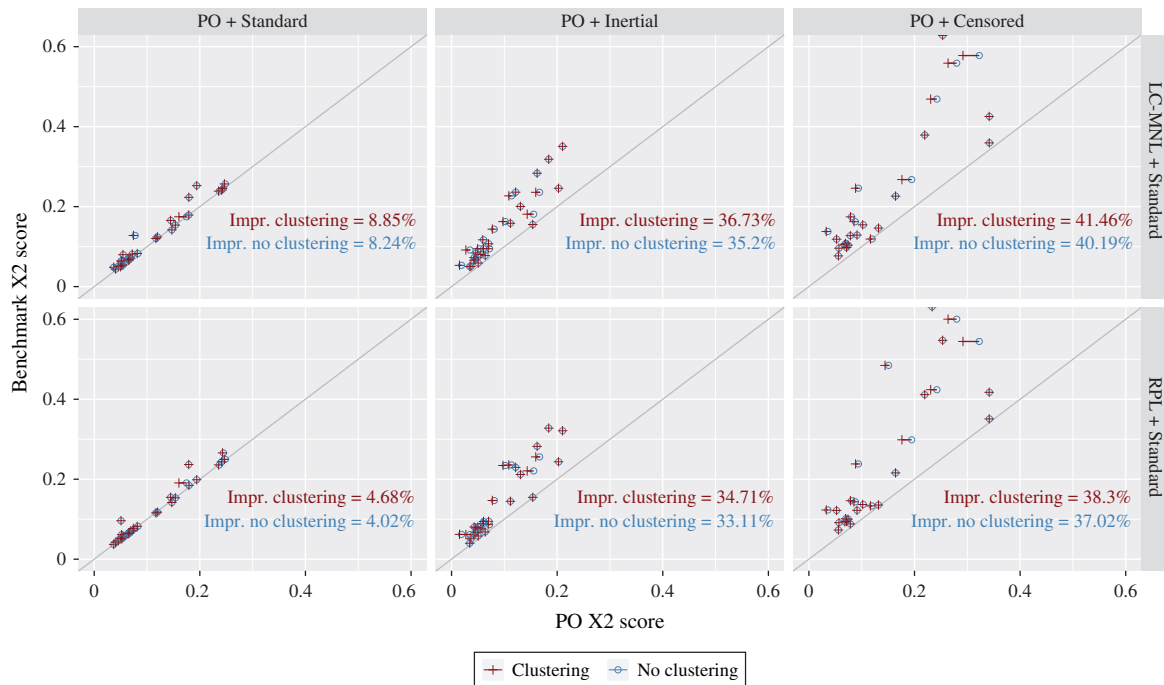
The above experimental setup is reflective of the practical prediction problem that a firm faces. Both measures are computed conditioning on a purchase event. Predicting the set of customers who will purchase in a given period t can be done separately using classical time-series models that capture any periodicity and seasonality in customer store visits.

5.4. Results and Discussion

We compare the predictive performance of the two benchmark models (LC-MNL and RPL under the standard consideration set definition) with three variants of our model class: PO-MNL (also labeled just PO from here onward) under the standard, inertial, and censored consideration set definitions.

Figure 3 presents scatterplots of the X2 scores of LC-MNL and RPL (under the standard model) versus the X2 scores of the three combinations of PO and the consideration set definitions, across the 29 product categories, including both single-class and clustering cases. We observe that our PO method dominates both (LC-MNL, standard) and (RPL, standard) benchmarks across the board. The charts allow to isolate three concurrent effects that favor our proposal. First, when considering the left column in Figure 3, we are focusing on the standard consideration set definition and comparing our PO versus both benchmarks. For the no-clustering case, the X2 score from PO exhibits an average improvement of 8.24% over LC-MNL and 4.02% over RPL. This improvement can be attributed to the effectiveness of the DAGs capturing partial preferences of the customers, which is remarkable because both benchmark models have more parameters and explicitly account for customer-level heterogeneity. It takes only around 10 seconds to fit the PO model as opposed to 67 minutes to fit the RPL model (which in turn dominates LC-MNL). The key reason for the superior performance of PO is that

Figure 3. (Color online) Scatterplot of the X2 Scores of All 29 Categories Under the PO and the (LC-MNL, Standard) and (RPL, Standard) Benchmarks for Both Single-Class and Clustering Cases



Note. Lower is better; therefore, we outperform the benchmark for points above the 45° line.

even without explicit clustering, the model accounts for heterogeneous customer preferences through their partial orders, making more efficient use of the limited data.

The second effect stems from the behavioral biases captured by nonstandard consideration set definitions: the (PO, inertial) combination allows to boost the X2 score performance to the range 33%–35%, and the (PO, censored) pair pushes it further to 40%. This outstanding improvement reveals the significant explanatory power displayed by partial preferences (represented by DAGs) jointly with the behavioral biases that determine the construction of such DAGs.

The third effect is captured by the clustering IP in Section 4 that groups together users with similar preferences. It can further boost the performance by 0.6 percentage points (pp) under the standard model, to around 1.4 pp under the behavioral (i.e., inertial and censored) models. Our results indicate that identifying similar customers has a more significant impact under the behavioral models, where DAGs are more informative since they are built on fewer spurious edges. Thus, at the expense of some additional computational effort, clustering into only a few classes can further boost our performance.

In Figure 4, we report the *loyalty score* of each category computed on the training data (left panel) and explore possible correlations between X2 scores and loyalty scores (right panel). To compute the loyalty

score of a category, we calculate the fraction of the total purchases coming from the most frequently purchased product (i.e., vendor) of each customer buying from that category and take the average of those fractions across customers purchasing from the category. There is a clear positive correlation between X2 scores and loyalty scores. That is, the X2 performance leveraged by the DAGs (versus the most competitive benchmark; i.e., (RPL, standard)) is further emphasized for high-brand-loyalty categories, and it is even more emphasized under behavioral models that highlight the stickiness principle (i.e., both inertial and censored models show a higher slope). This observation is further investigated in Online Appendix A2.2 (see Figure A2 therein). For categories with high loyalty indices, we can obtain a huge improvement in the X2 score by accounting for individual partial preferences under the three consideration set definitions.

Figure 5 presents scatterplots of the miss rates following a display format similar to Figure 3. Recall that this is a more stringent predictive measure, since it rewards or penalizes the sum of individual transaction assessments (see Definition (5)). We observe that our model combinations obtain improvements of between 2% and 6% over the benchmarks in four of the six panels, when using the clustering procedure. In the remaining two (lower right) panels, our performance is similar to the benchmarks (most points lie on the 45° line). The same three aforementioned

Figure 4. (Color online) Loyalty Scores Across Categories, Sorted in Descending Order (Left), and Scatter Plot and Linear Regression of Percentage Improvement of the X2 Score of PO Over (RPL, Standard), for Each Consideration Set Definition (Right)





Figure 5. (Color online) Scatterplot of the Miss Rates of All 29 Categories Under the PO and the (LC-MNL, Standard) and (RPL, Standard) Benchmarks for Both Single-Class and Clustering Cases



Note. Lower is better; therefore, we outperform the benchmark for points above the 45° line.

Downloaded from informs.org by [190.195.71.189] on 24 April 2017, at 17:39. For personal use only, all rights reserved.

effects play a role here (although the behavioral bias effect is diminished when compared with (RPL, standard)). The charts in the left column reveal the value delivered by capturing partial preferences under the standard consideration set definition, providing gains between 2.6% and 5.6%. Clustering still plays a role in boosting performance, resulting in improvements close to 1 pp in some cases. As noted, the behavioral models are just as good as the (RPL, standard) model in terms of the miss rates. This may be explained by noting that the stickiness assumption can lead to double penalization: missing first when the customer deviates from purchasing the “loyal” brand, and missing again when the customer returns to the “loyal” brand because the sticky product is updated to the latest purchase. This observation suggests that if the objective is to increase accuracy of immediate predictions, it may be best to increase the chances of a hit by expanding the consideration set to the entire offer set as done on (PO, standard), to account for possible deviations from the inertial assumption.

We also investigated the correlation between percentage improvement of PO over (RPL, standard) with respect to the miss rates (see Figure A3 in Online Appendix A2.2). As we said, the miss rate measure is quite stringent, but better results are obtained under PO for high loyalty categories, mainly under the standard but also under the censored consideration set definitions.

Overall, the combination (PO-MNL, censored) seems to provide the best performance under X2 scores, and a better or comparable performance with respect to miss rates, which is quite remarkable since, when building individual DAGs under the censored consideration set definition, we could describe the behavior of 71% of the customer base with nonempty DAGs. The performance is further boosted when we apply (PO-MNL, censored) over categories with high loyalty score. However, if the focus is mainly on miss rates, the combination (PO-MNL, standard) appears to be the best option.

5.5. Comparison with the Guadagni–Little Model

The final benchmark study is to compare our method with the proposal by Guadagni and Little (1983), denoted GL onward. Guadagni and Little calibrated an MNL model of brand choice, based on panel data, over the “regular ground coffee” category. They formulated a linear-in-parameters utility model that captures covariates such as full price, indicator for brand promotion, price cut for products on promotion, indicator to keep track of the promotion status of the last two purchases, and two additional covariates to capture loyalty effects: in terms of coffee brand and packaging size. For them, loyalty is taken to be the exponentially weighted average of past purchases of the brand and the packaging size, treated as 0–1 variables. However, their loyalty scores do not reflect stockouts: the non-purchase of a product because of a stockout is treated

as a regular switch to a different brand. For the purposes of our study, we implemented a variation of the GL model, adapted to the variables present in our data set (for details, see Online Appendix A2.1).

Even though the GL model is tailored to capture loyalty, as our inertial and censored models are, there is a fundamental difference between both approaches: GL maintains a richer state via the utility function of their MNL model, which is updated dynamically as promotions arise, both in the training and holdout samples, whereas we maintain a limited state via the consideration set definition, also in both the training and holdout samples (though in the holdout sample we do not update the DAGs). More specifically, GL maintains the promotion status of the current and the previous two purchases and longer-term loyalty information through an exponential smoothing reflecting recency and frequency of purchases. Instead, the inertial and censored models maintain the promotion status of current purchase and only short-term loyalty information in the form of the previous purchase. Therefore, the comparison should be taken with caution.¹⁷

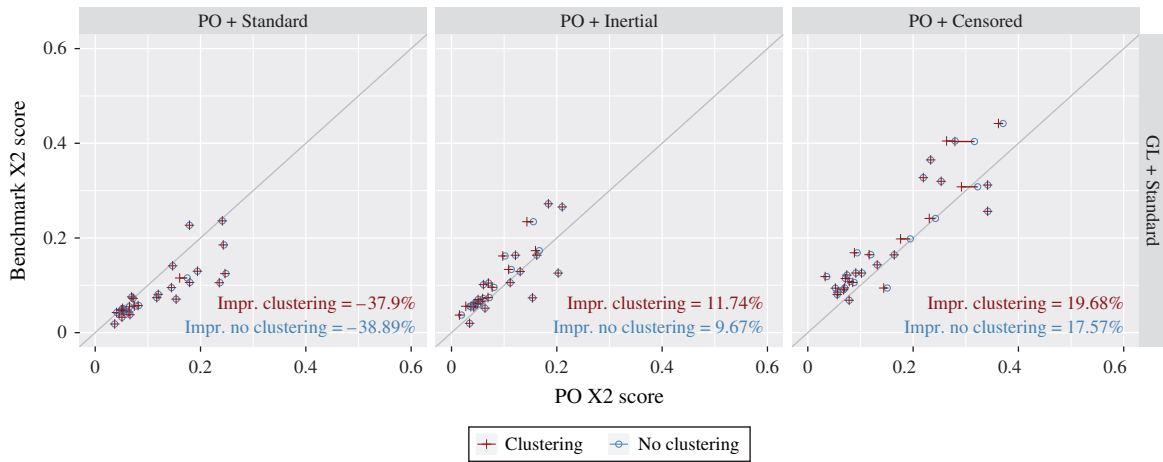
Guadagni and Little (1983) provide evidence for the remarkable performance of their method in predicting product market shares for “regular ground coffee.” The X2 score captures the accuracy of a method in predicting these market shares. Figure 6 shows that GL generally beats the performance of the standard model (which is not aimed at capturing brand loyalty) but also highlights the good performance that we achieve with (PO, inertial) and (PO, censored), with our X2 scores undercutting GL’s X2 scores by 9.7% and 17.6%, respectively. In both cases, the clustering procedure boosts the performance by extra 2 pp.

When evaluating miss rates (see Figure 7), our performance is slightly better under (PO, standard), worse by 3.8% under (PO, inertial), and better by 3.6% under (PO, censored). Recall that miss rate is a very stringent measure, where stockouts or promotions may affect a specific prediction. The sharp contrast that we observe between the inertial and censored performances may be explained by the fact that the set of customers with nonempty DAGs that are described by the censored consideration set definition is a big superset of the inertial customers with nonempty DAGs. The GL model is less accurate in predicting the purchases of these additional customers because they are not as strongly loyal as the inertial customers and are facing a changing environment with promotions and stockouts.

6. Conclusions

Estimating individual customer preferences for the purposes of making personalized demand predictions is becoming an increasingly important problem given the volume of fine-grained data that businesses are

Figure 6. (Color online) Scatterplot of the X2 Scores of All 29 Categories Under the PO and the GL+Standard Benchmark for Both Single-Class and Clustering Cases




Note. Lower is better; therefore, we outperform the benchmark for points above the 45° line.

able to collect. We proposed a nonparametric framework to estimate individual preferences from panel data on the transaction history of customers. Our technique is data driven, accommodates both parametric and nonparametric choice models, and is aimed at improving the accuracy of individual-level choice predictions. It exhibits three distinguishing features: (1) it allows each customer to have a general partial order, with the customer’s preferences being always consistent with her partial order rather than being independent across choice instances; (2) it incorporates a behavioral model of choice that captures individual bounded rationality to construct customized partial orders from purchase records; and (3) it allows for grouping customers by preference similarities. We demonstrate on real-world data that accounting for

partial-order and behavioral aspects of choice significantly improves the accuracy of individual-level predictions over existing state-of-the-art methods.

Our work opens the door for many exciting future research directions. A key theoretical contribution of our work is deriving the likelihood expressions for general classes of partial orders under the MNL/Plackett–Luce model. Existing work has mainly considered classes of partial orders (such as top-*k* rankings and partitioned preferences) that commonly arise when data are collected on the web. However, onsite transactions often yield choice data, resulting in partial orders that lack the types of structure already studied. Our work has introduced a new class of partial orders—forest of directed trees with unique roots—and has shown that it affords closed-form expression for likelihood under

Figure 7. (Color online) Scatterplot of the Miss Rates of All 29 Categories Under the PO and the GL+Standard Benchmark for Both Single-Class and Clustering Cases



Note. Lower is better; therefore, we outperform the benchmark for points above the 45° line.

the MNL/Plackett–Luce model. Identifying other general classes of partial orders for which likelihoods can be computed in a tractable manner under different popular choice models (such as variants of the MNL) is an exciting future direction. In addition, investigating properties of the clustering IP to derive theoretical guarantees for general classes of partial orders can result in significant contributions to this problem.

Finally, the comparison of the performance of our method versus the classical Guadagni and Little (1983) approach suggests that capturing state dependence via a parameterized utility function could be a means to further boost the potential of partial-order-based preferences.

Acknowledgments

The authors thank department editor Vishal Gaur, the associate editor, and three anonymous referees for their detailed and thoughtful comments, which greatly improved this paper. The authors are also grateful to the seminar participants at the University of Southern California, Universidad Torcuato di Tella (Argentina), the University of Chile, Universidad Adolfo Ibáñez (Chile), Duke University, Santa Clara University, IBM Thomas J. Watson Research Center, Johns Hopkins University, Georgetown University, University of Minnesota, University of Chicago, and Kellogg School of Management for their constructive comments that have helped improve the paper.

Endnotes

¹For instance, a field study in the airline RM practice suggests that a 20% reduction of forecast error can translate into 1% additional revenues (Pölt 1998).

²In the rank-based choice model of demand, the modeler should estimate a discrete probability mass function (pmf) on the set of customer types. The number of parameters (i.e., the number of customer types with nonzero probabilities) grows with the number of products and customers, and the volume of data; hence, the labeling as a *nonparametric* model.

³In fact, we use a variant of the model by Guadagni and Little (1983) as one of our benchmarks in Section 5.

⁴This example could be used to model the preferences of a coffee customer who always prefers decaf (product 1 and 3) to nondecaf (product 2). One may readily construct other such examples.

⁵We explain later in Section 5 how we aggregate different Universal Product Codes (UPCs) in the data set.

⁶A *directed tree* is a connected and directed graph which would still be an acyclic graph if the directions on the edges were ignored.

⁷Even though we use the term *distance*, its interpretation should be taken with caution because it could lead to negative values. This measure is a linear transformation of the Kendall tau distance, which counts the number of pairs of elements that are ranked opposite in two total orders. See Stanley (2000).

⁸The Mallows model has been extensively used in machine learning and directly specifies distributions over rankings. In particular, when the underlying customer DAGs have a particular structure (denoted *partitioned preferences*), maximizing the log-likelihood function of the DAGs under the Mallows model is equivalent to solving our IP (2). When the conditions above are not satisfied, our IP formulation can be treated as performing approximate MLE. See Jagabathula and Vulcano (2017) for further details.

⁹The RPL model is also referred to in the literature as the random coefficients model (e.g., see Train 2009, Chap. 6.2).

¹⁰The data consist of 5 K unique customers/panelists whose purchases span the 29 categories. Because we analyzed the categories separately, we treat each customer-category combination as a “customer.” There were two more categories in the data set, “photography supplies” and “razors,” that we ignore due to data sparsity.

¹¹For instance, Sabre Airline Solutions, one of the leading RM software providers for airlines, with a tradition of high-quality R&D, implemented a proprietary version of the procedure described in Ratliff et al. (2008) for a single-class MNL model.

¹²The reachability function $\Psi_i(\cdot)$ may contain the no purchase option, whose utility is normalized to be $\beta_{h_0} = 0$ for all h .

¹³Here, we use k to refer to the number of latent classes, which could be different from K —the number of clusters from Section 4.

¹⁴For this reason, the value of k for which the performance is reported could be different for different metrics and different categories.

¹⁵Since the offer sets change and predictions are conditioned on them, there random variables are not identically distributed across time.

¹⁶Our smoothing is similar to the popular add-1/2 smoothing that is used to deal with zero counts when estimating a multinomial distribution from a finite sample, where a pseudo count of 1/2 is added to the count of each type (Falahatgar et al. 2016)

¹⁷The previous benchmarks LC-MNL and RPL, where we did not consider covariates either, are more suitable for comparison. Those results allowed us to show the value of building the DAGs over a MNL model, allowing us to quantify the value of superimposing such incremental structure.

References

- Aho A, Garey M, Ullman J (1972) The transitive reduction of a directed graph. *SIAM J. Computing* 1(2):131–137.
- Ali A, Melià M (2012) Experiments with Kemeny ranking: What works when? *Math. Soc. Sci.* 64(1):28–40.
- Ben-Akiva M, Lerman S (1994) *Discrete Choice Analysis: Theory and Applications to Travel Demand*, 6th ed. (MIT Press, Cambridge, MA).
- Brightwell G, Winkler P (1991) Counting linear extensions is #P-complete. *STOC '91 Proc. Twenty-Third Annual ACM Sympos. Theory Comput.* (Kluwer Academic Publishers, Netherlands), 175–181.
- Bronnenberg B, Kruger M, Mela C (2008) The IRI academic data set. *Management Sci.* 27(4):745–748.
- Chandukala SR, Kim J, Otter T, Rossi PE, Allenby GM (2008) Choice models in marketing: Economic assumptions, challenges and trends. *Foundations and Trends® in Marketing* 2(2):97–184.
- Chaneton J, Vulcano G (2011) Computing bid-prices for revenue management under customer choice behavior. *Manufacturing Service Oper. Management* 13(4):452–470.
- Chen L, Homem-de Mello T (2010) Mathematical programming models for revenue management under customer choice. *Eur. J. Oper. Res.* 203(2):294–305.
- Chintagunta P (1993) Investigating purchase incidence, brand choice and purchase quantity decisions on households. *Marketing Sci.* 12(2):184–208.
- Clifford S (2012) Shopper alert: Price may drop just for you alone. *New York Times* (August 9), <http://www.nytimes.com/2012/08/10/business/supermarkets-try-customizing-prices-for-shoppers.html>.
- Danaher P, Smith M, Ranasinghe K, Danaher T (2015) Where, when, and how long: Factors that influence the redemption of mobile phone coupons. *J. Marketing Res.* 52(5):710–725.
- Falahatgar M, Ohannessian MI, Orlitsky A (2016) Near-optimal smoothing of structured conditional probability matrices. *Advances in Neural Information Processing Systems*, 4853–4861.

- Farias V, Jagabathula S, Shah D (2013) A new approach to modeling choice with limited data. *Management Sci.* 59(2):305–322.
- Guadagni P, Little J (1983) A logit model of brand choice calibrated on scanner data. *Marketing Sci.* 2(3):203–238.
- Gupta S (1988) Impact of sales promotions on when, what, and how much to buy. *J. Marketing Res.* 25(4):342–355.
- Haensel A, Koole G (2011) Estimating unconstrained demand rate functions using customer choice sets. *J. Revenue Pricing Management* 10(5):438–454.
- Hauser JR (2014) Consideration-set heuristics. *J. Business Res.* 67(8):1688–1699.
- Honhon D, Gaur V, Seshadri S (2010) Assortment planning and inventory decisions under stockout-based substitution. *Oper. Res.* 58(5):1364–1379.
- Honhon D, Jonnalagedda S, Pan X (2012) Optimal algorithms for assortment selection under ranking-based consumer choice models. *Management Sci.* 12(2):279–289.
- Hunter DR (2004) MM algorithms for generalized Bradley–Terry models. *Ann. Statist.* 32(1):384–406.
- Jagabathula S, Rusmevichientong P (2014) A two-stage model of consideration set and choice: Learning, revenue prediction, and applications. Working paper, Leonard N. Stern School of Business, New York.
- Jagabathula S, Vulcano G (2017) A partial-order-based Mallows model of discrete choice. Working paper, New York University, New York.
- Jeuland A (1979) Brand choice inertia as one aspect of the notion of brand loyalty. *Management Sci.* 25(7):671–682.
- Kök AG, Fisher ML (2007) Demand estimation and assortment optimization under substitution: Methodology and application. *Oper. Res.* 55(6):1001–1021.
- Kunnumkal S (2014) Randomization approaches for network revenue management with customer choice behavior. *Production Oper. Management* 23(9):1617–1633.
- Mahajan S, van Ryzin GJ (2001) Stocking retail assortments under dynamic consumer substitution. *Oper. Res.* 49(3):334–351.
- Marden J (1995) *Analyzing and Modeling Rank Data* (Chapman & Hall, London).
- Musalem A, Olivares M, Bradlow E, Terwiesch C, Corsten D (2010) Structural estimation of the effect of out-of-stocks. *Management Sci.* 56(7):1180–1197.
- Newman J, Ferguson M, Garrow L, Jacobs T (2014) Estimation of choice-based models using sales data from a single firm. *Manufacturing Service Oper. Management* 16(2):184–197.
- Pöhl S (1998) Forecasting is difficult—Especially if it refers to the future. *Revenue Management and Distribution Study Group Annual Meeting* (AGIFORS, Melbourne, Australia).
- Ratliff R, Rao B, Narayan C, Yellepeddi K (2008) A multi-flight recapture heuristic for estimating unconstrained demand from airline bookings. *J. Revenue Pricing Management* 7(2):153–171.
- Rusmevichientong P, van Roy B, Glynn P (2006) A nonparametric approach to multiproduct pricing. *Oper. Res.* 54(1):82–98.
- Smith J, Lim C, Alptekinoglu A (2009) New product introduction against a predator: A bilevel mixed-integer programming. *Naval Res. Logist.* 56(8):714–729.
- Stanley R (2000) *Enumerative Combinatorics*, Vol. 1 (Cambridge University Press, New York).
- Strauss D (1979) Some results on random utility models. *J. Math. Psych.* 20(1):35–52.
- Train K (2009) *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, New York).
- van Ryzin G, Vulcano G (2014) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Sci.* 61(2):281–300.
- Vulcano G, van Ryzin G, Chaar W (2010) Choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing Service Oper. Management* 12(3):371–392.
- Wierenga B (2008) *Handbook of Marketing Decision Models* (Springer Science+Business Media, New York).
- Zhang D, Cooper W (2005) Revenue management for parallel flights with customer-choice behavior. *Oper. Res.* 53(3):415–431.