**Submitted to *Management Science***
**manuscript MS-0001-1922.65**



# The Click-Based MNL Model: A Novel Framework for Modeling Click Data in Assortment Optimization

**(Authors' names blinded for peer review)**

We introduce the click-based MNL choice model, a novel framework for capturing customer purchasing decisions in e-commerce settings. We augment the classical Multinomial Logit choice model with the assumption that customers only consider the items they have clicked on before they proceed to compare their random utilities. We propose a simple estimation framework that leverages clickstream data and machine learning classification algorithms. We study the resulting assortment optimization problem, where the objective is to select a subset of products, made available for purchase, to maximize the expected revenue. Our main algorithmic contribution comes in the form of a polynomial-time approximation scheme (PTAS) for this problem, showing that the optimal expected revenue can be efficiently approached within any degree of accuracy. In the course of establishing this result, we develop novel technical ideas, including enumeration schemes and stochastic inequalities, which may be of broader interest. Using data acquired in collaboration with Alibaba, we fit click-based MNL and Mixed MNL models to historical sales and click data in a setting where the online platform must present customized six-product displays to users. We show that our approach significantly outperforms the Mixed MNL models in terms of out-of-sample predictive accuracy, and the computational cost of its estimation process is smaller by an order of magnitude.

*Key words*: Multinomial Logit model, consideration sets, clickstream data, approximation algorithms

## 1. Introduction

Assortment optimization is a fundamental computational challenge faced by online platforms, where the set of products displayed to customers can be seamlessly varied and personalized at essentially no cost. Existing prediction models and decision support tools in this context generally take advantage of historical transaction data to infer customers' preferences, as each purchase transaction reveals a choice made by the corresponding customer amongst the available product alternatives. Given this probabilistic prior about customer preferences, the selection of an assortment to maximize expected revenue can be framed as a combinatorial optimization problem, which has received a great deal of attention in the operations research and revenue management literature, as further elaborated in Section 1.3.

In recent years, online platforms increasingly collect data about additional forms of customer-product interactions. The primary source of such data is *clickstream* information, corresponding to the customers' search and click behavior. Empirical studies in the quantitative marketing literature suggest that clickstream data is highly valuable to predict customer purchases, brand loyalty, and churn rates (Laudon and Traver 2013). For example, by taking into account the path of webpages browsed by users of an online bookseller, Montgomery et al. (2004) show that the accuracy with which they can predict purchase conversions is increased to 40%, up from a mere 7% without integrating this piece of information. Nevertheless, despite marked research interest around discrete-choice modeling in the past decade, the marginal value of clickstream data in decision-making models still remains mostly unquantified in operational settings. In particular, leveraging these data sources within the framework of the assortment optimization problem is an apparent blind spot of the existing literature, from both modeling and computational standpoints.

*Overview of our research contributions.* Motivated by this state of affairs, the present paper aims to augment the predictive and prescriptive abilities of traditional choice models in e-commerce settings by incorporating click signals. Our main modeling idea is to assume that clicks coincide with the set of alternatives considered by a customer during a product screening phase preceding her final choice decision. This point of view leads us to introduce a choice model whose underlying dynamics unfold in two sequential stages. In the first stage, the customer forms her consideration set according to the click propensities. Then, in the second stage, she makes a purchasing decision from among the products in this restricted consideration set according to the classical multinomial logit (MNL) choice model. This modeling approach is inspired by the well-established literature in quantitative marketing and operations around two-stage choice models, which were originally conceived by Howard and Sheth (1969) and by Hauser (1978), and later empirically validated by numerous works, such as those of Jeuland (1979) and Crompton and Ankomah (1993).

We depart from the existing literature in viewing the consideration set as being explicitly defined by the click behavior, rather than being an unobserved latent parameter of the choice-making process. Using data from Alibaba's retail platform, we provide empirical evidence that our modeling approach explains customers' choice behavior significantly better than traditional choice models. From a computational standpoint, the random formation of the consideration sets creates algorithmic and probabilistic hurdles that necessitate the design of novel methodological frameworks for the resulting assortment optimization problem. In this context, our main technical contribution is a polynomial-time approximation scheme, proving that the optimal expected revenue can be efficiently approached within any degree of accuracy.

In what follows, we provide a brief review of the most relevant papers on assortment optimization, which sets the stage for a precise description of the modeling approach developed in this paper. A more detailed overview of the related literature is provided in Section 1.3.

*Brief review of existing literature.* Starting with the seminal work of Talluri and van Ryzin (2004), a growing line of research has explored assortment optimization problems under various choice models and operational constraints. When customer choice is governed by the MNL model (Luce 1959, McFadden 1974), and when there are no restrictions on the set of products to be offered, Talluri and van Ryzin (2004) show that there exists an optimal assortment consisting of all products priced above a certain threshold. Subsequent research papers have focused on constrained versions of this problem, extending the breadth of retail scenarios where assortment optimization is readily applicable (Rusmevichientong et al. 2010, Sumida et al. 2020).

One major limitation of the MNL choice model is its inability to capture customers' heterogeneity; hence, considerable research efforts have focused on formulating assortment optimization problems with choice models that account for latent customer heterogeneity. For example, a common approach for modeling general substitution patterns is to segment the customers into heterogeneous classes. Since the segment to which each customer belongs is not directly observable, the latter approach gives rise to probabilistic mixture models, such as the extensively-studied mixture of MNLs (see, e.g., Rusmevichientong et al. (2014), Désir et al. (2014), Feldman and Topaloglu (2015a)).

Another popular approach for capturing customers' choice heterogeneity is based on the notion of consideration sets. Most random-utility maximization choice models presume that customers consider *all* offered products before comparing their relative utilities. In reality, however, customers may be led to disregard certain products based on prominent features such price, ratings and reviews. For example, they may only consider purchasing products of a certain quality that are priced below some threshold, as in the models considered by Aouad et al. (2020) and by Jagabathula and Rusmevichientong (2016). Consequently, several recent papers have focused on consideration set-based assortment optimization problems (Davis et al. 2015, Gallego et al. 2020, Aouad and Segev 2020, Gallego and Li 2017). In these settings, the consideration sets are not directly observable; their distribution can only be inferred probabilistically from the final purchasing actions. Moreover, in order to develop polynomial-time algorithms for assortment optimization, these studies place additional structural restrictions either on the considerations sets or on the relative rankings over products' utilities. For example, the model of Feldman and Topaloglu (2018) assumes that the consideration sets are nested; the papers by Honhon et al. (2012), Aouad et al. (2020) and Gallego and Li (2017) are all based on the assumption of a *unique ranking*, meaning that all customers share an overarching ranking over the products' utilities. These modeling assumptions imply that only a very small fraction of all potential consideration sets and rankings occur with positive likelihood during the two-stage choice-making process.

*Click-based MNL modeling approach.* Our main modeling idea is to assume that the "clicks" observed by e-commerce platforms explicitly describe the customers' consideration sets, a key distinctive feature in relation to the existing literature. Specifically, we define each customer's consideration set precisely as the set of products she has clicked on. This assumption is quite natural, considering that customers click on products either to purchase them or to collect additional product-specific information. Following this main assumption, the customers' purchasing process unfolds in a two-step process. In the first step, the customer independently decides whether or not to click on each of the offered products, thereby forming a consideration set. In the second step, the customer makes a purchasing decision by ranking the utilities of the products she has clicked on; such ranking decisions are captured via a standard MNL choice model. We refer to the resulting choice model as the *click-based MNL* model, whose dynamics are formalized in Section 1.1.

This modeling approach has a number of distinct advantages. First, clickstream data is massively available in e-commerce; thus, our operational use case of click data for designing product recommendation engines is implementable in many e-commerce settings. Second, by assuming that the clicks follow independent Bernoulli outcomes, we develop a choice model that is both parsimonious and flexible. Specifically, the property that *any potential consideration set occurs with positive likelihood* is consistent with the lack of explicit structure in customers' clicks, as observed in real-world clickstream data such as those provided by Alibaba for our numerical experiments. This property is in sharp contrast with other consideration set-based choice models utilized in earlier literature. Lastly, the platform's ability to fully observe all click events makes the estimation of the model parameters quite straightforward. As explained in Section 1.2, the estimation strategy proposed in this paper is very efficient from a computational standpoint, and it can easily take advantage of state-of-the-art machine learning algorithms.

Nonetheless, at the core of our model lies the assumption that the observed "clicks" coincide with the customers' consideration sets. This notion leads to two fundamental research questions:

1. How practical is the assumption that customers' consideration sets are described by their clicks? Specifically, how does the predictive performance of the click-based MNL model fare against state-of-the-art choice models on real-world choice data?

2. Is the assortment optimization problem computationally tractable under the click-based MNL model? In particular, can we develop efficient algorithmic methods despite the absence of an explicit combinatorial structure on the customers' consideration sets and ranking preferences?

### 1.1. Problem formulation

In what follows, we formally describe the click-based MNL model and formulate its corresponding assortment optimization problem. We consider a setting where a retailer has access to a collection of $n$ items. Each item $i$ is associated with a selling price of $r_i$, an MNL-based preference weight of $w_i$, and a consideration probability of $\lambda_i$. Given the assortment $S \subseteq [n]$, the customers' purchasing decisions are made in two steps:

1. *Generating a consideration set.* First, a random subset $C_S \subseteq S$ is generated, by independently picking each item $i \in S$ with probability $\lambda_i$. We refer to $C_S$ as the (random) consideration set induced by the assortment $S$, whose total weight is denoted by $w(C_S) = \sum_{i \in C_S} w_i$.
2. *Picking from the consideration set.* For any given realization of the consideration set $C_S$, a single representative customer makes a purchase from among this subset according to the MNL choice model. Namely, each item $i \in C_S$ is chosen with probability $\frac{w_i}{1+w(C_S)}$, while the no-purchase option is chosen with the residual probability, $\frac{1}{1+w(C_S)}$.

In order to define the expected revenue function, let $\pi(i,S)$ be the probability of picking item $i$, given that the assortment $S$ is offered. Based on the preceding discussion, this purchase probability is clearly zero when $i \notin S$. In the more interesting scenario when $i \in S$, the latter probability can be written as

$$\pi(i,S) = \Pr\left[i \in C_S\right] \cdot \mathbb{E}\left[\frac{w_i}{1+w(C_S)} \middle| i \in C_S\right] = \lambda_i \cdot \mathbb{E}\left[\frac{w_i}{1+w_i+w(C_S^{-i})}\right] , \tag{1}$$

where the expectation above is taken over the randomness in generating the consideration set $C_S$, with the shorthand notation $C_S^{-i} = C_S \setminus \{i\}$. The objective is to identify an assortment $S \subseteq [n]$ that maximizes the resulting expected revenue $\mathcal{R}(S) = \sum_{i \in S} r_i \cdot \pi(i,S)$. This assortment optimization formulation is termed the *click-based MNL assortment problem*, for short.

*Related choice models.* It is worth noting that the click-based MNL model can be viewed as a special case of the classical mixed-MNL model with exponentially many customer segments. More specifically, in Appendix A, we present a formal description of how any click-based MNL model can be converted to an equivalent mixed-MNL model. Unfortunately, this connection provides little practical value, as the reduction requires specifying a mixed-MNL model over $2^n$ customer segments, meaning that existing algorithms for the assortment problem under the mixed-MNL cannot be efficiently applied to these instances. It is worth observing that even with polynomially many segments, the mixed-MNL-based assortment optimization problem is known to be $\Omega(n^{1-\epsilon})$-hard to approximate, as shown by Désir et al. (2014).

As explained in the literature review of Section 1.3, highly structured special cases of the click-based MNL model have surfaced in the recent literature. For example, the random consideration set model introduced by Manzini and Mariotti (2014) and studied by Gallego and Li (2017) can be viewed as a limiting case of the click-based MNL where the second-stage choice decisions follow a single ranking over the items.[1] In fact, the model of Gallego and Li (2017) is a special case of the Markov chain choice model (Blanchet et al. 2016). In contrast, our hardness result of Section 2.1 implies that, conditional on the standard complexity assumption $P \neq NP$, it is not possible to represent an instance of the click-based MNL model as an equivalent polynomially sized instance of the Markov chain model. The latter property implies a clear separation between the click-based MNL model in its general form and special cases examined in earlier literature.

To summarize, while our modeling approach is related to choice models studied in previous literature, these connections cannot be exploited to efficiently solve the click-based MNL assortment problem. As we proceed to show in this paper, the probabilistic structure of the consideration sets creates fundamental challenges that require new technical insights and algorithmic tools.

### 1.2. Main contributions

This paper studies the application of the click-based MNL model to assortment optimization, both at a fundamental algorithmic level and in practice. In terms of theory, we tightly characterize the computational status of the resulting assortment optimization problem. We show that this problem is NP-hard even for very basic click-based MNL settings. On the positive side, we develop a polynomial-time approximation scheme (PTAS), showing that provably near-optimal assortments can be efficiently computed. This result builds on the development of algorithmic and probabilistic tools to approximate the outcomes of independent Bernoulli trials, representing how customers' clicks generate random consideration sets. We believe that this fundamental contribution may find applications in additional revenue management problems.

Beyond our theoretical investigations, we demonstrate the significance of our modeling approach for practitioners. Through a collaboration with the online retailer Alibaba, we have acquired large-scale data sets describing customers' purchase history on mobile platforms, where a customized selection of products is displayed to end users. These data sets come as close as possible to those in use by data science teams in leading retail platforms (scale, feature engineering, data quality, etc.). We develop a practical estimation strategy that pieces together the use of clickstream data, machine learning methods, and maximum-likelihood estimation. We demonstrate that utilizing the click-based MNL model leads to noticeable improvements over MNL and mixed-MNL models in

[1] Any instance of the model by Gallego and Li (2017) can be efficiently reduced to an instance of click-based MNL assortment problem, while losing only an $O(\epsilon)$-fraction of the optimal revenue, for every $\epsilon \geq 0$. The model of Gallego and Li (2017) is based on the assumption that the customers' second-stage decisions follow a strict preference ranking $\sigma : [n] \to [n]$. By instantiating the click-based MNL model with preference weights $w_{\sigma(n)} = 1/\epsilon$ and $w_{\sigma(i)} = (1/\epsilon) \cdot w_{\sigma(i+1)}$ for every $i \in [n-1]$, it is not difficult to show that the expected revenue function of the resulting click-based MNL model approximates that in Gallego and Li (2017) up to an $O(\epsilon)$-factor.

term of prediction accuracy, while dramatically reducing the computational cost of the estimation process.

In what follows, we present our contributions in greater detail, with pointers to the relevant sections. We conclude by discussing the managerial implications of our work.

*Hardness results and estimation of choice probabilities.* In Section 2.1, we show that assortment optimization under the click-based MNL model is NP-hard, by relating this setting to the well-known set partition problem. Our reduction shows that the former problem is intractable even when there is only one item whose inclusion in the consideration set is random, whereas all remaining items are deterministically included. That said, as a preliminary indication that the click-based MNL model can still be rigorously analyzed, we show in Section 2.2 that the expected revenue $\mathcal{R}(S)$ of a given assortment $S \subseteq [n]$ can be efficiently estimated. Specifically, even though the screening phase of this two-stage choice model generates a distribution over exponentially-many consideration sets $C_S$, we devise a fully polynomial time approximation scheme (FPTAS) for computing the choice probabilities $\pi(i, S)$.

*Approximation scheme.* Our main algorithmic contribution comes in the form of a polynomial time approximation scheme (PTAS) for the assortment optimization problem. Formally, for any accuracy level $\epsilon > 0$, our algorithm constructs an assortment whose expected revenue is within a factor of $1-\epsilon$ of optimal, running in $O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$ time, as shown in Sections 2.3 and 4. At a high-level, our algorithmic approach proceeds from a distinction between two regimes of parameters: *unlikely items* corresponding to low click probabilities, and *likely items* corresponding to high click probabilities. Our algorithm builds on the intuition that assortment optimization becomes easier to handle when each of these settings is studied independently of the other. For unlikely items, the probability that two items are simultaneously considered is small, and thus, the substitution effects are limited. From an algorithmic standpoint, this means that, in this regime, the assortment problem can be approximately "linearized" to obtain knapsack-like formulations. In contrast, for likely items, products have non-negligible probabilities of being clicked on. Here, the instance structure is closer in spirit to that of a standard MNL-assortment optimization problem, and therefore, we utilize a revenue-ordered-like policy. Going from this intuition to a rigorous analysis requires significant technical developments, culminating in enumeration schemes and stochastic inequalities between sums of independent Bernoulli variables using Poisson-distributed surrogates, which may be of broader interest. For ease of exposition, we devote Section 3.1 to a thorough discussion of our algorithmic approach.

*Practical implementation.* In view of its running time guarantees, our approximation scheme still requires certain enhancements to be readily implemented in practice. Hence, in Section 5, we distill these ideas into three practical algorithms, ranging from a simple greedy heuristic to an enhanced version of our PTAS that can also handle cardinality-constrained instances. We then conduct an extensive set of numerical experiments in which we test the efficacy of these approaches on randomly generated instances consisting of 20 to 125 products. In this setting, for instances with 40 or fewer products, the optimal assortment can be recovered via a complete enumeration over all feasible assortments. We find that all three approaches recommend assortments that are within 1% of optimal on average; however only the aforementioned greedy heuristic and our PTAS admit practical running times. Concurrently, we examine instances with a number of products ranging from 80 to 125 to show that our PTAS can be employed on instances of practical scale. We observe that our PTAS admits running times that are under a minute on average, even when the algorithm is executed with a small accuracy level ($\epsilon \leq 5\%$).

*Case study on Alibaba's data.* In Section 6, we conduct a case study that focuses on quantifying the value of clickstream data in choice modeling settings. We make use of real historical sales data from Tmall.com, one of Alibaba's largest online marketplaces that connects third party-sellers to customers. At a high level, customers landing on any particular seller's homepage are often presented with an option to select a discount coupon. Upon clicking on this coupon, the customer is transferred to a coupon sub-page that contains six products, each of which can be purchased at a discount using the acquired coupon. For the top ten sellers (traffic-wise) across Tmall.com, our data set consists of all information related to the clicks and purchases of every customer landing on a coupon sub-page over a two-week selling period. Beyond these records, our data set includes a collection of 25 feature values associated with each offered product.

We benchmark the click-based MNL model against the standard MNL and mixed-MNL models. The latter choice models are estimated via standard maximum likelihood estimation (MLE) methods. Similarly, for the click-based MNL model, the utility parameters of the second-stage MNL instance are estimated using standard MLE. On the other hand, the prediction of click probabilities can be cast as a binary classification problem. As such, we leverage generic machine learning algorithms to estimate its click probabilities.

Consequently, we measure the fitting accuracy of these models through their respective out-of-sample log-likelihoods. We ultimately find that the click-based MNL yields 0.5% to 2.5% improvements in fitting accuracy. Given the connection between the mixed-MNL and click-based MNL models (see Appendix A), the magnitude of these gains is quite striking, and suggests that click behavior seems to provide a useful signal for the subset of products considered by customers. It is worth observing that the click-based MNL model also has an advantage from a computational standpoint. The estimation of click-based MNL models runs in less than 20 minutes, while the estimation of mixed-MNL models exceeds three hours.

We also study the extent to which the click-based MNL model leads to assortment recommendations that differ from those of the MNL model. Using semi-synthetic instances based on Alibaba's choice data, we show that the *cost of a model misspecification* can be significant. This notion quantifies the loss of revenue incurred by using an MNL-optimal assortment in place of an assortment that is optimal for the click-based MNL model, when the latter is the ground truth. To keep the paper as concise as possible, this empirical study is presented in the online companion (Appendix EC.4.3).

*Managerial implications.* While our empirical findings are tied to a specific application domain, our research generates managerial insights that could benefit retail platforms more broadly. Recent literature in revenue management has focused on developing increasingly sophisticated approaches to capture heterogeneity in customers' choice behavior using *latent processes*. This line of work is exemplified by choice models such as mixed-MNL, Markov chain, and classical consideration set-based models. In the context of the mixed-MNL model, the segment of the population to which a customer belongs is never explicitly observed in the data. Similarly, for the Markov chain model, the path of product substitutions followed by a customer before reaching to an offered product or to the no-purchase option serves to parametrize the customers' rankings, albeit not being an observable signal in practice. Lastly, as discussed in Section 1.3, the probabilistic structure of how consideration sets are formed is generally inferred from the final choices made by customers in the second stage of the choice making process.

In contrast, our work promotes a different path towards capturing customers' heterogeneity: We propose the use of clickstream data to *explicitly* define and estimate the customers' heterogeneous consideration sets. Surprisingly, this approach is competitive from a predictive standpoint even against mixed-MNL models, a family known to approximate any desired random-utility maximization choice model, while being much less costly to estimate. In other words, once customers whittle down the offered assortments to a small number of options, the ranking decisions over clicked products are sufficiently well-captured by a simple choice model such as MNL. As shown in previous literature, capturing latent heterogeneity comes at a substantial cost in terms of estimation and computation. For example, estimating Markov chain models over a large universe of featurized products remains an open question. Instead, the estimation of click-based models can be efficiently parametrized and easily combined with machine learning methods. *The main takeaway for choice modeling practitioners and researchers is that data generated by customers during their search process (clickstream data) provide powerful signals to estimate their heterogeneous preferences.* We believe that the conceptual simplicity of click-based MNL makes it a good candidate to fulfill this approach in data science practice.

### 1.3. Related work

In this section, we review directly related research, with a primary focus on papers that utilize consideration sets in modeling customers' purchasing patterns. This body of work includes assortment optimization formulations that incorporate an initial phase of consideration set formation, in addition to those that study operational decisions under consideration-set-based choice models.

*Traditional assortment optimization.* In the most well-studied version of the assortment optimization problem, the underlying choice model governing customer purchasing patterns falls under the random utility maximization (RUM) framework. In such models, customers associate a random utility with all offered products, and purchase the highest utility product. The most widespread RUM choice models under which assortment optimization problems have been considered include the MNL model (Talluri and van Ryzin 2004, Rusmevichientong et al. 2010, Davis et al. 2013), the nested logit model (Davis et al. 2014, Gallego and Topaloglu 2014, Feldman and Topaloglu 2015b), and the mixed-MNL model (Désir et al. 2014, Feldman and Topaloglu 2015a).

As previously mentioned, one apparent downside of these traditional RUM models is the underlying assumption that each customer considers purchasing all offered products. In reality, as noted by Gilbride and Allenby (2004), customers generally narrow down the products they are willing to consider based on prominent product features such as brand or price. Then, from among this smaller consideration set of products, the customer makes a purchasing decision.

*Assortment optimization with consideration sets.* The notion of a consideration set was first introduced by Howard and Sheth (1969); since then, numerous papers have found empirical evidence in support of this model for customer purchasing behavior (Lapersonne et al. 1995, Hauser and Wernerfelt 1990, Shocker et al. 1991, Mehta et al. 2003, Hauser 2014). Silk and Urban (1978), for example, measures the average size of consideration sets for everyday items such as laundry detergent, yogurt, and coffee; they find that this average size is generally between 2 and 8 products.

To our knowledge, the first to implicitly incorporate the notion of consideration sets within assortment optimization were Honhon et al. (2012). Here, the customer population is partitioned into classes, each distinguished by a unique preference list that describes a ranking on some subset of the available products. An arriving customer from a particular class will purchase the highest ranked offered product that appears on her respective preference list. Honhon et al. (2012) show that, when each preference list corresponds to a path in a binary tree of products, the assortment optimization problem can be recast as a shortest path problem. Aouad et al. (2020) build on this work by providing a general dynamic programming framework to efficiently compute optimal assortments under ranking-based models with various structural assumptions. Jagabathula and Rusmevichientong (2016) expand this version of the ranking-based model by adding a price threshold to the characterization of each customer class. In this setting, customers belonging to a particular class only consider purchasing products in their preference list that are priced below their respective price threshold. Feldman and Topaloglu (2018) consider an assortment optimization problem along similar lines, in which customers first form their consideration set based on some product feature threshold, and then make a purchase according to the MNL choice model.

A very recent work with certain common features to the present paper is that of Gallego and Li (2017), who consider assortment optimization under a ranking-based random consideration set model initially proposed by Manzini and Mariotti (2014). In this model, customers first form a random consideration set through a process that mirrors that of the click-based MNL model, namely, each product $i \in [n]$ is independently included in the customer's consideration set with probability $\lambda_i$. However, once the consideration set has been formed, customers are assumed to make a purchase according to a *single universal ranking* over all products. In other words, out of the products offered, customers purchase the highest ranked one from this single universal ranking that appears in their consideration set. Gallego and Li (2017) show how to formulate the MLE problem induced by this model as a mixed-integer nonlinear program, and provide a simple heuristic approach for computing local optima. Additionally, they prove that revenue-ordered assortments are optimal when the universal ranking contains no ties among products. When ties are allowed, the authors propose a $\frac{1}{2}$-approximation.

As explained in Section 1.1, any instance of the ranking-based model of Gallego and Li (2017) can be casted into an instance of the click-based MNL model, while preserving the expected revenue of any assortment and its related choice probabilities within any degree of accuracy. This statement holds even with ties between products, meaning in particular that we obtain a polynomial-time approximation scheme for this case as well. Furthermore, one fundamental difference between the model of Gallego and Li (2017) and our current work is that a single universal ranking cannot sensibly utilize click behavior as a proxy for the set of products that will be considered by customers. To verify this claim, consider a very simple setting, where we have only two products and only two sales data points. For both data points, both products were clicked. However, the first data point specifies that product 1 was purchased, while the second data point reveals that product 2 was purchased. Clearly, there is no universal ranking and consideration probabilities that lead to a non-zero likelihood for these two purchase events.

Turning the spotlight to papers that are directly motivated by e-commerce setting, Davis et al. (2015), Aouad and Segev (2020), and Gallego et al. (2020) consider how to optimally rank products that are relevant to a particular search query. Problems in this spirit have also been considered by Wang and Sahin (2017), Chu et al. (2020) and Derakhshan et al. (2018), who associate an explicit search cost for each product being considered. In this setting, the typical assumption is that products placed higher up on the displayed results page (i.e., ranked higher) are more likely to be clicked or purchased. Consequently, such choice models indirectly describe a certain distribution over consideration sets, each corresponding to a subset of top ranked products. This way, one captures the notion that customers start their search from the top result and generally proceed sequentially through the list of displayed results until their patience expires.

*Data-driven operations management.* Finally, our work contributes to a growing stream of research that leverages the increasing availability and granularity of historical data to inform operational decisions, such as those related to data-driven inventory management (Kunnumkal and Topaloglu 2008, Huh and Rusmevichientong 2009, Huh et al. 2011, Huh and Rusmevichientong 2009, Ban and Rudin 2019) and feature-based pricing (Cohen et al. 2020, Javanmard and Nazerzadeh 2019, Qiang and Bayati 2016). In the context of demand prediction, machine learning methods have gained wide popularity, given their ability to leverage data from a wide variety of sources to enhance their predictive efficacy. For example, Farias and Li (2019) have shown that the recovery of customer preferences via matrix completion techniques can be accelerated using comprehensive sources of customer-product interactions, such as clicks on products, clicks on the "like" button, inclusions to the shopping cart, etc. However, more accurate predictions do not necessarily yield more effective decisions. For example, Feldman et al. (2018) conducted a pilot study, comparing two algorithmic approaches for selecting the set of products displayed to customers landing on Alibaba's online marketplaces. Despite benefiting from additional features and achieving more accurate predictions, the current practice, based on sophisticated machine learning methods such as regularized logistic regression and gradient boosted decision trees, falls short against a standard MNL-based formulation, in terms of revenue performance. This study illustrates the value of a modeling framework that accurately captures key decision tradeoffs, such as product substitution effects. From this perspective, our modeling approach can be viewed as a way of augmenting the classical MNL choice model using clickstream data.

## 2. Main Theoretical Results

In what follows, we present our main theoretical results regarding the click-based MNL assortment problem. We first prove in Section 2.1 that this computational problem is NP-hard even in seemingly simple settings. In Section 2.2, we present our main algorithmic result in simplified form, by providing an efficient approximation scheme for the click-based MNL assortment problem under two auxiliary assumptions. Section 4 will be dedicated to showing how these assumptions can be completely eliminated, thus attaining a PTAS for the assortment problem in its utmost generality.

### 2.1. Hardness result

In what follows, we prove that the assortment optimization problem, as formally defined in Section 1.1, is NP-hard. Our approach leverages ideas that are similar in spirit to the reduction proposed by Feldman and Topaloglu (2018) for an MNL-based choice model with nested consideration sets. To this end, we focus on the feasibility version of the former problem, where an additional revenue threshold $K$ is specified. The goal of the *assortment feasibility problem* is to decide whether there exists an assortment with an expected revenue of at least $K$.

THEOREM 1. *The assortment feasibility problem is NP-complete.*

The proof is based on a reduction from set partition, which is one of Karp's 21 NP-complete problems (Karp 1972). This reduction is presented in Appendix B.1. Interestingly, we map the set partition instances to click-based MNL settings in which all products are clicked on with probability 1, except for a single product, which is clicked on with probability 1/2. Hence, this construction proves that the computational hardness streams from the basic probabilistic structure of the click-based MNL choice model, where the customer's random consideration set is formed through independent Bernoulli random variables.

### 2.2. FPTAS for computing choice probabilities

In addition to the computational hardness of the assortment optimization problem, the challenges surrounding the click-based MNL model start with the estimation of its customers' choice probabilities. Indeed, each assortment $S$ induces a distribution over $2^{|S|}$ consideration sets, each of which occurs with positive probability. Therefore, computing the choice probability $\pi(i,S)$ of an item $i$ in any given assortment $S$ is by no means straightforward, and we are not aware of any efficient method for computing $\pi(i,S)$ in an exact way. Nevertheless, as stated in the next claim, this quantity can be deterministically estimated within any degree of accuracy in polynomial time. In what follows, we let $w_{\min}$ and $w_{\max}$ denote the minimal and maximal MNL weights, respectively.

THEOREM 2. *For any $\epsilon \in (0,1)$, there is a deterministic $O(\frac{n^2}{\epsilon} \cdot \log(n \frac{w_{\max}}{w_{\min}}))$ time algorithm that computes an estimate $\tilde{\pi}(i,S)$ for the choice probability $\pi(i,S)$ satisfying*

$$\pi(i,S) \leq \tilde{\pi}(i,S) \leq (1+\epsilon) \cdot \pi(i,S) \ .$$

To establish Theorem 2, we formulate in Appendix B.2 an approximate dynamic program that estimates the purchase probabilities up to a factor of $1+\epsilon$. In this formulation, items are sequentially processed in an arbitrary order with a state variable describing the cumulative preference weights of all items considered before reaching the current state. In the remainder of the paper, we use the dynamic programming method of Theorem 2 as a black-box to estimate the choice probabilities under the click-based MNL model. This positive result provides an initial indication that the click-based MNL assortment problem may be computationally tractable, despite the NP-hardness reduction stated in Section 2.1.

### 2.3. Approximation scheme in a simplified setting

Here, we discuss our main algorithmic result for the click-based MNL assortment problem. We begin by stating a couple of structural assumptions regarding the click-based MNL instance in question. These assumptions allows us to focus on a slightly more structured setting that simplifies the design of an approximation scheme. We emphasize that these assumptions are by no means required; we explain in Section 4 how to bypass them using additional algorithmic ideas.

*Structural assumptions.* Let $\epsilon > 0$ be an extra input parameter to our approximation scheme, standing for its desired accuracy. For every integer $p$, we denote by $\Lambda_p$ the collection of items $i \in [n]$ with a consideration probability $\lambda_i \in [\epsilon \cdot (1+\epsilon)^p, \epsilon \cdot (1+\epsilon)^{p+1})$. Similarly, for every integer $q \geq 1$, we denote by $\mathcal{W}_q$ the collection of items $i \in [n]$ with a preference weight $w_i \in (\frac{\epsilon}{n} \cdot (1+\epsilon)^{q-1}, \frac{\epsilon}{n} \cdot (1+\epsilon)^q]$, while $\mathcal{W}_0 = \{i \in [n] : w_i \leq \frac{\epsilon}{n}\}$ is the collection of remaining items.

Assumption 1. *For every $(p, q) \in \mathbb{Z} \times \mathbb{Z}^+$, each item $i \in \Lambda_p \cap \mathcal{W}_q$ has a consideration probability of exactly $\lambda_i = \epsilon \cdot (1+\epsilon)^p$ and a preference weight of exactly $w_i = \frac{\epsilon}{n} \cdot (1+\epsilon)^q$.*

Assumption 1 essentially states that the click probabilities $\{\lambda_i\}_{i \in [n]}$ and preference weights $\{w_i\}_{i \in [n]}$ take values in a restricted set that explicitly depends on the accuracy parameter $\epsilon > 0$. It is not difficult to see that Assumption 1 can be enforced by slightly rounding the input parameters. However, this approach requires us to analyze how slight modifications of the parameters affect the expected revenue function. In other words, a *sensitivity analysis* is a crucial ingredient to justify Assumption 1. We provide such line of analysis in Section 4, where we present our approximation scheme for the click-based MNL assortment problem in full generality.

Our second assumption imposes further restrictions on the range of indices $q$ for which the weight class $\mathcal{W}_q$ is non-empty. Here, $Q_{\min}$ is the minimal index $q \in \mathbb{Z}^+$ for which $\mathcal{W}_q \neq \emptyset$ and let $Q_{\max}$ be the maximal such index.

Assumption 2. *There are only $O(\frac{1}{\epsilon^2} \cdot \log n)$ non-empty weight classes, meaning that $Q_{\max} - Q_{\min} = O(\frac{1}{\epsilon^2} \cdot \log n)$.*

The above assumption restricts our attention to a setting we refer to as the *bounded-ratio* regime. Intuitively, the bounded-ratio property ensures that there is limited variability in the preference weights of the underlying products, specifying that the ratio between extremal preference weights $\frac{w_{\max}}{w_{\min}}$ is polynomially bounded with respect to the number of items $n$. Clearly, this condition is likely to be satisfied by instances encountered in practice. Nevertheless, we develop in Section 4 algorithmic methods for the assortment optimization problem in its utmost generality, without the bounded-ratio assumption.

*Main algorithmic result.* Focusing on the setting described by Assumptions 1 and 2, we present our main algorithmic result. One notable nuance is that the algorithm we devise constructs a random assortment, rather than a deterministic one. In other words, the assortment decision is sampled from a distribution over assortments; this random experiment is independent from the customer's purchasing decision, as prescribed by the click-based MNL choice model. Moreover, we restrict attention to the class of *decomposable assortments*, which are generated by specifying a certain probability for each product to be included in the assortment or not. That is, given a sequence of independent Bernoulli random variables $B = (B_1, \dots, B_n)$, the realization of the decomposable assortment $B$ is given by $S = \{i \in [n] : B_i = 1\}$. In what follows, we denote probabilities and expectations with regards to this distribution by $\Pr_{S\sim B}[\cdot]$ and $\mathbb{E}_{S\sim B}[\cdot]$, respectively. Furthermore, we denote by $S^*$ an optimal (deterministic) assortment, which is fixed throughout our exposition. With these definitions at hand, we next state our main theorem.

THEOREM 3. *Suppose that Assumptions 1 and 2 are satisfied. For any $\epsilon > 0$, we can determine a decomposable assortment $B$ such that $\mathbb{E}_{S\sim B}[\mathcal{R}(S)] \geq (1-\epsilon)\cdot\mathcal{R}(S^*)$. Our algorithm runs in time $O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$.*

The theoretical significance of Theorem 3 resides in showing that the click-based MNL assortment problem can be approximated within any degree of accuracy, when Assumptions 1 and 2 are in place. At this stage, it is important to point out that decomposable assortment decisions (as stated in Theorem 3) can be derandomized via the standard method of conditional expectations (Alon and Spencer 2016, Chap. 16). For completeness, we explain in Appendix B.3 how to convert any decomposable assortment $B$ into a deterministic assortment $S$ without any loss in optimality. We devote Section 3 to an exposition of the algorithmic approach used to establish Theorem 3.

## 3. Proof of Theorem 3: Algorithmic Approach

In what follows, we present our approximation scheme for the click-based MNL assortment problem, as stated in Theorem 3. While its performance guarantee is formally established in Appendix D, here, we develop the main technical insights that motivate our algorithmic approach. For ease of exposition, we begin by providing a technical outline of this approach.

### 3.1. Technical outline

*Decomposition method.* The starting point of our algorithm is to decompose the collection of items based on the magnitude of their consideration probabilities. Let $P_{\min}$ be the minimal index $p \in \mathbb{Z}$ for which $\Lambda_p \neq \emptyset$ and let $P_{\max}$ be the maximal such index. For simplicity of notation, we use $\mathcal{Q} = [Q_{\min}, Q_{\max}]$, with $\{a_x\}_X$ as an abbreviation for $\{a_x\}_{x\in X}$.

With this notation at hand, we distinguish between *unlikely* and *likely* items:

- An item $i \in [n]$ is said to be unlikely if $\lambda_i \in [0, \epsilon]$. Thus, the unlikely items are precisely those in classes $\{\Lambda_p\}_{[P_{\min},0]}$. For every $q \in \mathcal{Q}$, we will use $\mathcal{U}_q = \bigcup_{p \in [P_{\min},0]} (\Lambda_p \cap \mathcal{W}_q)$ to denote the set of unlikely items with a preference weight of $w^q = \frac{\epsilon}{n} \cdot (1+\epsilon)^q$. In addition, $\mathcal{N}_{\text{unlike}} = \bigcup_{q \in \mathcal{Q}} \mathcal{U}_q$ will denote collection of all unlikely items.
- An item $i \in [n]$ is said to be likely if $\lambda_i \in (\epsilon, 1]$. The likely items are precisely those in classes $\{\Lambda_p\}_{[P_{\max}]}$. Consequently, $\mathcal{N}_{\text{likely}} = \bigcup_{p=1}^{P_{\max}} \Lambda_p$ will denote the collection of all likely items.

In what follows, we explain why the distinction between unlikely and likely items is very useful in approximating the click-based MNL assortment problem. The main implication of this decomposition is that we develop a specialized algorithm for each subset of items $\mathcal{N}_{\text{unlike}}$ and $\mathcal{N}_{\text{likely}}$. In particular, we argue that each of these two regimes (unlikely and likely items) have distinctive properties that simplify the formation of near-optimal assortment decisions.

*Sketch of the approximation scheme.* For the unlikely items, we argue that the click-based MNL assortment problem can be approximately *linearized* and consequently cast as a *min-knapsack* problem. The corresponding algorithm, referred to as `MinKnapsack`$(\cdot)$, is described in Section 3.2. For the likely items, we introduce the property of *revenue-ordered by class* assortments, which is shown to be satisfied by optimal assortments. By exploiting this property, we argue that likely items can be selected using a greedy procedure. The corresponding algorithm, referred to as `Greedy`$(\cdot)$, is described in Section 3.3. Finally, in Section 3.4, we provide a succinct summary of our approximation scheme, which combines `MinKnapsack`$(\cdot)$ and `Greedy`$(\cdot)$; we also discuss its computational performance.

*Guessing of input parameters.* Each of the specialized algorithms `MinKnapsack`$(\cdot)$ and `Greedy`$(\cdot)$ requires us to specify certain *input parameters*. Intuitively, these parameters serve to coordinate both algorithms and to capture a first-order approximation of how the optimal assortment $S^*$ is structured. To this end, for every $q \in \mathcal{Q}$, we define $c_q^* = \sum_{i \in S^* \cap \mathcal{U}_q} r_i \cdot \pi(i, S^*)$ as the *revenue contribution* of the class of items $\mathcal{U}_q$ in the optimal assortment $S^*$. Similarly, for every $(q,p) \in \mathcal{Q} \times [P_{\min}, P_{\max}]$, we define $c_{p,q}^* = \sum_{i \in S^* \cap \mathcal{W}_q \cap \Lambda_p} r_i \cdot \pi(i, S^*)$ as the revenue contribution of the class of items $\mathcal{W}_q \cap \Lambda_p$. In addition, for every $S \subseteq [n]$ and $q \in \mathcal{Q}$, let $\alpha(S,q) = \mathbb{E}[\frac{w^q}{1+w^q+w(C_S)}]$ be the *coefficient* of the class of items $\mathcal{U}_q$ in the assortment $S$. While the meaning of these coefficients will be revealed later on, these quantities will be critical to approximate the choice probabilities of unlikely items.

Clearly, the quantities $\{(c_q^*, \alpha(S^*, q)\}_{\mathcal{Q}}$ and $\{c_{p,q}^*\}_{[P_{\max}] \times \mathcal{Q}}$ are not known to our algorithm since they depend on the optimal assortment $S^*$. Nevertheless, our specialized algorithms will use input parameters that "closely approximate" these quantities. To be specific, `MinKnapsack`$(\cdot)$ is executed with input parameters $\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}}$ as a proxy for $\{(c_q^*, \alpha(S^*, q)\}_{\mathcal{Q}}$, and `Greedy`$(\cdot)$ is executed with $\{\hat{c}_{p,q}\}_{[P_{\max}] \times \mathcal{Q}}$ as a proxy for $\{c_{p,q}^*\}_{[P_{\max}] \times \mathcal{Q}}$.

At first glance, it is unclear how we may efficiently "guess" these input parameters. Indeed, from a computational standpoint, there are exponentially many possible values for $\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}}$ and $\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}$. However, the important observation is that the number of input parameters is relatively small, i.e., $|\mathcal{Q}| = Q_{\max} - Q_{\min} + 1 = O(\frac{1}{\epsilon^2}\cdot \log n)$ by Assumption 2, and $P_{\max} = O(\frac{1}{\epsilon}\cdot\log\frac{1}{\epsilon})$ since the consideration probabilities of likely items are within the range $(\epsilon, 1]$. Thus, basic counting arguments show that the input parameters can be tuned by enumerating over polynomially many vectors, as formally stated in the next claim, whose proof is deferred to Appendix C.1. In what follows, we use $\mathcal{L} = |\mathcal{Q}|\cdot(P_{\max}+1)$ for ease of notation.

Claim 1. *There is an efficiently constructible collection of input parameters* $\Omega$ *with* $|\Omega| = O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$ *such that there exists* $(\{\hat{c}_q,\hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}) \in \Omega$ *satisfying:*

1. $\hat{c}_{p,q} \le c^*_{p,q} < \hat{c}_{p,q} + \frac{2\epsilon}{\mathcal{L}}\cdot\mathcal{R}(S^*)$ *for every* $(p,q)\in[P_{\max}]\times\mathcal{Q}$ ,
2. $\hat{c}_q \le c^*_q < \hat{c}_q + \frac{2\epsilon}{\mathcal{L}}\cdot\mathcal{R}(S^*)$ *for every* $q\in\mathcal{Q}$ ,
3. $(1-\epsilon)^2\cdot\hat{\alpha}_q \le \alpha(S^*,q) \le (1-\epsilon)\cdot\hat{\alpha}_q$ *for every* $q\in\mathcal{Q}$ .

The above claim states that by enumerating over $\Omega$, we will consider in particular input parameters $\{\hat{c}_q,\hat{\alpha}_q\}_{\mathcal{Q}}$ and $\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}$ that closely approximate $\{(c^*_q, \alpha(S^*,q)\}_{\mathcal{Q}}$ and $\{c^*_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}$. Indeed, Property 1 and 2 ensure that the revenue contribution parameters are accurate, while Property 3 ensures that the coefficient parameters are accurate; the latter have a slight negative bias that is required in our subsequent analysis.

### 3.2. Algorithm for unlikely items

We begin by providing technical insights and intuition about our approach for handling unlikely items. In particular, we explain why the assortment decisions over $\mathcal{N}_{\text{unlike}}$ can be approximately captured as min-knapsack problems. Next, we formally describe the resulting algorithm `MinKnapsack`$(\cdot)$.

*Linearization ideas.* Let us focus our attention on the class of unlikely items $\mathcal{U}_q$ with a preference weight of $w^q$. Generally speaking, the tradeoff in assortment decisions is between generating "more revenue" by selecting additional items and mitigating the "cannibalization" of the items' choice probabilities. In what follows, we argue that this tradeoff can be approximately formulated through an instance of the min-knapsack problem.

We first focus on the "cannibalization" component in the above tradeoff. Based on representation (1) of the choice probabilities, for every assortment $S\subseteq[n]$ and item $i\in S$, we have:

$$\pi(i,S) = \lambda_i\cdot\mathbb{E}\left[\frac{w_i}{1+w_i+w(C_S^{-i})}\right] .$$

Intuitively, the amount of "cannibalization" exerted by the items $S\setminus\{i\}$ on item $i$'s choice probability scales with the random variable $w(C_S^{-i})$. To be specific, the quantity $\mathbb{E}[\frac{w_i}{1+w_i+w(C_S^{-i})}]$ is convex non-increasing with respect to $w(C_S^{-i})$. The items of $S \cap \mathcal{U}_q$ have a total contribution to $w(C_S^{-i})$ of exactly $w^q \cdot |C_{S\cap\mathcal{U}_q}^{-i}|$. Here, the random variable $|C_{S\cap\mathcal{U}_q}^{-i}|$ is a sum of independent Bernoulli outcomes, representing the number of items considered amongst $S \cap \mathcal{U}_q$ other than item $i$.

The main technical insight for unlikely items is that the distribution of $|C_{S\cap\mathcal{U}_q}^{-i}|$ can be approximated using a single *summary statistic*, which is its expectation $\mathbb{E}[|C_{S\cap\mathcal{U}_q}^{-i}|]$. This surprising insight proceeds from the next lemma, showing that sums of independent Bernoulli outcomes are "well-approximated" by Poisson surrogates. In what follows, for $\lambda \in [0,1]$, we denote by $P(\lambda)$ a Poisson random variable with parameter $\lambda$. Two random variables $Z_1$ and $Z_2$ are said to be in the convex non-increasing order $Z_1 \succeq_{\text{cni}} Z_2$ if, for every convex non-increasing function $\phi$, we have $\mathbb{E}[\phi(Z_1)] \geq \mathbb{E}[\phi(Z_2)]$, provided these expectations exist (Shaked and Shanthikumar 2007, Chap. 3).

Lemma 1 **(Poissonization)**. *Suppose that $\epsilon \in [0, \sqrt{2}-1]$. Let $Y = \sum_{i=1}^{k} X_i$ where $\{X_i\}_{[k]}$ are independent Bernoulli random variables with success probabilities $\{\lambda_i\}_{[k]} \in [0,\epsilon]^k$. Letting $\lambda = \sum_{i=1}^{k} \lambda_i$, we have $P(\lambda) \succeq_{\text{cni}} Y \succeq_{\text{cni}} P((1+\epsilon)\cdot\lambda)$.*

This claim shares similarities with the Poissonization method (Le Cam 1960). That said, rather than using a concentration inequality with respect to a Poisson limiting distribution, our proof proceeds from basic properties of stochastic orders, which are presented in Appendix C.2. An important requirement of Lemma 1 is that the success probabilities of the Bernoulli random variables are at most $\epsilon$, which is in perfect alignment with the consideration probabilities of the unlikely items. Based on Lemma 1, we will be able to argue that the cannibalization effects due to the assortment decisions $S \subseteq \mathcal{U}_q$ can be "mitigated" by ensuring that $\mathbb{E}[|C_{S\cap\mathcal{U}_q}|]$ is sufficiently small, noting that $\mathbb{E}[|C_{S\cap\mathcal{U}_q}|]$ and $\mathbb{E}[|C_{S\cap\mathcal{U}_q}^{-i}|]$ are $\epsilon$-close from each other. Moreover, the summary statistic $\mathbb{E}[|C_{S\cap\mathcal{U}_q}|] = \sum_{j\in S\cap\mathcal{U}_q} \lambda_j$ is a linear function with respect to the assortment decisions $S \cap \mathcal{U}_q$.

Next, we turn our attention to the "revenue" component in the above-mentioned tradeoff. How do we quantify the revenue contribution of an individual unlikely item? By equation (1), in order to compute the choice probability $\pi(i,S)$ for an item $i \in S \cap \mathcal{U}_q$, we need to estimate the quantity $\mathbb{E}[\frac{w_i}{1+w_i+w(C_S^{-i})}]$, which clearly depends on the assortment $S$ and the precise values of $q \in \mathcal{Q}$ and $p \in [P_{\min}, 0]$ that govern the values of $w_i$ and the distribution of $w(C_S^{-i})$, respectively. Interestingly, in the next claim, we argue that the dependency on $p \in [P_{\min}, 0]$ plays a minor role for unlikely items. To this end, recall that $\alpha(S,q) = \mathbb{E}[\frac{w^q}{1+w^q+w(C_S)}]$ is the coefficient of the class of items $\mathcal{U}_q$ in the assortment $S$, a quantity that only depends on $S$ and $q$.

Lemma 2. *For all $S \subseteq [n]$ and $i \in S \cap \mathcal{U}_q$, we have $(1-\epsilon)\cdot\pi(i,S) \leq \lambda_i \cdot \alpha(S,q) \leq \pi(i,S)$.*

The proof is provided in Appendix C.3. Consequently, the contribution of items $i \in S \cap \mathcal{U}_q$ to the expected revenue $\mathbb{E}[\mathcal{R}(S)]$ is well-approximated by the expression $\sum_{i\in S\cap\mathcal{U}_q} r_i\lambda_i \cdot \alpha(S,q)$. The important observation is that, should we use the input parameter $\hat{\alpha}_q$ as a constant proxy for $\alpha(S,q)$, the resulting expression $\sum_{i \in S \cap \mathcal{U}_q} r_i \lambda_i \cdot \hat{\alpha}_q$ is linear with respect to $S$. This linearization of the revenue contribution (i.e., viewing the coefficient $\hat{\alpha}_q$ as fixed rather than dependent on the assortment decisions) will be justified in our subsequent analysis.

All in all, the above observations suggest that the assortment decisions within $\mathcal{U}_q$ can be approximated by a min-knapsack problem with the following linear objective and constraint: Minimize $\mathbb{E}[|C_{S \cap \mathcal{U}_q}|]$ by selecting $S \cap \mathcal{U}_q$ subject to $\sum_{i \in S \cap \mathcal{U}_q} r_i \lambda_i \cdot \hat{\alpha}_q \geq \hat{c}_q$, where $\hat{\alpha}_q > 0$ is the chosen coefficient for the class of items $\mathcal{U}_q$ and $\hat{c}_q \geq 0$ is the desired revenue contribution from this class of items. Here, the minimization of $\mathbb{E}[|C_{S \cap \mathcal{U}_q}|]$ allows us to mitigate the cannibalization effects, while the linear constraint $\sum_{i \in S \cap \mathcal{U}_q} r_i \lambda_i \cdot \hat{\alpha}_q \geq \hat{c}_q$ ensures that a sufficient revenue contribution is generated.

*Algorithm for unlikely items.* To summarize the above discussion, we formally present our specialized algorithm `MinKnapsack`$(\cdot)$, which proceeds by solving a collection of min-knapsack problems. Here, for every $q \in \mathcal{Q}$, we assume that a coefficient $\hat{\alpha}_q$ and a revenue contribution $\hat{c}_q$ for the class of items $\mathcal{U}_q$ are given according to the enumeration procedure of Claim 1. Given these parameters, the algorithm `MinKnapsack`$(\cdot)$ constructs a decomposable assortment $B_{\text{unlike}}$ over the collection of unlikely items, i.e., $B_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$.

First, for each $q \in \mathcal{Q}$, we solve the following min-knapsack instance: Find a subset of items $U \subseteq \mathcal{U}_q$ that minimizes $\mathbb{E}[|C_U|] = \sum_{i \in U} \lambda_i$ subject to the linear constraint $\sum_{i \in U} \gamma_i \geq (1-\epsilon) \cdot \frac{n}{\epsilon}$, where $\gamma_i = \lfloor \frac{n \lambda_i r_i \hat{\alpha}_q}{\epsilon \hat{c}_q} \rfloor$. To gain some insight into the latter constraint, note that by eliminating floors, it can be rearranged as $\hat{\alpha}_q \cdot \sum_{i \in U} r_i \lambda_i \geq (1-\epsilon) \cdot \hat{c}_q$. Hence, the constraint precisely ensures that the desired revenue contribution of $\hat{c}_q$ is approximately generated by the items $U \subseteq \mathcal{U}_q$. From a computational perspective, the min-knapsack instance constructed above can be solved to optimality in $O(\frac{n^2}{\epsilon})$ time via dynamic programming since the $\gamma_i$-parameters are integral and upper-bounded by $\frac{n}{\epsilon}$; e.g., see Vazirani (2013, Chap 8.1).

Now, let $U_q \subseteq \mathcal{U}_q$ be the resulting optimal subset of items; if the min-knapsack instance is infeasible, we simply pick $U_q = \emptyset$. Next, we define the decomposable assortment $B^q$, where each item $i \in U_q$ is independently picked with probability $1-\epsilon$; namely, $B^q = (B^q_i)_{i \in U_q}$ is a collection of IID Bernoulli random variables with probability of success of $1-\epsilon$. Finally, our algorithm returns the decomposable assortment $B_{\text{unlike}}$, defined by concatenating the assortments $B^q$ over $q \in \mathcal{Q}$.

### 3.3. Algorithm for likely items

We now turn our attention to consider likely items. We explain why the optimal assortment decisions over $\mathcal{N}_{\text{likely}}$ can be approximated using a greedy procedure. Next, we formally describe the resulting algorithm `Greedy`$(\cdot)$.

*Revenue-ordered property.* Following Assumption 1, we establish a generalized revenue-ordered-like property for the click-based MNL assortment problem. We say that an assortment $S$ is *revenue-ordered by class* if, for every class $\mathcal{W}_q \cap \Lambda_p$, the subset $S \cap \mathcal{W}_q \cap \Lambda_p$ consists of the $|S \cap \mathcal{W}_q \cap \Lambda_p|$ highest-revenue items in $\mathcal{W}_q \cap \Lambda_p$, breaking ties arbitrarily.

LEMMA 3. *There exists an optimal assortment $S^*$ which is revenue-ordered by class.*

It is worth observing that, although we subsequently make use of Lemma 3 to develop an algorithm for likely items, the revenue-ordered by class property holds for all items, irrespective of whether they are likely or unlikely. The proof is presented in Appendix C.4. We next describe how this property is utilized from an algorithmic perspective. Suppose that we are given the desired revenue contribution $\hat{c}_{p,q} \geq 0$ within each class of likely items $\mathcal{W}_q \cap \Lambda_p$. By Lemma 3, higher price items should always be offered before lower price items within each class. Consequently, our algorithm exploits this observation by greedily selecting items within each class of items $\mathcal{W}_q \cap \Lambda_p$ by order of decreasing prices, until reaching the desired revenue contribution $\hat{c}_{p,q}$. This greedy approach is described in formal terms next.

*Algorithm for likely items.* Suppose that a revenue contribution $\hat{c}_{p,q}$ is given for each class of likely items $\mathcal{W}_q \cap \Lambda_p$. We also assume that the decomposable assortment $B_{\rm unlike}$ over the unlikely items is already determined using the algorithm $\texttt{MinKnapsack}(\cdot)$ of Section 3.2. Given these inputs, the algorithm $\texttt{Greedy}(\cdot)$ constructs a decomposable assortment $B_{\rm likely}$ over the set of likely items, i.e., $B_{\rm likely} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}, B_{\rm unlike})$.

Starting from the decomposable assortment $B_{\rm unlike}$, by order of decreasing revenues, we greedily add a single likely item from any class $\Lambda_p \cap \mathcal{W}_q$ that violates the following termination criterion. Specifically, the algorithm terminates when, for every class $\Lambda_p \cap \mathcal{W}_q$, either we have added all of its items, or discover that the items already chosen from $\Lambda_p \cap \mathcal{W}_q$ contribute at least $(1-\epsilon)\cdot\hat{c}_{p,q}$ to the expected revenue. Here, the customers' choice probabilities are estimated using the FPTAS of Theorem 2 with an accuracy level of $\epsilon$. To sum-up, the algorithm $\texttt{Greedy}(\cdot)$ proceeds as follows:

1. We start with $S_0 = \emptyset$.
2. Given $S_t$, our algorithm halts if the following termination criterion is met: For every $p \in [P_{\max}]$ and $q \in [Q_{\min}, Q_{\max}]$, either $S_t \cap \Lambda_p \cap \mathcal{W}_q = \Lambda_p \cap \mathcal{W}_q$ or
$$\sum_{i \in S_t \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S\sim B_{\rm unlike}}[\tilde{\pi}(i, S \cup S_t)] \geq (1-\epsilon)\cdot \hat{c}_{p,q} \ . \tag{2}$$
Otherwise, let $\Lambda_p \cap \mathcal{W}_q$ be a class for which this criterion does not hold. Letting $i_t$ be the maximum price item in $(\Lambda_p \cap \mathcal{W}_q) \setminus S_t$, we define the next assortment as $S_{t+1} = S_t \cup \{i_t\}$.
3. Let $S_T$ be the assortment obtained in step 2, at the last iteration $T$. Our algorithm returns the decomposable assortment $B_{\rm likely} = (B_i)_{i \in S_T}$, where each $B_i$ is a Bernoulli random variable with a probability of success of 1.

*Additional remarks.* From a computational perspective, the number of iterations in step 2 is at most $n$, and thus, this algorithm makes $O(n^2)$ calls to the FPTAS of Theorem 2 to estimate the choice probabilities, which results in an overall running time of $O(\frac{1}{\epsilon^2} \cdot n^4 \log n)$ by Assumption 2. It is important note that this FPTAS was devised for deterministic assortments, while $\texttt{Greedy}(\cdot)$ generates decomposable assortments. That said, our FPTAS is equally applicable to decomposable assortments using the derandomization method described in Appendix B.3.

### 3.4. Summary of the approximation scheme

We are now ready to fully describe our approximation scheme for the click-based MNL assortment problem in the setting of Theorem 3. This algorithm proceeds in three steps:

1. Generate the collection of input parameters $\Omega$.
2. For every $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}] \times \mathcal{Q}}) \in \Omega$:
    (a) Compute $B_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$.
    (b) Compute $B_{\text{likely}} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}] \times \mathcal{Q}}, B_{\text{unlike}})$.
3. Return the decomposable assortment $B = B_{\text{unlike}} \cup B_{\text{likely}}$ of maximum expected revenue out of those generated in Step 2.

Hence, this approximation scheme combines our specialized algorithms for unlikely and likely items in Steps 2(a) and 2(b). In Sections 3.2 and 3.3, we have provided technical insights and intuitive explanations as to why these specialized algorithms, $\texttt{MinKnapsack}(\cdot)$ and $\texttt{Greedy}(\cdot)$, are employed for each class of items. This line of reasoning is formalized in Appendix D, where we rigorously show that the resulting decomposable assortment $B$ indeed satisfies the guarantees stated in Theorem 3, i.e., $\mathbb{E}_{S \sim B}[\mathcal{R}(S)] = (1 - O(\epsilon)) \cdot \mathcal{R}(S^*)$.

From a computational standpoint, our algorithm runs in $O(n^{O(\frac{1}{\epsilon^4} \log \frac{1}{\epsilon})})$ time. Indeed, as explained in Sections 3.2 and 3.3, Steps 2(a) and 2(b) are very efficient for every instantiation of the input parameters, with an overall running time of $O(\frac{1}{\epsilon^2} \cdot n^4 \log n)$. Thus, the main computational bottleneck comes from the enumeration over all input parameters in $\Omega$. Specifically, the running time bound of Theorem 3 immediately follows from the upper bound on $|\Omega|$ established in Claim 1.

This computational analysis suggests that our algorithmic approach could potentially be ineffective in practice in terms of speed. Nevertheless, we will demonstrate in Section 5 that practical variants of our enumeration strategy can be implemented at very large scale, while preserving our performance guarantees. Before delving into a computational study of our approximation scheme, we explain in Section 4 why Assumptions 1 and 2 can be considered without loss of generality.

## 4. Approximation Scheme in a General Setting

To complete our theoretical investigation of the click-based MNL assortment problem, we explain how to relax Assumptions 1 and 2 and derive an unconditional polynomial-time approximation scheme. Hence, our most general algorithmic result is stated in the following theorem.

THEOREM 4. *For every $\epsilon > 0$, we can compute an assortment $S$ such that $\mathcal{R}(S) \geq (1-\epsilon)\cdot\mathcal{R}(S^*)$. The running time of our algorithm is $O((\frac{n}{\epsilon})^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$.*

Theorem 4 states that there exists a polynomial-time approximation scheme for the click-based MNL assortment problem, without any assumption whatsoever. As such, this result tightly complements the NP-hardness result presented in Section 2.1, and quite surprisingly holds for arbitrary instances of the problem. Due to lengthy technical details, the proof of Theorem 4 is presented in Appendix EC.2 of the online companion. At a high level, our line of reasoning shows that, given an instance of the click-based MNL assortment problem, we can construct a nearly equivalent instance in which Assumptions 1 and 2 are satisfied, and thus, the algorithmic approach of Theorem 3 is applicable. In the remainder of this section, we heuristically describe the main ingredients of this proof.

*Informal proof ideas.* As mentioned in Section 2.3, eliminating Assumption 1 requires a careful sensitivity analysis. Note that the choice probabilities prescribed by the click-based MNL model can be expressed as polynomials of degree $n$ with respect to the parameters $\{\lambda_i\}_{i\in[n]}$ by expanding the probabilities associated with distinct realizations of the consideration set $\mathcal{C}_S$. At first glance, this functional form is not well-behaved with respect to small perturbations of the consideration probabilities $\{\lambda_i\}_{i\in[n]}$. That said, we able to derive strong sensitivity bounds by exploiting the probabilistic structure of the choice model.

In what follows, we outline our sensitivity analysis of the expected revenue function. Suppose we are given preference weights $\{w_i\}_{i\in[n]}$ and consideration probabilities $\{\lambda_i\}_{i\in[n]}$ that possibly violate Assumption 1. Recall that, for every integer $p$, we define $\Lambda_p$ as the collection of items $i \in [n]$ with a consideration probability $\lambda_i \in [\epsilon\cdot(1+\epsilon)^p, \epsilon\cdot(1+\epsilon)^{p+1})$, and, for every integer $q \geq 1$, $\mathcal{W}_q$ is the collection of items $i \in [n]$ with a preference weight $w_i \in (\frac{\epsilon}{n}\cdot(1+\epsilon)^{q-1}, \frac{\epsilon}{n}\cdot(1+\epsilon)^q]$, while $\mathcal{W}_0 = \{i \in [n] : w_i \leq \frac{\epsilon}{n}\}$. Next, we slightly round these parameters in order to conform them with Assumption 1. Namely, for each item $i$, its rounded weight $\tilde{w}_i$ is defined as $\tilde{w}_i = \frac{\epsilon}{n}\cdot(1+\epsilon)^q$ if $w_i \in \mathcal{W}_q$ for some $q \geq 1$, and $\tilde{w}_i = \frac{\epsilon}{n}$ if $w_i \in \mathcal{W}_0$. The rounded consideration probability $\tilde{\lambda}_i$ of each item $i$ is defined by $\tilde{\lambda}_i = \epsilon\cdot(1+\epsilon)^p$, where $p \in \mathbb{Z}$ is the unique index for which $\lambda_i \in \Lambda_p$. Finally, the rounded selling price of item $i$ is set to $\tilde{r}_i = \frac{w_i}{\tilde{w}_i}\cdot r_i$.

*How does this rounding procedure affect the click-based MNL assortment problem?* To answer this question, let $\tilde{\mathcal{R}}(\cdot)$ denote the expected revenue function with respect to the rounded inputs $\{\tilde{\lambda}_i\}_{i\in[n]}$, $\{\tilde{w}_i\}_{i\in[n]}$ and $\{\tilde{r}_i\}_{i\in[n]}$. For every assortment $S \subseteq [n]$, we consider the decomposable assortment $B^S$, where each item $i \in S$ is picked with probability $1-\epsilon$ whereas items $i \in [n]\setminus S$ are not picked at all. The next claim establishes our sensitivity bounds for $\tilde{\mathcal{R}}(S)$ with regards to $\mathcal{R}(S)$ and $\mathbb{E}_{S'\sim B^S}[\mathcal{R}(S')]$.

LEMMA 4. *For every $S \subseteq [n]$, we have $(1-2\epsilon)\cdot\mathcal{R}(S) \leq \tilde{\mathcal{R}}(S) \leq (1+3\epsilon)\cdot\mathbb{E}_{S'\sim B^S}[\mathcal{R}(S')]$.*

This claim implies that the click-based MNL assortment problem is well-behaved with respect to our alterations of the input parameters. Indeed, given an $\alpha$-approximation $\tilde{S}$ of the modified instance, the corresponding decomposable assortment $B^{\tilde{S}}$ satisfies

$$\mathbb{E}_{S' \sim B^{\tilde{S}}}\left[\mathcal{R}\left(S'\right)\right] \geq (1-3\epsilon) \cdot \tilde{\mathcal{R}}(\tilde{S}) \geq (1-3\epsilon)\,\alpha \cdot \tilde{\mathcal{R}}(S^*) \geq (1-5\epsilon)\,\alpha \cdot \mathcal{R}(S^*) \ ,$$

where the first and last inequalities immediately follow from Lemma 4. Hence, restricting attention to modified instances, which satisfy Assumption 1, results in only an $O(\epsilon)$ loss of optimality.

The missing piece to complete the proof of Theorem 4 is to argue that the bounded-ratio regime (Assumption 2) can be similarly enforced with a negligible loss of optimality. The proof ideas here are significantly more complex. Unlike Assumption 1, we are not aware of any simple transformation that makes arbitrary instances compatible with the bounded-ratio setting, while preserving optimality up to an $O(\epsilon)$ factor. Our approach is to decompose the original instance of the click-based MNL assortment problem into independent subproblems that satisfy Assumption 2. This decomposition is achieved by means of dynamic programming. We refer the interested reader to Sections EC.1 and EC.2 of the online companion, where the formal proof of Theorem 4 is presented.

## 5. Computational Experiments

In this section, we implement our PTAS on a wide variety of randomly generated test instances with up to 125 products. Additionally, we distill the ideas and findings of our PTAS into two heuristic approaches for the assortment problem under the click-based MNL model.

### 5.1. Instance generator

We randomly generate cardinality-constrained instances of the click-based MNL assortment problem with $n \in \{20, 30, 40\}$. Specifically, we enforce that each recommended assortment must include exactly six products. The motivation for the addition of this constraint is two-fold. First, in many e-commerce settings, product recommendation pages are composed of displays that consist of a fixed number of products. This is the case, for example, in the Alibaba discount coupon setting considered in Section 6.1, where Alibaba is tasked with presenting personalized six product displays of discounted products to arriving customers. Second, by enforcing this exact cardinality constraint, we limit both the total number of feasible assortments to $\binom{n}{6}$, and the number of of potential consideration sets to $2^6$ for any fixed assortment. As such, when $n$ is reasonably small, i.e. $n \leq 40$, we can use a complete enumeration over all feasible assortments to recover the optimal assortment, which enables us to report an exact optimality gap as our measure of each recommended assortment's profitability. Additionally, we conduct a separate set of experiments with $n \in \{80, 100, 125\}$ to demonstrate that our PTAS scales to instances with large $n$. The prices, weights and consideration probabilities of each product are generated as follows:

- *Prices:* For each item $i \in [n]$, we generate $r_i$ from a log-normal distribution with location 0 and scale $\sigma \in \{0.1, 0.5\}$. Let $r_{\min}$ and $r_{\max}$ be the respective smallest and largest prices generated for a particular instance.
- *Preference weights:* The preference weight of item $i$ is set to $w_i = e^{\alpha_w - \beta_w r_i}$, where $\alpha_w = r_{\min} \cdot \beta_w$ and we vary $\beta_w \in \{0.5, 1, 2, 4\}$.
- *Consideration probabilities:* We first compute unnormalized consideration probabilities $\lambda'_i = e^{\alpha_c - \beta_c r_i}$ for items $i \in [n]$, and then set the true consideration probability of item $i$ to be $\lambda_i = 0.1 \cdot (\lambda'_i / \sum_{j\in[n]} \lambda'_j)$. We set $\alpha_c = \beta_c \cdot r_{\min}$ and we vary $\beta_c \in \{0.5, 1, 2, 4\}$.

We choose $\alpha_w$ and $\alpha_c$ as described above to ensure that the preference weights and click probabilities do not vary too widely. Furthermore, the scaling of the consideration probabilities is motivated by the idea that click-through rates are generally quite small in e-commerce settings. For example, in Section 6.3 we fit click-based MNL models to historical click/sales data from Alibaba and find that the predicted click probabilities never exceed 0.1.

This set-up allows us to characterize each test case through a particular configuration of the four parameters $(n, \sigma, \beta_w, \beta_c)$. Since we vary $n \in \{20, 30, 40\}$, $\sigma \in \{0.1, 0.5\}$, and $\beta_w, \beta_c \in \{0.5, 1, 2, 4\}$, there are a total of 96 unique test cases. For each test case, we generate 20 unique streams of prices, which ultimately yields 1,920 distinct assortment problem instances to consider. We solve each of these instances using the three approaches detailed next.

### 5.2. Tested algorithms

In this section, three algorithms are developed in order to compute assortment recommendations for the click-based MNL instances generated in Section 5.1. The first is a version of our PTAS that we mold to the cardinality constrained instances at-hand. It is important to note that this updated version of our approximation scheme continues to generates $\epsilon$-optimal assortments in a polynomial running time. In addition, we devise two heuristic approaches that draw inspiration from our PTAS to varying degrees, but come with no performance guarantees.

Each of the three algorithms begins with rounding the preference weights, revenues, and consideration probabilities, as described in Section 4. Consequently, the items are partitioned into the classes $\Lambda_p \cap \mathcal{W}_q$, where $p \in [P_{\min}, P_{\max}]$ and $q \in [Q_{\min}, Q_{\max}]$. In what follows, since we are choosing assortments consisting of six products, the choice probabilities can be computed exactly for every assortment by enumerating over all $2^6$ consideration sets.

*Coupon Display PTAS.* Our first approach is a slightly modified variant of our PTAS that is adapted to the cardinality constrained setting that we consider, the full details of which can be found in Appendix EC.3.1. We demonstrate that this implementation incurs running times that are less than a minute on average, even for the instances in which $n = 125$. These results demonstrate that our PTAS scales to problems of typical size as seen in related literature; see, for example, the recent numerical studies of Bertsimas and Mišić (2019), Berbeglia et al. (2018), and Désir et al. (2020). The PTAS was implemented in Java 1.8 on an Intel Core i5 with 3.2GHz CPU and 32GB of RAM.

*Greedy heuristic.* This approach is motivated by the revenue-ordered by class property highlighted in Section 3.2. Specifically, we have shown that there exists an optimal assortment $S^*$ that picks the $|S^* \cap \Lambda_p \cap \mathcal{W}_q|$ largest revenue products in $\Lambda_p \cap \mathcal{W}_q$, for every $p \in [P_{\min}, P_{\max}]$ and $q \in [Q_{\min}, Q_{\max}]$. We develop a greedy heuristic that constructs assortments satisfying this revenue-order-by-class property, by iteratively adding the product that generates the largest marginal increase in the expected revenue. More formally, we initially start with an empty assortment $S_0 = \emptyset$, and in each step, we add exactly one product to our assortment until the required capacity of 6 products is reached. Using $S_t$ to denote our current assortment, we construct the collection of candidate assortment $\mathcal{S}_{t+1} = \{S_t \cup \{i_{t,p,q}\} : p \in [P_{\min}, P_{\max}], q \in [Q_{\min}, Q_{\max}]\}$, where $i_{t,p,q}$ is the largest revenue product out of $(\Lambda_p \cap \mathcal{W}_q) \setminus S_t$. Consequently, we define $S_{t+1} = \arg\max_{S \in \mathcal{S}_{t+1}} \mathcal{R}(S)$. Ultimately, our heuristic returns the assortment $S_6$ hence defined. It is worth noting that, since Alibaba is required to display precisely 6 products, items are added up until reaching this capacity, even if the expected revenue decreases. The greedy heuristic was implemented in Python 3.6 on an Intel Core i5 with 3.2GHz CPU and 32GB of RAM.

*SAA heuristic.* This approach uses a popular simulation-optimization technique, known as sample average approximation (SAA); see, for example, the paper by Kleywegt et al. (2002). In this approach, for each product $i \in [n]$, we first generate $K$ samples $C^1_{[n]\setminus\{i\}}, \ldots, C^K_{[n]\setminus\{i\}}$ of the random consideration set $C_{[n]\setminus\{i\}}$, which are utilized as follows. For every assortment $S \subseteq [n]$, we use

$$\lambda_i r_i w_i \cdot \frac{1}{K} \sum_{k \in [K]} \frac{1}{1 + w_i + w(S \cap C^k_{[n]\setminus\{i\}})}$$

as an approximation of the expected revenue contribution of product $i \in S$. Hence, for a fixed collection of sampled consideration sets, our goal is to solve the following nonlinear integer program:

$$\max_{x \in \mathcal{F}} \sum_{i \in [n]} \lambda_i r_i w_i x_i \cdot \left( \frac{1}{1 + w_i + \sum_{j \in C^k_{[n]\setminus\{i\}}} w_j x_j} \right), \tag{3}$$

where $\mathcal{F} = \{x \in \{0,1\}^n : \sum_{i \in [n]} x_i = 6\}$ denotes all feasible six-product assortments. It is not difficult to reformulate the latter problem (3) as a linear integer program (SAA-IP), which is presented in Appendix EC.3.2. This heuristic was implemented in Python 3.6 on an Intel Core i5 with 3.2GHz CPU and 32GB of RAM; all integer programs were solved in Gurobi 6.5.1.

### 5.3. Results

In this section, we compare the performance of the three algorithms described above, by varying the accuracy level $\epsilon \in \{0.05, 0.04, 0.03\}$. For the remainder of this section, we use PTAS as a shorthand for the updated approximation scheme, GR for the greedy heuristic, and SAA for the SAA-based integer programming approach. Furthermore, for each algorithm $\mathcal{A} \in \{\text{PTAS}, \text{GR}, \text{SAA}\}$, let $S^{\mathcal{A}}$ be the assortment returned by algorithm $\mathcal{A}$, and let $S^*$ be the optimal assortment. We define the optimality gap of algorithm $\mathcal{A}$ for a single instance as $\frac{\mathcal{R}(S^*)-\mathcal{R}(S^{\mathcal{A}})}{\mathcal{R}(S^*)}$, and we drop from our analysis all "easy" instances deemed to be those for which all of the tested approaches recover the optimal assortment.

We begin by noting that we were unable to carry out SAA for all problem instances due to prolonged running times for solving (SAA-IP), which occurred even when this integer program was formulated with $K = 20$ samples. For example, when $\epsilon = 0.05$, there were test cases for which it took upwards of 25 minutes to solve (SAA-IP). Nevertheless, we carried out a smaller set of experiments to give some insight into the performance of SAA, where we varied $n \in \{20, 30\}$, $\sigma \in \{0.1, 0.5\}$, and $\beta_w, \beta_c \in \{0.5, 1, 2\}$, and then generated only 10 problem instances for each test case. Table 1 shows the average optimality gap and running time of SAA and GR. These results show that GR universally dominate SAA on the instances considered; producing significantly more profitable assortments in fractions of the running time.

| | | Avg. % Opt. Gap | | Avg. Running Time (s) | |
|---|---|---|---|---|---|
| $n$ | $\epsilon$ | GR | SAA | GR | SAA |
| 20 | 0.05 | 0.25 | 1.06 | 0.0037 | 31.08 |
| 20 | 0.1 | 0.70 | 1.11 | 0.0037 | 24.38 |
| 30 | 0.05 | 0.23 | 1.20 | 0.0059 | 193.79 |
| 30 | 0.1 | 0.76 | 2.98 | 0.0059 | 131.05 |

**Table 1 Results of smaller experiments, which show the relative performance of SAA**

Given the discussion above, we only analyze the performance of PTAS and GR on the full testbed of instances. The results for this full set of experiments are presented in Table 2, where columns 3 and 4 respectively show the average optimality gaps of PTAS and GR over all test instance. Furthermore, columns 5 and 6 show the largest optimality gaps of both approaches across all test cases. This statistic is meant to gives a sense of the worst-case performance of PTAS and GR. Finally, columns 6 and 7 present the average running times of PTAS and GR over all instances.

With only a superficial scan of the results in Table 2, it is clear that both PTAS and GR are quite efficacious. For example, when $\epsilon = 0.05$, both approaches have average optimality gaps below 0.25%. Furthermore, it is important to remember that the optimality gaps reported in Table 2

| | | Avg. % Opt. Gap | | Max. % Opt. Gap | | Avg. Running Time (s) | |
|---|---|---|---|---|---|---|---|
| $n$ | $\epsilon$ | PTAS | GR | PTAS | GR | PTAS | GR |
| 20 | 0.03 | 0.012 | 0.14 | 1.66 | 11.79 | 0.040 | 0.0038 |
| 20 | 0.04 | 0.026 | 0.16 | 1.53 | 11.79 | 0.029 | 0.0037 |
| 20 | 0.05 | 0.065 | 0.21 | 3.91 | 11.79 | 0.022 | 0.0037 |
| 30 | 0.03 | 0.015 | 0.13 | 1.29 | 6.41 | 0.39 | 0.0060 |
| 30 | 0.04 | 0.022 | 0.18 | 1.29 | 6.41 | 0.26 | 0.0060 |
| 30 | 0.05 | 0.063 | 0.21 | 4.27 | 6.41 | 0.17 | 0.0060 |
| 40 | 0.03 | 0.014 | 0.15 | 0.68 | 6.53 | 1.61 | 0.0084 |
| 40 | 0.04 | 0.030 | 0.20 | 1.22 | 6.53 | 1.00 | 0.0084 |
| 40 | 0.05 | 0.082 | 0.24 | 4.47 | 6.53 | 0.61 | 0.0084 |

**Table 2 Optimality gaps and running times of GR and PTAS**

| $n$ | Avg. Running Time (s) |
|---|---|
| 80 | 11.73 |
| 100 | 25.20 |
| 125 | 51.53 |

**Table 3 Performance of PTAS for instances with large $n$ and $\epsilon = 0.05$.**

are pessimistic measures of performance, since these metrics are computed after ignoring instances where both PTAS and GR recover the optimal assortment. Interestingly, we also observe that PTAS appears to be more robust than GR in terms of its worst case performance. More specifically, when $\epsilon = 0.03$, PTAS had a worst-case optimality gap of 1.66% compared to the 11.79% observed for GR. Hence all-in-all, it seems that PTAS provides the best balance between fast running times and high quality assortments. That said, GR is a suitable alternative for massive-scale e-commerce applications, where the computational requirements can be very restrictive. Finally, we note that the results seen in Table 3 indicate that PTAS can scale to instances with 125 products, finishing on average in less than a minute when $\epsilon = 0.05$.

## 6. Case Study with Alibaba Sales Data

In this section, we present a case study in which we fit click-based MNL models and mixed-MNL models to historical sales and click data from Alibaba, a Chinese online and mobile commerce company that has recently surpassed Walmart as the world's largest retailer (Lim 2016). Our primary goal is to understand and quantify the benefits of using click behavior as a proxy for the set of products considered by each customer. First, we demonstrate that utilizing the click-based MNL model leads to significant improvements over the mixed-MNL model in term of prediction accuracy. Moreover, the computation time needed to fit the click-based MNL is multiple orders of magnitude smaller than the time required to fit the mixed-MNL models. Next, we quantify the revenue impact of model misspecification, which occurs when the presumed choice model does not match the model that governs reality. Specifically, we consider the profitability of assortment decisions made under an MNL model, when a click-based MNL model is in fact the ground truth. The assortment instances we consider are generated using the MNL and click-based MNL models that were fit to the Alibaba sales data. All-in-all, our results indicate that click behavior provides a strong signal for the set of products ultimately considered by customers, and hence, a choice model that incorporates this behavioral premise will likely capture customer purchasing patterns more accurately.

*Outline.* We begin in Section 6.1 by presenting the exact retailing setting at Alibaba, as considered in our case study, which is followed by an inclusive summary of the available historical sales data in Section 6.2. Then, we describe in Sections 6.3 and 6.4 the methods employed to fit both the click-based MNL model and the mixed-MNL model respectively. The fitting accuracies of these two models are benchmarked against each other as well as a standard MNL model in Section 6.5. It is worth noting that we also conduct semi-synthetic experiments to study the cost of model misspecification, where the assortment recommendation of the MNL and click-based MNL models are compared in terms of revenue. To keep the paper concise, these results are presented in the online companion, Appendix EC.4.3.

### 6.1. General setting

The sales data presented in this section were collected at Tmall.com, one of Alibaba's largest online marketplaces that connects third-party sellers to customers. More specifically, Tmall.com is China's largest third-party business-to-consumer platform for branded goods, such as Nike and Adidas. Throughout this section, the term "seller" refers to any business that offers its products on the marketplace considered.

Customers arriving to the front page of Tmall.com can then navigate to seller-specific pages, which only offer products sold by that particular seller. Upon visiting a particular seller's front page, the customer can acquire a seller-specific coupon by clicking on a coupon icon at the top of that page. Clicking this icon brings the customer to a "coupon sub-page" that contains a personalized display of six products, each of which can now be purchased at a discount price. Alibaba selects the six products to display with the goal of maximizing the aggregate revenue generated by coupon-claiming customers. As formalized in Appendix EC.4.3, this decision is analogous to an assortment optimization problem. Figure 1 demonstrates how a customer progresses from a seller's front page to the coupon sub-page and finally to the six displayed products. For each customer who arrives to one of these coupon sub-pages, our data set includes the set of offered products as well as the subset of products clicked and ultimately purchased by that customer.


**Figure 1 The process of landing of the seller-specific coupon sub-page**

## 6.2. Available sales data

In what follows, we describe the sales data shared by Alibaba. We begin by discussing the make-up of the available historical sales data used to fit our customer choice models.

*Seller statistics.* Alibaba provided us with two weeks of sales data acquired from the top-10 sellers on Tmall.com in terms of overall traffic. Due to confidentiality agreements, the precise identities of these sellers cannot be revealed. Table 4 provides key sales statistics for each of these ten sellers, including the category of products sold by the particular seller in addition to information on the number of products clicked and purchased over the two week selling horizon. The rightmost column of this table specifies the conversion rate, which is the fraction of customers who arrived to a coupon sub-page with six displayed products and who made a purchase. Typically, in the context of e-commerce, this statistic is between two and five percent, which is precisely what we observe. It should be noted that, since multiple products were purchased in approximately 0.01% of customer visits, the conversion rate does not exactly evaluate to the number of purchases (column five) divided by the number of customer arrivals (column six).

*Available sales data.* For each seller, we let $\mathcal{N} = \{1, \ldots, n\}$ be the set of all potential products that could be displayed on coupon sub-pages (column three in Table 4). Further, we assume the sales data is composed of historical sales data from $\tau$ customers (column six). To help unclutter notation, we do not indicate the dependency of $\mathcal{N}$ and $\tau$ on the specific seller, even though these quantities clearly differ from one seller to another, as indicated in Table 4. For each arriving customer $t$, we denote by $S_t \subseteq \mathcal{N}$ the set of six displayed products, where each product $i \in S_t$ is stored as a vector

| Seller | Product Category | # products | # clicks | # purchases | # customers | conversion % |
|---|---|---|---|---|---|---|
| 1 | Electronics | 169 | 8,338 | 2,045 | 41,765 | 4.88 |
| 2 | Men's Apparel | 1,047 | 11,508 | 1,956 | 213,678 | 0.88 |
| 3 | Diapers | 132 | 10,296 | 2,979 | 90,467 | 3.01 |
| 4 | Bed Linens | 115 | 6,975 | 1,767 | 39,494 | 4.43 |
| 5 | Perfume | 103 | 32,535 | 8,478 | 131,822 | 6.16 |
| 6 | Women's Apparel | 501 | 7,267 | 2,127 | 63,466 | 3.23 |
| 7 | Furniture | 49 | 4,949 | 1,937 | 33,579 | 5.75 |
| 8 | Cooking Appliances | 82 | 4,220 | 1,448 | 40,108 | 3.59 |
| 9 | Women's Apparel | 118 | 17,792 | 2,163 | 139,853 | 1.49 |
| 10 | Cooking Appliances | 38 | 3,376 | 2,180 | 37,925 | 5.75 |

**Table 4 Key seller statistics**

of representative feature values $X_{it}$. More specifically, we use the 25 features with the highest importance scores according to the machine learning approaches utilized by Alibaba's engineers for demand estimation purposes. Among these top 25 features are product-specific features such as price, the image quality of the associated picture displayed to each customer, and various measures of historical click-through rates and customer sentiment towards the product as reflected by past reviews. In addition, we use customer-specific features such as the given customer's spending and total number of products added to the shopping cart both in the last week and in the last month. Beyond these rather straightforward product/customer features, we also have access to joint features that are specific to each customer and product pair. For example, one such joint feature is a collaborative filtering score, that takes into account past purchase and click behavior from the given customer and other customers who are deemed to have similar purchasing preferences to compute a single score signifying the extent to which the particular product will appeal towards the particular customer. Once again, due to confidentiality agreements, we cannot disclose the complete list of all 25 features.

The sales data for each customer $t$ can be described by a 4-tuple $(S_t, X_t, \text{Click}_t, z_t)$, where $X_t = \{X_{it} : i \in S_t\}$ specifies the feature values of each offered product, $\text{Click}_t = \{c_{it} \in \{0,1\} : i \in S_t\}$ is the subset of displayed products that were clicked, and $z_t$ indicates the purchased product. When customer $t$ did not make a purchase, we set $z_t = 0$; in the rare event that $k \geq 2$ different products were purchased during a single visit, we create $k$ copies of the 4-tuple $(S_t, X_t, \text{Click}_t, z_t)$, one for each of the unique values for $z_t$. With regards to $\text{Click}_t$, we set $c_{it} = 1$ if product $i \in S_t$ is clicked by customer $t$, and $c_{it} = 0$ otherwise. Within the scope of our click-based MNL model, for each customer $t$, we assume that $C_{S_t} = \{i \in S_t : c_{it} = 1\}$. In other words, we use the click events to indicate whether each product in the displayed assortment is part of the given customer's consideration set. For each seller, we represent the full set of historical sales as $\text{PurchaseHistory} = \{(S_t, X_t, C_t, z_t) : t = 1, \ldots, \tau\}$.

### 6.3. Fitting the click-based MNL model

In this section, we describe the methods used to estimate the parameters of the click-based MNL using maximum likelihood estimation (MLE). To begin, for each seller, we randomly select 75% of its sales data to be used for fitting the choice models and hold-out the remaining 25% of the data to test the accuracy of these models. Stated more concisely, we employ a 75/25 train/test split of the sales data. We then combine the training data for each seller and use this aggregated data set to seed the MLE problem. Once the optimal sets of parameters have been derived, we measure the accuracy of the fitted model using the log-likelihood computed on the testing sets of each of each individual seller. The series of steps described above – 75/25 train/test split, solving the MLE problems, computing the log-likelihood on the test data sets – make up what we refer to as a single trial. We ultimately report the average out-of-sample log-likelihood for each seller over 20 trials.

*Model parameterization.* Under the MNL model, the random utility that customer $t$ associates with product $i \in S_t$ is given by $U_{it} = V_{it} + \epsilon_{it}$, where $V_{it}$ is the known deterministic component of this utility and $\epsilon_{it}$ is an i.i.d. Gumbel random variable that captures heterogeneity in the customer population. In order to incorporate product and customer features, one can write the deterministic component as a linear combination of the feature values, meaning that $V_{it} = \beta^T X_{it}$, where $X_{it}$ is the vector of feature values that customer $t$ associated with product $i \in S_t$. In this case, the probability that customer $t$ purchases product $i \in S_t$ is given by

$$\pi\left(i, S_t, X_t\right) = \frac{e^{\beta^T X_{it}}}{1 + \sum_{j \in S_t} e^{\beta^T X_{jt}}} ,$$

whereas the no-purchase option is selected with probability

$$\pi\left(0, S_t, X_t\right) = \frac{1}{1 + \sum_{j \in S_t} e^{\beta^T X_{jt}}} .$$

Having fully described the second stage of our click-based MNL model, we now formalize the initial consideration set formulation step. Initially, we assume that the click probabilities do not depend on the feature sets of arriving customers. As a result, we let $\lambda_i$ be the probability $\Pr[c_{it} = 1]$ that any given customer $t$ clicks on product $i \in S_t$. We use $\lambda = \{\lambda_i : i \in \mathcal{N}\}$ to denote the set of all click probabilities.

With this notation, we formulate the log-likelihood, written as a function of PurchaseHistory, under the click-based MNL model as follows.

$$\begin{aligned}
&\mathcal{LL}\left(\beta, \lambda \middle| \text{PurchaseHistory}\right) \\
&\quad = \log\left(\prod_{t=1}^{\tau}\left(\prod_{i \in C_{S_t}} \lambda_i \prod_{i \in S_t \setminus C_{S_t}} (1 - \lambda_j)\right) \pi\left(z_t, C_{S_t}, X_t\right)\right)
\end{aligned}$$

$$= \sum_{t=1}^{\tau} \left( \sum_{i \in C_{S_t}} \log \lambda_i + \sum_{i \in S_t \setminus C_{S_t}} \log (1-\lambda_i) \right) + \sum_{t=1}^{\tau} \log \left( \pi \left( z_t, C_{S_t}, X_t \right) \right)$$

$$= \sum_{t=1}^{\tau} \left( \sum_{i \in C_{S_t}} \log \lambda_i + \sum_{i \in S_t \setminus C_{S_t}} \log (1-\lambda_i) \right) + \sum_{t=1}^{\tau} \left( \beta^T X_{z_t t} - \log \left( 1 + \sum_{i \in C_{S_t}} e^{\beta^T X_{it}} \right) \right). \tag{4}$$

Therefore, the MLE problem of interest, used to derive estimates of the feature weights $\beta$ and the click probabilities $\lambda$, is given by:

$$\max_{\lambda \in [0,1]^n, \beta} \mathcal{LL} \left( \beta, \lambda \middle| \, \text{PurchaseHistory} \right). \qquad \text{(MLE Click)}$$

Clearly, given the log-likelihood function in (4), the MLE problem decomposes into independent optimizations problems over the weights $\beta$ and click probabilities $\lambda$. Hence, (MLE Click) can be rewritten as:

$$\max_{\lambda \in [0,1]^n, \beta} \mathcal{LL} \left( \beta, \lambda \middle| \, \text{PurchaseHistory} \right) = \max_{\lambda \in [0,1]^n} \sum_{t=1}^{\tau} \left( \sum_{i \in C_{S_t}} \log \lambda_i + \sum_{i \in S_t \setminus C_{S_t}} \log (1-\lambda_i) \right)$$
$$+ \max_{\beta} \sum_{t=1}^{\tau} \left( \beta^T X_{z_t t} - \log \left( 1 + \sum_{i \in C_{S_t}} e^{\beta^T X_{it}} \right) \right).$$

*Solving the MLE problem.* Fortunately, the problem of finding the optimal feature weights $\beta^*$ is exactly an MLE problem under the standard MNL model, which is a well-known convex optimization problem (see Chapter 3 of Train (2009)). To numerically solve this problem, we use Tensorflow (Abadi et al. 2015), through which we implement a highly optimized version of stochastic gradient descent. Since the click probabilities follow a multinomial distribution, it is easy to verify that the optimal click probabilities are given by $\lambda_i^* = \sum_{t=1}^{\tau} c_{it} / \sum_{t=1}^{\tau} \mathbb{1}_{\ i \in S_t}$. In other words, each optimal click probabilities $\lambda_i^*$ is merely the empirical click probabilities, i.e., the proportion of customers who clicked product $i$ out of those to whom this product was offered.

While this closed-form expression for $\lambda^*$ is easy to compute, its accuracy in reflecting the true click probabilities will critically depend on the amount of available click data. Since the majority of the sellers we consider carry over 100 products, we are estimating over 100 click probabilities. In light of this characteristic of our setting, we find that this version of the click-based MNL model is significantly outperformed by the classic MNL model, a pattern that is most likely a consequence of insufficient click data. This insight leads us to develop a more sophisticated procedure for estimating the click probabilities, in which we featurize the click probabilities and use sophisticated machine learning methods to understand how these features shape click behaviors.

*Featurized click probabilities.* In order to avoid estimating a unique click probability for each product, we featurize these probabilities with the same set of 25 features that were used to featurize the deterministic utility components $V_{it}$ for the MNL choice model. More formally, we assume that the click probabilities take the form $\lambda_i = g(X_{it})$, where we fit a gradient boosted classification tree to determine the function $g(\cdot)$. Specifically, we use a Python implementation of catboost (Prokhorenkova et al. 2018); a new gradient boosting toolkit. Our exact implementation, which includes the hyperparameters that we tune, is provided in Appendix EC.4.1

### 6.4. Fitting the mixed-MNL model

The mixed-MNL models are also fit to the historical Alibaba sales data using an MLE-based approach. To ensure a fair comparison between the mixed-MNL and click-based MNL models, we utilize the same training and testing sets to fit and assess both model. In what follows, we first formalize the mixed-MNL model and its corresponding MLE problem. We then describe how we utilize a sequential fitting procedure to determine the optimal number of customer segments for each fitted mixed-MNL model.

*The mixed-MNL model.* Under the mixed-MNL model, the customer population is partitioned into $G$ customer segments, who each make purchasing decisions according to a unique MNL model. As such, in our Alibaba setting, the probability that customer $t$ purchases product $i \in S_t$ is given by

$$\pi\left(i, S_t, X_t\right) = \sum_{g\in[G]} \theta_g \cdot \frac{e^{\beta_g^T X_{it}}}{1+\sum_{j\in S_t} e^{\beta_g^T X_{jt}}},$$

whereas the no-purchase option is selected with probability

$$\pi\left(0, S_t, X_t\right) = \sum_{g\in[G]} \theta_g \cdot \frac{1}{1+\sum_{j\in S_t} e^{\beta_g^T X_{jt}}} \ .$$

In the above choice probability expressions, $\theta_g$ and $\beta_g$ respectively signify the arrival probability and vector of features weights associated with customer segment $g$. We use $\theta = \{\theta_g : g \in [G]\}$ to denote the set of all arrival probabilities. Finally, it is important to note that these choice probabilities also reflect the fact that the segment to which each arriving customer belongs is latent, since the sales data does not reveal this information.

Given the specification described above, the log-likelihood under the mixed-MNL model is as follows.

$$\begin{aligned}\mathcal{LL}\left(\beta_1, \ldots, \beta_G, \theta \middle| \, \text{PurchaseHistory}\right) &= \log\left(\prod_{t=1}^{\tau} \pi\left(i, S_t, X_t\right)\right) \\ &= \sum_{t=1}^{\tau} \log\left(\sum_{g\in[G]} \theta_g \cdot \frac{e^{\beta_g^T X_{z_t t}}}{1+\sum_{j\in S_t} e^{\beta_g^T X_{jt}}}\right).\end{aligned}$$

Therefore, the MLE problem of interest, used to derive estimates of the feature weights $\beta_1, \ldots, \beta_G$ and the click probabilities $\theta$, is given by:

$$\max_{\theta \in [0,1]^n, \beta_1, \ldots, \beta_G} \mathcal{LL}\left(\beta_1, \ldots, \beta_G, \theta \middle| \, \text{PurchaseHistory}\right), \tag{MLE MMNL}$$

with the addition constraint that $\sum_{g \in [G]} \theta_g = 1$.

*Solving the MLE problem and choosing G.* We solve problem (MLE MMNL) by directly maximizing the log-likelihood using MATLAB's constrained non-linear solver *fmincon* (MATLAB Optimization Toolbox). Although this log-likelihood function is non-convex, we found that *fmincon* generally converged to a local optima.

The optimal number of customer segments is determined by solving problem (MLE MMNL) in sequential stages, where in stage $k$, we set $G = k$. Furthermore, in stage $k$, we seed the MLE problem with the optimal estimates from stage $k-1$ so as to ensure that the training likelihood strictly increases as we add more customer segments. To determine when to stop this process, we employ a train/validate approach that is a common machine learning practice. More specifically, we split the training set into a smaller training set and a validation data set that will be used to assess the benefits of additional customer segments. The smaller training set that seeds the MLE problem now consists of 60% of the data, while the validation set consists of the remaining 15% that made-up the initial training set. The testing set remains unchanged. In each stage, after fitting the mixed-MNL as described above, we compute the fitted model's likelihood on the validation set. If this validation likelihood is an improvement over the previous stage's validation likelihood, then we increment $G$ by one, and move to the next stage. Otherwise, we stop the fitting procedure and output the fitted model with the highest validation likelihood. We set a time limit of three hours for fitting each mixed-MNL model, but allow any stage that has commenced to be carried out to completion before halting the procedure.

### 6.5. Fitting Accuracy Results

We begin by noting that the three hour time limit allocated to our procedure for fitting the mixed-MNL models was exhausted in every case, whereas fitting the click-based MNL model never required more than 20 minutes. In fact, each mixed-MNL model ended up being fit with $G = 3$ customer segments. To check if the three hour time limit had been too restrictive, we conducted an addition set of 20 trials, where we lifted the time cap, and fit only mixed-MNL models with $G \in \{1, 2, 3, 4, 5\}$. We present the results of these experiments in Appendix EC.4.2, where it can be seen that the mixed-MNL models fit with $G = 3$ was generally the most accurate model, and so the time limit had little to no effect on the performance of the mixed-MNL fits.

The results for our main set of experiments are presented in Table 5 and Figure 2. These results provide average measures of fitting accuracy based on 20 trials, where recall that a single trial consists of the following three steps: randomly splitting the data into 75/25 train and test sets, solving the corresponding MLE problem on train data, and computing the out-of-sample log-likelihood on test data. The first column of Table 5 identifies the seller number and the second gives the average log-likelihood on the testing set of the fitted click-based MNL. The third and fourth columns give the percentage improvements of the click-based MNL fits as compared to the fits of a standard MNL model and the mixed-MNL model.

| Seller # | Log Likelihood | Improvement over MNL | Improvement over mixed-MNL |
|---|---|---|---|
| 1 | -1791.18 | 13.21% | -0.13% |
| 2 | -2369.30 | 4.86% | 1.88% |
| 3 | -2920.93 | 3.11% | 0.38% |
| 4 | -1595.40 | 2.66% | 2.48% |
| 5 | -6633.32 | 1.92% | 0.87% |
| 6 | -2070.23 | 6.80% | 0.56% |
| 7 | -1667.22 | 3.25% | 0.067% |
| 8 | -1295.56 | 2.74% | -0.21% |
| 9 | -2468.47 | 0.99% | 0.94% |
| 10 | -1698.77 | 3.44% | 1.10% |

**Table 5 Predictive performance of the fitted models.**

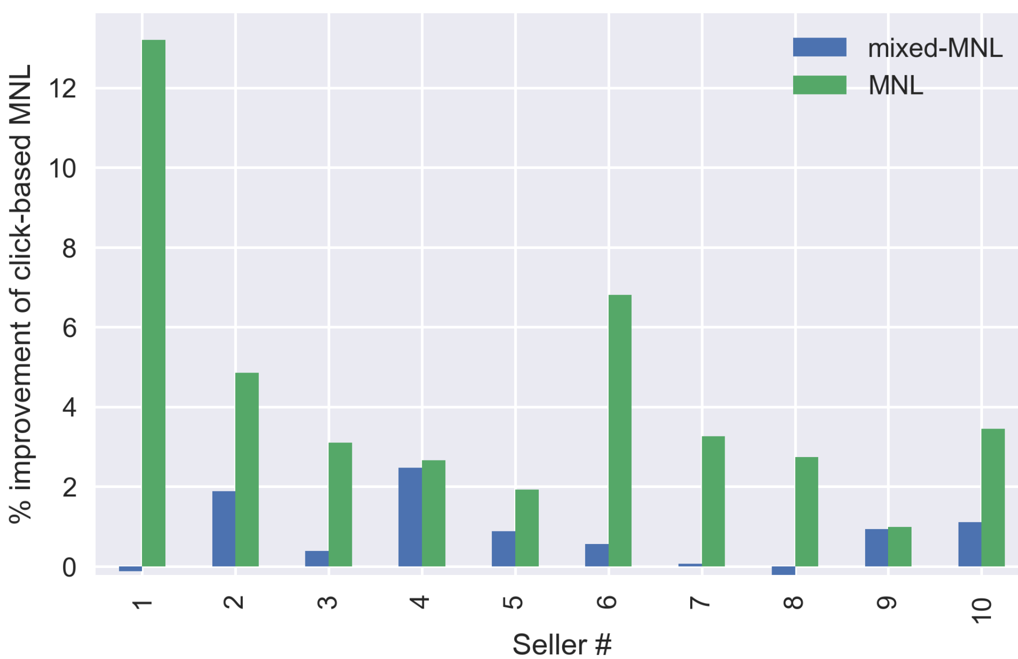


**Figure 2 Visualization of results from Table 5.**

The first salient feature of the results presented in Table 5 is the noticeable improvement of the click-based MNL over the traditional MNL model. In fact, for eight out of the ten sellers, this percentage improvement is at least 2.5%, and even tops 4.8% for Sellers 1,2 and 6. These gains appear especially impressive in light of other empirical studies that seek to benchmark recently developed choice models against the MNL model. For example, Aouad et al. (2018) and Şimşek and Topaloglu (2018) respectively benchmark the Exponomial and Markov chain choice models against the MNL model, but report gains in out-of-sample likelihood measures of only 1-2%. Second, and somewhat surprisingly, we observe that the fitting accuracy of the click-based MNL model exceeds that of the mixed-MNL model for eight of the ten sellers. Moreover, for six of the ten sellers, these gains exceed 0.5%. Altogether, these results suggest that clickstream data contains valuable information to refine the predictions of customer choices, and that the click-based MNL seems to incorporate this information in an effective and efficient manner.

## References

Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X (2015) TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.

Alon N, Spencer JH (2016) *The Probabilistic Method* (Wiley), 4th edition.

Aouad A, Farias V, Levi R (2020) Assortment optimization under consider-then-choose choice models. *Management Science* .

Aouad A, Feldman J, Segev D (2018) The exponomial choice model: Algorithmic frameworks for assortment optimization and data-driven estimation case studies. *Available at SSRN 3192068* .

Aouad A, Segev D (2020) Display optimization for vertically differentiated locations under multinomial logit preferences. *Management Science* .

Ban GY, Rudin C (2019) The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1):90–108.

Berbeglia G, Garassino A, Vulcano G (2018) A comparative empirical study of discrete choice models in retail operations. Working paper. Available online as SSRN report #3136816.

Bertsimas D, Mišić VV (2019) Exact first-choice product line optimization. *Operations Research* 67(3):651–670.

Blanchet J, Gallego G, Goyal V (2016) A markov chain approximation to choice modeling. *Operations Research* 64(4):886–905.

Boutsikas MV, Vaggelatou E (2002) On the distance between convex-ordered random variables, with applications. *Advances in Applied Probability* 34(2):349–374.

Chu LY, Nazerzadeh H, Zhang H (2020) Position ranking and auctions for online marketplaces. *Management Science* .

Cohen MC, Lobel I, Paes Leme R (2020) Feature-based dynamic pricing. *Management Science* .

Crompton JL, Ankomah PK (1993) Choice set propositions in destination decisions. *Annals of Tourism Research* 20(3):461–476.

Davis J, Gallego G, Topaloglu H (2013) Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper. Available online at: https://people.orie.cornell.edu/huseyin/publications/logit-const.pdf.

Davis J, Gallego G, Topaloglu H (2014) Assortment optimization under variants of the nested logit model. *Operations Research* 62(2):250–273.

Davis JM, Topaloglu H, Williamson DP (2015) Assortment optimization over time. *Operations Research Letters* 43(6):608–611.

Derakhshan M, Golrezaei N, Manshadi V, Mirrokni V (2018) Product ranking on online platforms. *Available at SSRN 3130378* .

Désir A, Goyal V, Jagabathula S, Segev D (2020) Mallows-smoothed distribution over rankings approach for modeling choice. *Operations Research (accepted)* .

Désir A, Goyal V, Zhang J (2014) Near-optimal algorithms for capacity constrained assortment optimization. Working paper. Available online as SSRN report #2543309.

Farias VF, Li AA (2019) Learning preferences with side information. *Management Science* 65(7):3131–3149.

Feldman J, Topaloglu H (2015a) Bounding optimal expected revenues for assortment optimization under mixtures of multinomial logits. *Production and Operations Management* 24(10):1598–1620.

Feldman J, Topaloglu H (2015b) Capacity constraints across nests in assortment optimization under the nested logit model. *Operations Research* 63(4):812–822.

Feldman J, Topaloglu H (2018) Capacitated assortment optimization under the multinomial logit model with nested consideration sets. *Operations Research* 66(2):380–391.

Feldman J, Zhang D, Liu X, Zhang N (2018) Customer choice models versus machine learning: Finding optimal product displays on alibaba. Working paper. Available online as SSRN report #3232059.

Gallego G, Li A (2017) Attention, consideration then selection choice model. Working paper. Available online as SSRN report #2926942.

Gallego G, Li A, Truong VA, Wang X (2020) Approximation algorithms for product framing and pricing. *Operations Research* 68(1):134–160.

Gallego G, Ratliff R, Shebalov S (2015) A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research* 63(1):212–232.

Gallego G, Topaloglu H (2014) Constrained assortment optimization for the nested logit model. *Management Science* 60(10):2583–2601.

Gilbride TJ, Allenby GM (2004) A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science* 23(3):391–406.

Hauser JR (1978) Testing the accuracy, usefulness, and significance of probabilistic choice models: An information-theoretic approach. *Operations Research* 26(3):406–421.

Hauser JR (2014) Consideration-set heuristics. *Journal of Business Research* 67(8):1688–1699.

Hauser JR, Wernerfelt B (1990) An evaluation cost model of consideration sets. *Journal of Consumer Research* 16(4):393–408.

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

Howard JA, Sheth JN (1969) *The Theory of Buyer Behavior* (New York, Wiley).

Huh WT, Levi R, Rusmevichientong P, Orlin JB (2011) Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research* 59(4):929–941.

Huh WT, Rusmevichientong P (2009) A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* 34(1):103–123.

Jagabathula S, Rusmevichientong P (2016) A nonparametric joint assortment and price choice model. *Management Science* 63(9):3128–3145.

Javanmard A, Nazerzadeh H (2019) Dynamic pricing in high-dimensions. *The Journal of Machine Learning Research* 20(1):315–363.

Jeuland AP (1979) Brand choice inertia as one aspect of the notion of brand loyalty. *Management Science* 25(7):671–682.

Karlin S, Novikoff A (1963) Generalized convex inequalities. *Pacific Journal of Mathematics* 13(4):1251–1279.

Karp RM (1972) Reducibility among combinatorial problems. *Complexity of Computer Computations*, 85–103 (Springer).

Kleywegt AJ, Shapiro A, Homem-de Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2):479–502.

Kunnumkal S, Topaloglu H (2008) Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research* 56(3):646–664.

Lapersonne E, Laurent G, Le Goff JJ (1995) Consideration sets of size one: An empirical investigation of automobile purchases. *International Journal of Research in Marketing* 49(1):55–66.

Laudon KC, Traver CG (2013) *E-commerce* (Pearson).

Le Cam L (1960) An approximation theorem for the poisson binomial distribution. *Pacific Journal of Mathematics* 10(4):1181–1197.

Lim J (2016) Alibaba group FY2016 revenue jumps 33%, EBITDA up 28%. Forbes, https://www.forbes.com/sites/jlim/2016/05/05/alibaba-fy2016-revenue-jumps-33-ebitda-up-28/#5f15a05f53b2.

Luce RD (1959) Individual choice behavior: A theoretical analysis. *Frontiers in Econometrics* 2:105–142.

Manzini P, Mariotti M (2014) Stochastic choice and consideration sets. *Econometrica* 82(3):1153–1176.

MATLAB Optimization Toolbox (2019) Matlab optimization toolbox. The MathWorks, Natick, MA, USA.

McFadden D (1974) Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics* 2:105–142.

Mehta N, Rajiv S, Srinivasan K (2003) Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science* 22(1):58–84.

Méndez-Díaz I, Miranda-Bront JJ, Vulcano G, Zabala P (2014) A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Applied Mathematics* 164:246–263.

Montgomery AL, Li S, Srinivasan K, Liechty JC (2004) Modeling online browsing and path analysis using clickstream data. *Marketing Science* 23(4):579–595.

Prokhorenkova L, Gusev G, Vorobev A, Dorogush AV, Gulin A (2018) Catboost: unbiased boosting with categorical features. *Advances in Neural Information Processing Systems*, 6639–6649.

Qiang S, Bayati M (2016) Dynamic pricing with demand covariates, working paper. Available online as SSRN report #2765257.

Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research* 58(6):1666–1680.

Rusmevichientong P, Shmoys D, Tong C, Topaloglu H (2014) Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23(11):2023–2039.

Shaked M, Shanthikumar JG (2007) *Stochastic orders* (Springer Science & Business Media).

Shocker AD, Ben-Akiva M, Boccara B, Nedungadi P (1991) Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions. *Marketing Letters* 2(3):181–197.

Silk AJ, Urban GL (1978) Pre-test-market evaluation of new packaged goods: A model and measurement methodology. *Journal of Marketing Research* 15(2):171–191.

Şimşek AS, Topaloglu H (2018) An expectation-maximization algorithm to estimate the parameters of the markov chain choice model. *Operations Research* 66(3):748–760.

Sumida M, Gallego G, Rusmevichientong P, Topaloglu H, Davis J (2020) Revenue-utility tradeoff in assortment optimization under the multinomial logit model with totally unimodular constraints. *Management Science* .

Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(24):15–33.

Train K (2009) *Discrete Choice Methods with Simulation* (Cambridge University Press).

Vazirani VV (2013) *Approximation algorithms* (Springer Science & Business Media).

Wang R, Sahin O (2017) The impact of consumer search cost on assortment planning and pricing. *Management Science* 64(8):3469–3970.

## Appendix A: The Connection between the Click-Based MNL and the Mixed-MNL Choice Models

The goal of this section is to formalize the connection between the click-based MNL and mixed-MNL models. More specifically, we first show how to represent any click-based MNL model in terms of an equivalent mixed-MNL model. We then argue that since the resulting mixed-MNL model requires specifying exponentially many customer segments, existing approaches for assortment optimization in this context are computational impractical, except possibly for very small scale instances.

*The reduction.* To begin, recall that under the mixed-MNL model, the customer population is partitioned into $K$ customer segments, where each segment $k \in [K]$ is associated with an arrival probability $\theta_k$ and a vector of MNL-based preference weights $w_1^k, \dots, w_n^k$. The arrival probability of each customer class can be interpreted as the fraction of customers belonging to this segment. For a given assortment $S \subseteq [n]$ and item $i \in S$ , the choice probability is

$$\pi(i,S) = \sum_{k\in[K]} \theta_k \cdot \frac{w_i^k}{1+w_k(S)},$$

where $w_k(S) = \sum_{i\in S} w_i^k$.

We can convert any click-based MNL model with consideration probabilities $\{\lambda_i\}_{i\in[n]}$ and MNL-based preference weights of $\{w_i\}_{i\in[n]}$ to an equivalent mixed-MNL model as follows. First, we introduce a customer segment for each subset of items $C \subseteq [n]$. The stochastic nature of the consideration set formation phase in the click-based MNL model is captured by defining the arrival probability of customer segment $C$ as

$$\theta_C = \prod_{i\in C} \lambda_i \cdot \prod_{i\notin C} (1-\lambda_i) .$$

The preference weights for customers of segment $C$ are defined as $w_i^C = w_i$ if $i \in C$, and $w_i^C = 0$ otherwise. Under this specification, note that the choice probabilities of item $i \in S$ is

$$\begin{aligned} \pi(i,S) &= \sum_{C\subseteq[n]:i\in C} \theta_C \cdot \frac{w_i^C}{1+w_C(S)} && (5)\\ &= \sum_{C\subseteq S:i\in C} \left( \prod_{j\in C} \lambda_j \cdot \prod_{j\notin S\setminus C} (1-\lambda_j) \right) \cdot \frac{w_i}{1+w(C)} && (6)\\ &= \lambda_i \cdot \sum_{C\subseteq S\setminus\{i\}} \Pr\left[C_S^{-i} = C\right] \cdot \frac{w_i}{1+w_i+w(C)} && (7)\\ &= \lambda_i \cdot \mathbb{E}\left[ \frac{w_i}{1+w_i+w(C_S^{-i})} \right]. && (8) \end{aligned}$$

The important observation is that the last expression is exactly the choice probabilities dictated by the click-based MNL model. Hence the click-based MNL model can be viewed as a parameterized mixed-MNL model in the sense that its $2n$ parameters fully specify a mixed-MNL model over $2^n$ customer segments.

Given that there are a handful of previously developed approaches to tackle the assortment problem under a mixed-MNL model, it is natural to wonder if any of these algorithms can be combined with our reduction to attain any non-trivial guarantees for the click-base MNL model. Unfortunately, there is little merit in this idea due to the fact that the click-based MNL model's representation as a mixed-MNL model requires creating $2^n$ customer classes. It is worth observing that even with polynomially many segments, the mixed-MNL-based assortment optimization problem is known to be $\Omega(n^{1-\epsilon})$-hard to approximate, as shown by Désir et al. (2014). Furthermore, the integer-programming-based approach presented by Méndez-Díaz et al. (2014) is also not tractable in our setting since the formulation also requires exponentially many decision variables.

## Appendix B: Proofs from Section 2

### B.1. Proof of Theorem 1

To prove the desired result, we present a reduction from set partition, which is one of Karp's 21 NP-complete problems (Karp 1972). Here, we are given a collection of positive integers $w_1, \dots, w_n$. The partition problem asks whether these numbers can be divided into two subsets, each summing to precisely $L = \frac{1}{2}\sum_{i=1}^n w_i$.

Given a partition instance of this form, we define an instance of the assortment feasibility problem as follows:

- There are $n+1$ products.
- For every $i \in [n]$, the revenue of product $i$ is given by $r_i = 2$. The revenue of product $n+1$ is $r_{n+1} = 7$.
- For every $i \in [n]$, the preference weight of product $i$ is $w_i$, while that of product $n+1$ is $4L$. In addition, the no-purchase option is associated with a preference weight of $L$.
- Each of the products $1, \dots, n$ has a consideration probability of 1, while product $n+1$ is considered with probability 1/2.
- Finally, the expected revenue threshold is $K = 3$.

In the remainder of this proof, we show that there exists a set of products to offer with an expected revenue of at least $K$ if and only if there exists a vector $x \in \{0,1\}^n$ for which $\sum_{i=1}^n w_i x_i = L$.

We first observe that, since product $n+1$ has the largest revenue among all products, it is not difficult to show that offering this product always increases the expected revenue of any set of products. Therefore, the residual question for the assortment feasibility problem is that of deciding whether one can pick a subset of the products $1, \dots, n$ to obtain, along with the preselected product $n+1$, an expected revenue of at least $K = 3$. For this purpose, we use the binary vector $x = (x_1, \dots, x_n)$ to indicate our product choices, meaning that $x_i = 1$ if product $i$ is offered and $x_i = 0$ otherwise. Letting $W(x) = \sum_{i=1}^n w_i x_i$ and recalling that product $n+1$ is preselected, the expected revenue of the assortment decisions represented by $x$ can be written as

$$\frac{1}{2} \cdot \frac{2W(x)}{L + W(x)} + \frac{1}{2} \cdot \frac{2W(x) + 28L}{L + W(x) + 4L} = \frac{W(x)}{L + W(x)} + \frac{W(x) + 14L}{5L + W(x)}.$$

Therefore, there exists a set of products that provides an expected revenue of at least $K = 3$ if and only if there exists a vector $x \in \{0,1\}^n$ for which

$$\frac{W(x)}{L + W(x)} + \frac{W(x) + 14L}{5L + W(x)} \geq 3.$$

The inequality above can equivalently be rewritten as $(W(x) - L)^2 \leq 0$. Therefore, an expected revenue of at least $K = 3$ can be obtained if and only if there exists a vector $x \in \{0,1\}^n$ for which $W(x) = L$, which is precisely what the partition problem is interested in.

### B.2. Proof of Theorem 2

*The probabilistic question.* In order to establish this result, we consider a slightly more general setting. Specifically, let $X_1, \ldots, X_n$ be a collection of $n$ independent Bernoulli random variables; we use $\lambda_i$ to denote the success probability of $X_i$. In addition, let $w_1, \ldots, w_n$ be non-negative weights, and consider the random variable $W = \sum_{i \in [n]} w_i X_i$. Given $\beta \geq w_{\min}$, we wish to approximately estimate the expectation $\mathbb{E}[\frac{1}{\beta+W}]$. It is easy to see that addressing the latter question would immediately provide the required estimate for $\pi(i,S)$. Indeed, due to representation (1) of the choice probabilities, where we have shown that $\pi(i,S) = \lambda_i w_i \cdot \mathbb{E}[\frac{1}{1+w_i+w(C_S^{-i})}]$, the connection is made by setting $\beta = 1 + w_i$ and $W = \sum_{j \in S \setminus \{i\}} w_j X_j$.

*The continuous dynamic program.* Now let us first introduce the following notation:

- For an integer $0 \leq t \leq n$, let $W_t = \sum_{i \in [t]} w_i X_i$.
- For $\alpha \in [\beta, \beta + \sum_{i=t+1}^n w_i]$, let $F(t,\alpha) = \mathbb{E}[\frac{1}{\alpha+W_t}]$.

Clearly, our objective is to estimate $F(n,\beta) = \mathbb{E}[\frac{1}{\beta+W_n}] = \mathbb{E}[\frac{1}{\beta+W}]$. With this notation, we argue that the value function $F$ can be expressed in a recursive way:

- Initialization: For $t = 0$ and $\alpha \in [\beta, \beta + \sum_{i=1}^n w_i]$, we have $F(0,\alpha) = \frac{1}{\alpha}$.
- General step: For $t = 1, \ldots, n$ and $\alpha \in [\beta, \beta + \sum_{i=t+1}^n w_i]$, we have

$$\begin{aligned}
F(t,\alpha) &= \mathbb{E}\left[\frac{1}{\alpha + W_t}\right] \\
&= \mathbb{E}\left[\frac{1}{\alpha + W_{t-1} + w_t X_t}\right] \\
&= \Pr[X_t = 1] \cdot \mathbb{E}\left[\left.\frac{1}{\alpha + W_{t-1} + w_t X_t}\right| X_t = 1\right] \\
&\quad + \Pr[X_t = 0] \cdot \mathbb{E}\left[\left.\frac{1}{\alpha + W_{t-1} + w_t X_t}\right| X_t = 0\right] \\
&= \lambda_t \cdot \mathbb{E}\left[\frac{1}{\alpha + w_t + W_{t-1}}\right] + (1-\lambda_t) \cdot \mathbb{E}\left[\frac{1}{\alpha + W_{t-1}}\right] \\
&= \lambda_t \cdot F(t-1, \alpha + w_t) + (1-\lambda_t) \cdot F(t-1,\alpha) \ . 
\end{aligned} \tag{9}$$

It is worth noting that the parameter $\alpha$ is continuous, meaning that the characterization above is still not an algorithmic result.

*Efficient discretization.* In order to discretize the continuous dynamic program $F$, we define an approximate version $\tilde{F}$. Here, the parameter $\alpha$ is restricted to the set $\mathcal{E}$, consisting of all numbers in $[\beta, \beta + \sum_{i=1}^n w_i]$ that can be written as $\beta \cdot (1+\delta)^k$ for some integer $k \geq 0$, where $\delta = \frac{\epsilon}{2(n+1)}$. It is easy to verify that $|\mathcal{E}| = O(\frac{1}{\delta} \cdot \log(n \frac{w_{\max}}{w_{\min}})) = O(\frac{n}{\epsilon} \cdot \log(n \frac{w_{\max}}{w_{\min}}))$, since $\beta \geq w_{\min}$.

To define our modified recursive equations, let $\lfloor \cdot \rfloor_{\mathcal{E}}$ be an operator that rounds its argument down to the nearest number in $\mathcal{E}$. We mention in passing that each occurrence of this operator in the upcoming analysis is indeed well-defined, due to applying it only to arguments in $[\beta, \beta + \sum_{i=1}^n w_i]$. With this notation, the approximate value function $\tilde{F}$ is defined as follows:

- Initialization: For $t = 0$ and $\alpha \in \mathcal{E}$, we set $\tilde{F}(0,\alpha) = \frac{1}{\alpha}$.
- General step: For $t = 1, \ldots, n$ and $\alpha \in \mathcal{E} \cap [\beta, \beta + \sum_{i=t+1}^n w_i]$, we set

$$\tilde{F}(t,\alpha) = \lambda_t \cdot \tilde{F}(t-1, \lfloor \alpha + w_t \rfloor_{\mathcal{E}}) + (1-\lambda_t) \cdot \tilde{F}(t-1,\alpha) \ .$$

In the next claim, we show that $\tilde{F}(n, \lfloor \beta \rfloor_{\mathcal{E}})$ over-estimates $F(n,\beta)$ by a factor of at most $1+\epsilon$.

LEMMA 5. $F(n,\beta) \le \tilde{F}(n, \lfloor \beta \rfloor_{\mathcal{E}}) \le (1+\epsilon) \cdot F(n,\beta)$.

*Proof* To establish the desired claim, we first show that for every $0 \le t \le n$ and $\alpha \in [\beta, \beta + \sum_{i=t+1}^{n} w_i]$,

$$F(t,\alpha) \le \tilde{F}(t, \lfloor \alpha \rfloor_{\mathcal{E}}) \le (1+\delta)^{t+1} \cdot F(t,\alpha) \ , \tag{10}$$

by induction on $t$. The base case is $t=0$, where $\tilde{F}(0, \lfloor \alpha \rfloor_{\mathcal{E}}) = \frac{1}{\lfloor \alpha \rfloor_{\mathcal{E}}}$ and $F(0,\alpha) = \frac{1}{\alpha}$, clearly indicating that $F(0,\alpha) \le \tilde{F}(0, \lfloor \alpha \rfloor_{\mathcal{E}}) \le (1+\delta) \cdot F(0,\alpha)$.

Now, for the general case of $t \ge 1$, note that

$$\begin{aligned}
\tilde{F}(t, \lfloor \alpha \rfloor_{\mathcal{E}}) &= \lambda_t \cdot \tilde{F}(t-1, \lfloor \lfloor \alpha \rfloor_{\mathcal{E}} + w_t \rfloor_{\mathcal{E}}) + (1-\lambda_t) \cdot \tilde{F}(t-1, \lfloor \alpha \rfloor_{\mathcal{E}}) \\
&\ge \lambda_t \cdot F(t-1, \lfloor \alpha \rfloor_{\mathcal{E}} + w_t) + (1-\lambda_t) \cdot F(t-1,\alpha) \\
&= \lambda_t \cdot \mathbb{E}\left[ \frac{1}{\lfloor \alpha \rfloor_{\mathcal{E}} + w_t + W_{t-1}} \right] + (1-\lambda_t) \cdot F(t-1,\alpha) \\
&\ge \lambda_t \cdot \mathbb{E}\left[ \frac{1}{\alpha + w_t + W_{t-1}} \right] + (1-\lambda_t) \cdot F(t-1,\alpha) \\
&= \lambda_t \cdot F(t-1, \alpha + w_t) + (1-\lambda_t) \cdot F(t-1,\alpha) \\
&= F(t,\alpha) \ ,
\end{aligned}$$

where the first inequality follows from the induction hypothesis, the second and third equalities hold by definition of $F(t-1,\cdot)$, and the last equality is precisely the recursive equation (9). Using similar arguments in the opposite direction, we have

$$\begin{aligned}
\tilde{F}(t, \lfloor \alpha \rfloor_{\mathcal{E}}) &= \lambda_t \cdot \tilde{F}(t-1, \lfloor \lfloor \alpha \rfloor_{\mathcal{E}} + w_t \rfloor_{\mathcal{E}}) + (1-\lambda_t) \cdot \tilde{F}(t-1, \lfloor \alpha \rfloor_{\mathcal{E}}) \\
&\le (1+\delta)^t \cdot (\lambda_t \cdot F(t-1, \lfloor \alpha \rfloor_{\mathcal{E}} + w_t) + (1-\lambda_t) \cdot F(t-1,\alpha)) \\
&= (1+\delta)^t \cdot \left( \lambda_t \cdot \mathbb{E}\left[ \frac{1}{\lfloor \alpha \rfloor_{\mathcal{E}} + w_t + W_{t-1}} \right] + (1-\lambda_t) \cdot F(t-1,\alpha) \right) \\
&\le (1+\delta)^t \cdot \left( \lambda_t \cdot (1+\delta) \cdot \mathbb{E}\left[ \frac{1}{\alpha + w_t + W_{t-1}} \right] + (1-\lambda_t) \cdot F(t-1,\alpha) \right) \\
&\le (1+\delta)^{t+1} \cdot (\lambda_t \cdot F(t-1, \alpha + w_t) + (1-\lambda_t) \cdot F(t-1,\alpha)) \\
&= (1+\delta)^{t+1} \cdot F(t,\alpha) \ .
\end{aligned}$$

Based on inequality (10), it follows that $F(n,\beta) \le \tilde{F}(n, \lfloor \beta \rfloor_{\mathcal{E}})$ regardless of the value of $\delta$. In addition, since we have previously chosen $\delta = \frac{\epsilon}{2(n+1)}$, an upper bound on our over-estimation error is obtained by observing that

$$\frac{\tilde{F}(n, \lfloor \beta \rfloor_{\mathcal{E}})}{F(n,\beta)} \le (1+\delta)^{n+1} \le e^{(n+1)\cdot\delta} = e^{\epsilon/2} \le 1+\epsilon \ ,$$

where the last inequality holds since $\epsilon \in (0,1)$. $\square$

To conclude our proof, it remains to upper bound the number of states needed in order to compute $\tilde{F}(n, \lfloor \beta \rfloor_{\mathcal{E}})$. The size of the state space is $O(n \cdot |\mathcal{E}|)$, and the approximate value function for each such state can be evaluated in $O(1)$ time. Consequently, the overall running time is $O(n \cdot |\mathcal{E}|) = O(\frac{n^2}{\epsilon} \cdot \log(n \frac{w_{\max}}{w_{\min}}))$, precisely as stated in Theorem 2. $\square$

### B.3. Derandomization

Given a decomposable assortment $B$, we show how to efficiently compute a deterministic assortment $S_d \subseteq [n]$ with an expected revenue of $\mathcal{R}(S_d) \geq \mathbb{E}_{S\sim B}[\mathcal{R}(S)]$. To this end, we apply the method of conditional expectation (Alon and Spencer 2016, Chap. 16) with respect to the sequence of random variables $(B_1,\dots,B_n)$. The idea is to iteratively replace each random variable by one of its possible realizations, without decreasing the conditional expected revenue in each iteration, given the set of deterministic values chosen thus far.

In order to apply this derandomization method, one should just be able to compute the expected revenue generated by a decomposable assortment. Indeed, for any binary sequence $(b_1,\dots,b_k) \in \{0,1\}^k$ with $k \in [0,n]$, we wish to compute the expected revenue $\mathbb{E}_{S\sim B^k}[\mathcal{R}(S)]$, where $B^k = (b_1,\dots,b_k,B_{k+1},\dots,B_n)$. In turn, we are left with the task of computing the choice probability of each item $i \in [n]$ under the decomposable assortment $B^k$, denoted by $\mathbb{E}_{S\sim B^k}[\pi(i,S)]$. The key observation is that our FPTAS for computing the choice probabilities with respect to deterministic assortments, described in Theorem 2, can be readily leveraged to estimate $\mathbb{E}_{S\sim B^k}[\pi(i,S)]$. Indeed, in Claim 2 below, we show that computing the choice probabilities of a decomposable assortment is equivalent to computing the choice probabilities of a deterministic assortment with modified consideration probabilities. Specifically, for every item $i \in [n]$, we define $\hat{\lambda}_i = \lambda_i \cdot \Pr[B_i^k = 1]$. Henceforth, we denote by $\hat{\pi}(i,S)$ the choice probability of item $i$ in the assortment $S$ with respect to the modified consideration probabilities $(\hat{\lambda}_1,\dots,\hat{\lambda}_n)$.

Claim 2. $\mathbb{E}_{S\sim B^k}[\pi(i,S)] = \hat{\pi}(i,[n])$.

*Proof.* The equivalence between the choice probabilities proceeds from the observation that:

$$\begin{aligned}
\Pr_{S\sim B^k}[i \in C_S] &= \lambda_i \cdot \Pr_{S\sim B^k}[i \in S] \\
&= \lambda_i \cdot \Pr\left[B_i^k = 1\right] \\
&= \hat{\lambda}_i \\
&= \Pr\left[i \in \hat{C}_{[n]}\right] ,
\end{aligned}$$

where $\hat{C}_{[n]}$ is the consideration set induced by the assortment $[n]$ under the modified consideration probabilities $\hat{\lambda}_i$. By remarking that $\mathbb{I}[1 \in C_S],\dots,\mathbb{I}[n \in C_S]$ and $\mathbb{I}[1 \in \hat{C}_{[n]}],\dots,\mathbb{I}[n \in \hat{C}_{[n]}]$ are sequences of independent Bernoulli random variables, we infer from the above equality that $w(C_S^{-i})$ and $w(\hat{C}_{[n]}^{-i})$ follow the same distribution when $S \sim B^k$, for every $i \in [n]$. Hence, based on representation (1) of the choice probabilities, we obtain that

$$\begin{aligned}
\mathbb{E}_{S\sim B^k}[\pi(i,S)] &= \Pr_{S\sim B^k}[i \in C_S] \cdot \mathbb{E}_{S\sim B^k}\left[\frac{w_i}{1 + w_i + w(C_S^{-i})}\right] \\
&= \Pr\left[i \in \hat{C}_{[n]}\right] \cdot \mathbb{E}\left[\frac{w_i}{1 + w_i + w(\hat{C}_{[n]}^{-i})}\right] \\
&= \hat{\pi}(i,[n]) .
\end{aligned}$$

□

## Appendix C: Proofs from Section 3

### C.1. Proof of Claim 1

The proof proceeds by defining the parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}})$ according to the properties stated in Claim 1. Next, we explain how these parameters can be determined through efficient enumeration.

*Defining the input parameters.* Let $\hat{\mathcal{R}}^*$ be an under-estimate of the optimal expected revenue $\mathcal{R}(S^*)$ such that $\hat{\mathcal{R}}^* \leq \mathcal{R}(S^*) \leq 2 \cdot \hat{\mathcal{R}}^*$. By noting that $\mathcal{R}(\{i^*\}) \leq \mathcal{R}(S^*) \leq n \cdot \mathcal{R}(\{i^*\})$, where $i^* \in [n]$ is the item that maximizes $\mathcal{R}(\{i\})$, we pick $\hat{\mathcal{R}}^*$ within the set $\{2^s \cdot \mathcal{R}(\{i^*\}) : 0 \leq s \leq \lceil \log_2 n \rceil\}$.

- *Revenue contribution vectors:* We approximately estimate the revenue contributions given by $\{c^*_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}$ and $\{c^*_q\}_{\mathcal{Q}}$, through the respective quantities $\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}]}$ and $\{\hat{c}_q\}_{\mathcal{Q}}$. These vectors are jointly defined as follows. Letting $\delta = \frac{\epsilon}{\mathcal{L}} \cdot \hat{\mathcal{R}}^*$, where $\mathcal{L} = (P_{\max}+1) \cdot |\mathcal{Q}|$, for every $p \in [P_{\max}]$ and $q \in \mathcal{Q}$ the quantity $\hat{c}_{p,q}$ is of the form $\hat{c}_{p,q} = n_{p,q} \cdot \delta$, where $n_{p,q} \in \mathbb{N}$ is the unique integer for which

$$\hat{c}_{p,q} = n_{p,q} \cdot \delta \leq c^*_{p,q} < (n_{p,q}+1) \cdot \delta = \hat{c}_{p,q} + \delta \ . \tag{11}$$

Similarly, for every $q \in \mathcal{Q}$, the approximate quantity $\hat{c}_q$ is of the form $\hat{c}_q = n_q \cdot \delta$, where $n_q \in \mathbb{N}$ is the unique integer for which

$$\hat{c}_q = n_q \cdot \delta \leq c^*_q < (n_q+1) \cdot \delta = \hat{c}_q + \delta \ . \tag{12}$$

- *Coefficient vector:* We approximately estimate the coefficients $\alpha(S^*, q) = \mathbb{E}[\frac{w^q}{1+w^q+w(C_{S^*})}]$ of the unlikely items $\mathcal{U}_q$ in the optimal assortment $S^*$ for every $q \in \mathcal{Q}$, through the quantities $\{\hat{\alpha}_q\}_{\mathcal{Q}}$. Letting $\beta = \frac{1}{1+(n+1)\cdot w^{Q_{\max}}}$, the estimated quantity $\hat{\alpha}_q$ is of the form $\hat{\alpha}_q = w^q \beta \cdot (1+\epsilon)^{r_q}$, where $r_q \in \mathbb{N}$ is the unique integer for which

$$(1-\epsilon)^2 \cdot \hat{\alpha}_q \leq \alpha(S^*, q) \leq (1-\epsilon) \cdot \hat{\alpha}_q \ . \tag{13}$$

Inequalities (11)-(13) immediately imply that the input parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}})$ satisfy Properties 1-3, as stated in Claim 1. Now, we argue these estimates can be determined by enumerating over polynomially many distinct vectors. To this end, we bound the cardinality of the set $\Omega$ formed by all possible input vectors of the form specified in (11)-(13).

*Efficient enumeration.* We first observe that $\hat{\mathcal{R}}^*$ is chosen within a set of $O(\log n)$ distinct values. In order to obtain an upper bound on the number of distinct contribution vectors, we observe that

$$\sum_{p=1}^{P_{\max}} \sum_{q=Q_{\min}}^{Q_{\max}} n_{p,q} + \sum_{q=Q_{\min}}^{Q_{\max}} n_q \leq \frac{2\mathcal{L}}{\epsilon} \ ,$$

where the inequality follows from (11) and (12), noting that $\sum_{p=1}^{P_{\max}} \sum_{q=Q_{\min}}^{Q_{\max}} c^*_{p,q} + \sum_{q=Q_{\min}}^{Q_{\max}} c^*_q = \mathcal{R}(S^*) \leq 2 \cdot \hat{\mathcal{R}}^*$. Consequently, the vector formed by $\{n_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}$ and $\{n_q\}_{\mathcal{Q}}$ can be viewed as a partition of some integer $\mu \in [0, \frac{2\mathcal{L}}{\epsilon}]$ into precisely $\mathcal{L} = (|\mathcal{Q}| \cdot (P_{\max}+1)$ integers. Therefore, the contributions vector can be recovered by enumerating over $O(2^{O(\frac{\mathcal{L}}{\epsilon})}) = O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$ distinct partitions, since $\mathcal{L} = O(\frac{1}{\epsilon^3}\log\frac{1}{\epsilon}\log n)$.

In order to derive an upper bound on the number of distinct coefficient vectors, we remark that the quantity $\frac{\alpha(S^*,q)}{w^q} = \mathbb{E}[\frac{1}{1+w^q+w(C_{S^*})}]$ is monotone non-increasing in $q \in \mathcal{Q}$, and it is contained in the range $[\beta, \frac{1}{1+w^{Q_{\min}}}]$, where $\beta = \frac{1}{1+(n+1)w^{Q_{\max}}}$. Thus, by inequality (13), the corresponding sequence of exponents $\{r_q\}_{q=Q_{\min}}^{Q_{\max}}$ within our estimates $\hat{\alpha}_q = w^q\beta \cdot (1+\epsilon)^{r_q}$ is monotone non-increasing, and contained in the interval

$[\lceil \frac{1}{\epsilon} \cdot (\log(n+1) + |\mathcal{Q}|) \rceil]$. Each such sequence is uniquely determined by a partition of an integer $\mu \in [\lceil \frac{1}{\epsilon} \cdot (\log(n+1) + |\mathcal{Q}|) \rceil]$ into precisely $|\mathcal{Q}|$ integers, meaning that there are only $O(2^{O(\frac{1}{\epsilon}\log n + \frac{1}{\epsilon^3}\log n)}) = O(n^{O(\frac{1}{\epsilon^3})})$ distinct candidates since $|Q| = O(\frac{1}{\epsilon^2}\log n)$ by Assumption 2.

All in all, we have shown that parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}})$ that satisfy Properties 1-3 in Claim 1 can be recovered by enumerating over a collection $\Omega$ of cardinality $|\Omega| = O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$.

### C.2. Proof of Lemma 1 (Poissonization)

For $\mu \in [0,1]$, we use $B(\mu)$ to denote a Bernoulli random variable with probability of success $\mu$. We remind the reader that two random variables $Z_1$ and $Z_2$ are in the (standard) convex order $Z_1 \succeq_{\text{cx}} Z_2$ if, for every convex function $\phi$, we have $\mathbb{E}[\phi(Z_1)] \geq \mathbb{E}[\phi(Z_2)]$, provided the expectations exist. We use $Z_1 \sim Z_2$ to indicate that $Z_1, Z_2$ follow the same distribution. It is worth observing that the convex order $Z_1 \succeq_{\text{cx}} Z_2$ implies in particular that the convex non-increasing relationship $Z_1 \succeq_{\text{cni}} Z_2$ holds.

Our proof proceeds from two basic claims relating $B(\mu)$ and $P(\mu)$ through stochastic orders. Claim 3 argues that these two random variables are convexly ordered.

Claim 3. *$P(\mu) \succeq_{\text{cx}} B(\mu)$, for every $\mu \in [0,1]$*

This inequality is well-known; see for example Lemma 3 in the paper by Boutsikas and Vaggelatou (2002) who invoke a general criterion due to Karlin and Novikoff (1963). Since the convex order is closed under convolution (e.g., see Thm. 3.A.12(d) in Shaked and Shanthikumar (2007)), Claim 3 implies that

$$P(\lambda) \sim \sum_{i=1}^{k} P(\lambda_i) \succeq_{\text{cx}} \sum_{i=1}^{k} B(\lambda_i) \sim \sum_{i=1}^{k} X_i \sim Y \ .$$

Since the stochastic convex order implies the convex non-increasing order, we conclude that $P(\lambda) \succeq_{\text{cni}} Y$.

It remains to show that $Y \succeq_{\text{cni}} P((1+\epsilon)\cdot\lambda)$. To this end, when the Poisson arrival rate is made slightly larger than $\lambda$, we establish in the next claim a reverse inequality to that of Claim 3, according to the usual stochastic order.

Claim 4. *For every $\epsilon \in [0, \sqrt{2}-1]$ and $\mu \in [0,\epsilon]$, we have $P((1+\epsilon)\cdot\mu) \succeq_{\text{st}} B(\mu)$.*

*Proof* To prove the claim, it suffices to show that $\Pr[P((1+\epsilon)\cdot\mu) = 0] \leq \Pr[B(\mu) = 0]$. For this purpose, note that

$$\begin{aligned}
\Pr[P((1+\epsilon)\cdot\mu) = 0] &= e^{-(1+\epsilon)\cdot\mu} \\
&\leq 1 - (1+\epsilon)\cdot\mu + \frac{1}{2}(1+\epsilon)^2\cdot\mu^2 \\
&\leq 1 - \mu \\
&= \Pr[B(\mu) = 0] \ ,
\end{aligned}$$

where the second inequality holds since $\epsilon \in [0, \sqrt{2}-1]$ and $\mu \in [0,\epsilon]$. $\square$

By noting that the usual stochastic order implies a convex non-increasing relationship in reverse order, we have

$$Y \sim \sum_{i=1}^{k} X_i \sim \sum_{i=1}^{k} B(\lambda_i) \succeq_{\text{cni}} \sum_{i=1}^{k} P((1+\epsilon)\cdot\lambda_i) \sim P((1+\epsilon)\cdot\lambda) \ ,$$

where the stochastic inequality follows from Claim 4 and the fact that the usual stochastic order is closed under convolution.

### C.3. Proof of Lemma 2

Focusing on an unlikely item $i \in S \cap \mathcal{U}_q$, and letting $p \in [P_{\min}, 0]$ be the unique index for which $i \in S \cap \mathcal{W}_q \cap \Lambda_p$, we first observe that

$$\begin{aligned}\pi(i,S) &= \lambda_i \cdot \mathbb{E}\left[\frac{w_i}{1+w_i+w(C_S^{-i})}\right] \\ &\geq \lambda_i \cdot \mathbb{E}\left[\frac{w^q}{1+w^q+w(C_S)}\right] \\ &= \lambda_i \cdot \alpha(S,q) \ ,\end{aligned}$$

where the inequality above holds since $w_i = w^q$ and $w(C_S) \succeq_{\rm st} w(C_S^{-i})$. In the opposite direction, we have

$$\begin{aligned}\lambda_i \cdot \alpha(S,q) &= \lambda_i \cdot \mathbb{E}\left[\frac{w^q}{1+w^q+w(C_S)}\right] \\ &\geq \lambda_i \cdot \Pr\left[i \notin C_S\right] \cdot \mathbb{E}\left[\left.\frac{w^q}{1+w^q+w(C_S^{-i})}\right| i \notin C_S\right] \\ &\geq \lambda_i \cdot (1-\epsilon) \cdot \mathbb{E}\left[\frac{w_i}{1+w_i+w(C_S^{-i})}\right] \\ &= (1-\epsilon) \cdot \pi(i,S) \ ,\end{aligned}$$

where the second inequality proceeds by noting that the random variables $\mathbb{I}[i \notin C_S]$ and $w(C_S^{-i})$ are independent, and that $\Pr\left[i \notin C_S\right] = 1 - \lambda_i \geq 1 - \epsilon$ since item $i$ is unlikely.

### C.4. Proof of Lemma 3

We assume without loss of generality that items are indexed such that $r_1 \geq \cdots \geq r_n$, and let $S^* \subseteq [n]$ be an assortment for which the quantity $\sum_{i \in S^*} i$ is minimized over all optimal assortments. We argue that $S^*$ is necessarily revenue-ordered by class. To arrive at a contradiction, suppose that there exists a class $\mathcal{W}_q \cap \Lambda_p$ and a pair of items $i_1, i_2 \in \mathcal{W}_q \cap \Lambda_p$ such that $i_1 < i_2$, with $i_2 \in S^*$ and $i_1 \notin S^*$. In this case, we construct the altered assortment $\tilde{S} = (S^* \setminus \{i_2\}) \cup \{i_1\}$ where $i_1$ is swapped into $S^*$ in place of $i_2$. The important observation is that, since these items have precisely the same consideration probabilities and preference weights, which govern the random choice outcomes, we have $\pi(i_1, \tilde{S}) = \pi(i_2, S^*)$ as well as $\pi(j, \tilde{S}) = \pi(j, S^*)$ for every other item $j \in S^* \setminus \{i_2\}$. Thus, we obtain $\mathcal{R}(\tilde{S}) - \mathcal{R}(S^*) = r_{i_1} \cdot \pi(i_1, \tilde{S}) - r_{i_2} \cdot \pi(i_2, S^*) \geq 0$, since $r_{i_1} \geq r_{i_2}$ and $\pi(i_1, \tilde{S}) = \pi(i_2, S^*)$, meaning that the assortment $\tilde{S}$ is also optimal. However, as $i_1 < i_2$, we have $\sum_{i \in \tilde{S}} i < \sum_{i \in S^*} i$, contradicting the choice of $S^*$.

## Appendix D: Proof of Theorem 3: Analysis

In what follows, we formally establish near-optimal performance guarantees for the approximation algorithm presented in Section 3. We first establish in Appendix D.1 various structural properties that mirror the technical insights of Section 3. Next, in Appendices D.2 and D.3, our analysis proceeds by bounding the revenue contributions of the unlikely and likely items selected by the specialized algorithms `MinKnapsack`$(\cdot)$ and `Greedy`$(\cdot)$, respectively. Finally, by combining these results in Appendix D.4, we complete the proof of Theorem 3.

*Preliminaries and notation.* Recall from Section 3.4 that $B = B_{\text{unlike}} \cup B_{\text{likely}}$ is the decomposable assortment returned in Step 3 of our approximation scheme. In what follows, we fix a collection of input parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}) \in \Omega$ that satisfy Properties 1-3 of Claim 1. We denote by $\hat{B} = \hat{B}_{\text{unlike}} \cup \hat{B}_{\text{likely}}$ the resulting decomposable assortment, generated in Steps 2 and 3 of our approximation scheme, i.e., $\hat{B}_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$ and $\hat{B}_{\text{likely}} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}, \hat{B}_{\text{unlike}})$. We let $\hat{U}_q$ be the collection of items generated by solving the min-knapsack instance over the class of items $\mathcal{U}_q$, and let $\hat{B}_q$ be the resulting decomposable assortment specified by $\texttt{MinKnapsack}(\cdot)$ within this class. Similarly, we let $\hat{S}_t \subseteq \mathcal{N}_{\text{likely}}$ be the subset of likely items constructed in the $t$-th iteration of $\texttt{Greedy}(\cdot)$.

### D.1. Main structural properties

We remind the reader that the decomposable assortment $\hat{B}$ is comprised of $\hat{B}_{\text{unlike}}$ and $\hat{B}_{\text{likely}}$, constructed by $\texttt{MinKnapsack}(\cdot)$ and $\texttt{Greedy}(\cdot)$ in Sections 3.2 and 3.3, respectively. In what follows, we establish key structural properties of these decomposable assortments, which validate the design principles that guide these specialized algorithms.

*Properties of the min-knapsack procedure (Step 2a).* Recall that, in our description of $\texttt{MinKnapsack}(\cdot)$ in Section 3.2, the objective of the min-knapsack instances was to construct a subset of items $S$ so as to minimize the quantity $\mathbb{E}[|C_{S\cap\mathcal{U}_q}|]$ within each class of items $\mathcal{U}_q$. This formulation was (informally) related to a mitigation of the "cannibalization effects" due to the inclusion of items in the assortment. The next claim essentially confirms that, for every $q \in \mathcal{Q}$, the choice probabilities are no more "cannibalized" by our assortment decisions $\hat{B}_q$ within $\mathcal{U}_q$ than by the optimal assortment decisions $S^* \cap \mathcal{U}_q$.

Claim 5. *For every positive random variable $W > 0$ independent of the consideration sets $\{C_{S\cap\mathcal{U}_q}\}_{q\in\mathcal{Q}}$ and $\{C_{S^*\cap\mathcal{U}_q}\}_{q\in\mathcal{Q}}$, where $S$ denotes a random assortment sampled according to the decomposable assortment $\hat{B}_{\text{unlike}}$, we have*

$$\mathbb{E}_{S\sim\hat{B}_{\text{unlike}}}\left[\frac{1}{W + \sum_{q\in\mathcal{Q}} w^q \cdot |C_{S\cap\mathcal{U}_q}|}\right] \geq \mathbb{E}\left[\frac{1}{W + \sum_{q\in\mathcal{Q}} w^q \cdot |C_{S^*\cap\mathcal{U}_q}|}\right] .$$

The proof of Claim 5 is presented in Appendix D.5, and it crucially relies on the Poissonization idea of Lemma 1.

*Properties of the greedy procedure (Step 2b).* In what follows, we establish an invariant of $\texttt{Greedy}(\cdot)$, arguing that at any point in time, the constructed assortment forms a subset of the likely items picked by the optimal assortment $S^*$.

Claim 6. $\hat{S}_t \subseteq S^* \cap \mathcal{N}_{\text{likely}}$, *for every $t \in [0, T]$.*

The proof is presented in Appendix D.6. At a high-level, we show that the termination criterion placed on $\texttt{Greedy}(\cdot)$ prevents our algorithm from over-selecting items in any given class $\mathcal{W}_q \cap \Lambda_p$.

### D.2. Bounding the contributions of unlikely items

Armed with the properties of Appendix D.1, we are ready analyze the contributions of the unlikely items picked by $\hat{B}_{\text{unlike}}$ to the expected revenue $\mathbb{E}_{S\sim\hat{B}}[\mathcal{R}(S)]$.

Claim 7. $\mathbb{E}_{S\sim\hat{B}}[\sum_{i\in S\cap\mathcal{N}_{\text{unlike}}} r_i \cdot \pi(i, S)] \geq (1 - 4\epsilon) \cdot \sum_{q=Q_{\min}}^{Q_{\max}} \hat{c}_q$.

*Proof.* Observe that, for every $q \in \mathcal{Q}$, the expected choice probability $\mathbb{E}_{S\sim\hat{B}}[\pi(i,S)]$ of every unlikely item $i \in \hat{U}_q$ can be lower-bounded as follows:

$$\begin{aligned}
&\mathbb{E}_{S\sim\hat{B}}\left[\pi(i,S)\right] \\
&\quad = \mathrm{Pr}_{S\sim\hat{B}}\left[i \in C_S\right] \cdot \mathbb{E}_{S\sim\hat{B}}\left[\frac{w_i}{1+w_i+w(C_S^{-i})}\right] \\
&\quad \geq (1-\epsilon)\cdot\lambda_i\cdot\mathbb{E}_{S\sim\hat{B}}\left[\frac{w_i}{1+w_i+w(C_S)}\right] && (14)\\
&\quad = (1-\epsilon)\cdot\lambda_i\cdot\mathbb{E}_{S\sim\hat{B}}\left[\frac{w_i}{1+w_i+\sum_{q'=Q_{\min}}^{Q_{\max}} w^{q'}\cdot|C_{S\cap\mathcal{U}_{q'}}|+w(C_{S\cap\mathcal{N}_{\text{likely}}})}\right] \\
&\quad \geq (1-\epsilon)\cdot\lambda_i w_i\cdot\mathbb{E}_{S\sim\hat{B}}\left[\frac{1}{1+w_i+\sum_{q'=Q_{\min}}^{Q_{\max}} w^{q'}\cdot|C_{S\cap\mathcal{U}_{q'}}|+w(C_{S^*\cap\mathcal{N}_{\text{likely}}})}\right] && (15)\\
&\quad \geq (1-\epsilon)\cdot\lambda_i w_i\cdot\mathbb{E}\left[\frac{1}{1+w_i+\sum_{q'\in\mathcal{Q}} w^{q'}\cdot|C_{S^*\cap\mathcal{U}_{q'}}|+w(C_{S^*\cap\mathcal{N}_{\text{likely}}})}\right] && (16)\\
&\quad = (1-\epsilon)\cdot\lambda_i\cdot\mathbb{E}\left[\frac{w^q}{1+w^q+w(C_{S^*})}\right] \\
&\quad = (1-\epsilon)\cdot\lambda_i\cdot\alpha\left(S^*,q\right) \\
&\quad \geq (1-3\epsilon)\cdot\lambda_i\cdot\hat{\alpha}_q\ . && (17)
\end{aligned}$$

Here, inequality (14) is obtained by noting that $\mathrm{Pr}_{S\sim\hat{B}}\left[i\in C_S\right]=(1-\epsilon)\cdot\lambda_i$ and that $w(C_S^{-i}) \preceq_{\text{st}} w(C_S)$ for any subset of items $S$. Inequality (15) holds since $w(C_{S^*\cap\mathcal{N}_{\text{likely}}}) \succeq_{\text{st}} w(C_{S_T})$ as implied by Claim 6, and by noting that $S\cap\mathcal{N}_{\text{likely}}$ is equal to $S_T$ with probability 1, since $S$ is sampled from the decomposable assortment $\hat{B}$ whose likely items are offered with probability 1. Inequality (16) follows from Claim 5, instantiated with $W=1+w_i+w(C_{S^*\cap\mathcal{N}_{\text{likely}}})$. Finally, inequality (17) proceeds from Property 3 of Claim 1, stating in particular that $\alpha(S^*,q)\geq(1-\epsilon)^2\cdot\hat{\alpha}_q$.

Consequently, for the expected revenue contribution due to likely items, we obtain:

$$\begin{aligned}
\mathbb{E}_{S\sim\hat{B}}\left[\sum_{i\in S\cap\mathcal{N}_{\text{unlike}}} r_i\cdot\pi(i,S)\right] &\geq (1-3\epsilon)\cdot\sum_{q=Q_{\min}}^{Q_{\max}}\sum_{i\in U_q} r_i\lambda_i\cdot\hat{\alpha}_q \\
&\geq (1-3\epsilon)\cdot\sum_{q=Q_{\min}}^{Q_{\max}}\left(\frac{\epsilon}{n}\cdot\hat{c}_q\cdot\sum_{i\in\hat{U}_q}\gamma_i\right) \\
&\geq (1-4\epsilon)\cdot\sum_{q=Q_{\min}}^{Q_{\max}}\hat{c}_q\ ,
\end{aligned}$$

where the second inequality proceeds from eliminating floors in our definition of $\gamma_i$, where $\gamma_i=\lfloor\frac{n\lambda_i r_i\hat{\alpha}_q}{\epsilon\hat{c}_q}\rfloor$, and the last inequality holds since $\sum_{i\in S^*\cap\mathcal{U}_q}\gamma_i\geq(1-\epsilon)\cdot\frac{n}{\epsilon}$, as shown by inequality (24) in the proof of Claim 5. $\square$

### D.3. Bounding the contributions of likely items

We now turn our attention to the contributions of the likely items picked by $\hat{B}_{\text{likely}}$ to the expected revenue $\mathbb{E}_{S\sim\hat{B}}[\mathcal{R}(S)]$. Specifically, we establish the following claim.

Claim 8. $\mathbb{E}_{S\sim\hat{B}}[\sum_{i\in S\cap\mathcal{N}_{\text{likely}}} r_i\cdot\pi(i,S)]\geq(1-\epsilon)\cdot\sum_{p=1}^{P_{\max}}\sum_{q=Q_{\min}}^{Q_{\max}}\hat{c}_{p,q}$.

*Proof.* We begin by noting that $S \cap \mathcal{N}_{\text{likely}}$ is equal to $\hat{S}_T$ with probability 1, since the decomposable assortment $\hat{B}_{\text{likely}}$ generated in the last iteration of `Greedy`$(\cdot)$ includes each item of $S_T$ with probability 1. Thus, in order to prove the desired claim, it suffices to show that

$$\sum_{i \in \hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \tilde{\pi}(i, S \cup \hat{S}_T) \right] \geq (1-\epsilon) \cdot \hat{c}_{p,q} \ ,$$

for every $p \in [P_{\max}]$ and $q \in \mathcal{Q}$. For this purpose, by Claim 6, we have in particular $(\hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q) \subseteq (S^* \cap \Lambda_p \cap \mathcal{W}_q)$. Therefore, given the termination criterion of `Greedy`$(\cdot)$, the inequalities $|\hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q| < |S^* \cap \Lambda_p \cap \mathcal{W}_q|$ and $\sum_{i \in \hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim B_{\text{unlike}}}[\tilde{\pi}(i, S \cup \hat{S}_T)] < (1-\epsilon) \cdot \hat{c}_{p,q}$ cannot simultaneously hold at the last iteration $T$. If the latter inequality does not hold, we are clearly done, noting that this inequality is exactly the opposite of what we want to show. We therefore consider the case where the former inequality does not hold, i.e., $|\hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q| \geq |S^* \cap \Lambda_p \cap \mathcal{W}_q|$. Clearly, in this case, we have $(\hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q) = (S^* \cap \Lambda_p \cap \mathcal{W}_q)$. Consequently, letting $S^*_{\text{likely}} = S^* \cap \mathcal{N}_{\text{likely}}$, it follows that

$$\begin{aligned}
\sum_{i \in \hat{S}_T \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \tilde{\pi}\left(i, S \cup \hat{S}_T\right) \right] &= \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \tilde{\pi}\left(i, S \cup \hat{S}_T\right) \right] \\
&\geq \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \pi\left(i, S \cup \hat{S}_T\right) \right] \\
&\geq \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \pi\left(i, S \cup S^*_{\text{likely}}\right) \right] \ , \qquad (18)
\end{aligned}$$

where the first inequality holds since the FPTAS of Theorem 2 ensures that $\mathbb{E}_{S \sim \hat{B}_{\text{unlike}}}[\tilde{\pi}(i, S \cup \hat{S}_T)] \geq \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}}[\pi(i, S \cup \hat{S}_T)]$, and the second inequality holds since $\hat{S}_T \subseteq S^*_{\text{likely}}$ by Claim 6. In the remainder of this proof, we show that the sum-expression on the right-hand side of inequality (18) can be lower-bounded by $\hat{c}_{p,q}$. To this end, we observe that

$$\begin{aligned}
&\sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \pi\left(i, S \cup S^*_{\text{likely}}\right) \right] \\
&\quad = \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \lambda_i w_i \cdot \mathbb{E}_{S \sim \hat{B}_{\text{unlike}}} \left[ \frac{1}{1 + w_i + w(C_S) + w(C^{-i}_{S^* \cap \mathcal{N}_{\text{likely}}})} \right] \\
&\quad \geq \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \lambda_i w_i \cdot \mathbb{E} \left[ \frac{1}{1 + w_i + w(C_{S^* \cap \mathcal{N}_{\text{unlike}}}) + w(C^{-i}_{S^* \cap \mathcal{N}_{\text{likely}}})} \right] \\
&\quad = \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) \\
&\quad = c^*_{p,q} \\
&\quad \geq \hat{c}_{p,q} \ .
\end{aligned}$$

The first inequality proceeds from Claim 5, instantiated with $W = 1 + w_i + w(C^{-i}_{S^* \cap \mathcal{N}_{\text{likely}}})$. The last inequality follows from Property 1 of Claim 1. $\square$

### D.4. Proving the performance guarantee of Theorem 3

To conclude the proof of Theorem 3, we show that the decomposable assortment $B$ generates an expected revenue within factor $1 - O(\epsilon)$ of the optimal expected revenue $\mathcal{R}(S^*)$. Specifically, by combining the respective bounds obtained for unlikely and likely items in Claims 7 and 8, we obtain

$$\mathbb{E}_{S \sim B} \left[ \mathcal{R}(S) \right]$$

$$
\begin{aligned}
&\geq (1-\epsilon)\cdot \mathbb{E}_{S\sim\hat{B}}\left[\mathcal{R}\left(S\right)\right]\\
&= (1-\epsilon)\cdot \mathbb{E}_{S\sim\hat{B}}\left[\sum_{i\in S\cap\mathcal{N}_{\mathrm{likely}}} r_i\cdot\pi\left(i,S\right)+\sum_{i\in S\cap\mathcal{N}_{\mathrm{unlike}}} r_i\cdot\pi\left(i,S\right)\right]\\
&\geq (1-5\epsilon)\cdot\left(\sum_{p=1}^{P_{\max}}\sum_{q=Q_{\min}}^{Q_{\max}}\hat{c}_{p,q}+\sum_{q=Q_{\min}}^{Q_{\max}}\hat{c}_q\right)\\
&\geq (1-5\epsilon)\cdot\left(\sum_{p=1}^{P_{\max}}\sum_{q=Q_{\min}}^{Q_{\max}}\left(c^*_{p,q}-\frac{2\epsilon}{\mathcal{L}}\cdot\mathcal{R}(S^*)\right)+\sum_{q=Q_{\min}}^{Q_{\max}}\left(c^*_q-\frac{2\epsilon}{\mathcal{L}}\cdot\mathcal{R}(S^*)\right)\right)\\
&\geq (1-7\epsilon)\cdot\mathcal{R}(S^*)\ . \qquad (19)
\end{aligned}
$$

Here, the first inequality follows by noting that the decomposable assortment $B$ dominates $\hat{B}$ in terms of revenue by Step 3 of our approximation scheme. The second inequality follows from Claims 7 and 8. The next inequality is due to Properties 1 and 2 of Claim 1. The last inequality follows by recalling that $\mathcal{L}=(P_{\max}+1)\cdot|\mathcal{Q}|$.

### D.5. Proof of Claim 5

Let $\mathcal{Q}=\{q_1,\ldots,q_{|\mathcal{Q}|}\}$ be an arbitrary indexing of the elements of $\mathcal{Q}$. The proof proceeds by an induction over $t\in[0,|\mathcal{Q}|]$ showing that:

$$
\begin{aligned}
&\mathbb{E}_{S\sim\hat{B}_{\mathrm{unlike}}}\left[\frac{1}{W+\sum_{q\in\mathcal{Q}}w^q\cdot|C_{S\cap\mathcal{U}_q}|}\right]\\
&\qquad\geq \mathbb{E}_{S\sim\hat{B}_{\mathrm{unlike}}}\left[\frac{1}{W+\sum_{k=1}^{t}w^{q_k}\cdot|C_{S^*\cap\mathcal{U}_{q_k}}|+\sum_{k=t+1}^{|\mathcal{Q}|}w^{q_k}\cdot|C_{S\cap\mathcal{U}_{q_k}}|}\right]\ .
\end{aligned}
$$

The base case of the induction $t=0$ is immediate, since this inequality is satisfied with an equality. Next, we show that the induction hypothesis propagates to every $1\leq t\leq|\mathcal{Q}|$. To this end, letting $V_{-t}=W+\sum_{k=1}^{t-1}w^{q_k}\cdot|C_{S^*\cap\mathcal{U}_{q_k}}|+\sum_{k=t+1}^{t}w^q\cdot|C_{S\cap\mathcal{U}_{q_k}}|$, we have:

$$
\begin{aligned}
&\mathbb{E}_{S\sim\hat{B}_{\mathrm{likely}}}\left[\frac{1}{W+\sum_{k=1}^{t}w^{q_k}\cdot|C_{S^*\cap\mathcal{U}_{q_k}}|+\sum_{k=t+1}^{|\mathcal{Q}|}w^{q_k}\cdot|C_{S\cap\mathcal{U}_{q_k}}|}\right]\\
&\quad=\mathbb{E}\left[\frac{1}{w^{q_t}\cdot|C_{S^*\cap\mathcal{U}_{q_t}}|+V_{-t}}\right]\\
&\quad=\mathbb{E}_{V_{-t}}\left[\mathbb{E}\left[\frac{1}{w^{q_t}\cdot|C_{S^*\cap\mathcal{U}_{q_t}}|+V_{-t}}\middle|V_{-t}\right]\right]\\
&\quad\leq\mathbb{E}_{V_{-t}}\left[\mathbb{E}\left[\frac{1}{w^{q_t}\cdot P(\sum_{i\in S^*\cap\mathcal{U}_{q_t}}\lambda_i)+V_{-t}}\middle|V_{-t}\right]\right] \qquad (20)\\
&\quad\leq\mathbb{E}_{V_{-t}}\left[\mathbb{E}\left[\frac{1}{w^{q_t}\cdot P(\sum_{i\in\hat{U}_{q_t}}\lambda_i)+V_{-t}}\middle|V_{-t}\right]\right] \qquad (21)\\
&\quad\leq\mathbb{E}_{V_{-t}}\left[\mathbb{E}_{S\sim\hat{B}_{\mathrm{likely}}}\left[\frac{1}{w^{q_t}\cdot|C_{S\cap\mathcal{U}_{q_t}}|+V_{-t}}\middle|V_{-t}\right]\right] \qquad (22)\\
&\quad=\mathbb{E}_{S\sim\hat{B}_{\mathrm{likely}}}\left[\frac{1}{w^{q_t}\cdot|C_{S\cap\mathcal{U}_{q_t}}|+V_{-t}}\right]\\
&\quad=\mathbb{E}_{S\sim\hat{B}_{\mathrm{likely}}}\left[\frac{1}{W+\sum_{k=1}^{t-1}w^{q_k}\cdot|C_{S^*\cap\mathcal{U}_{q_k}}|+\sum_{k=t}^{|\mathcal{Q}|}w^{q_k}\cdot|C_{S\cap\mathcal{U}_{q_k}}|}\right]
\end{aligned}
$$

$$\leq \mathbb{E}_{S\sim \hat{B}_{\text{likely}}}\left[\frac{1}{W+\sum_{q\in\mathcal{Q}} w^q \cdot |C_{S\cap \mathcal{U}_q}|}\right] , \tag{23}$$

where inequality (23) immediately follows from our induction hypothesis.

Inequality (20) proceeds from Lemma 1 and the independence between the random variables $V_{-t}$ and $|C_{S^*\cap\mathcal{U}_{q_t}}|$. To apply Lemma 1, observe that $|C_{S^*\cap\mathcal{U}_{q_t}}| = \sum_{i\in S^*\cap\mathcal{U}_{q_t}} X_i$ where $\{X_i\}_{i\in S^*\cap\mathcal{U}_{q_t}}$ is a collection of independent Bernoulli random variables with probabilities of success $\{\lambda_i\}_{i\in S^*\cap\mathcal{U}_{q_t}}$. Hence, the first inequality of Lemma 1 implies that $P(\sum_{i\in S^*\cap\mathcal{U}_{q_t}}\lambda_i) \succeq_{\text{cni}} |C_{S^*\cap\mathcal{U}_{q_t}}|$. This relationship yields (20) since the function $x\mapsto \frac{1}{w^{q_t}\cdot x+v}$ is convex non-increasing over $\mathbb{R}^+$ for any $v>0$.

Inequality (21) is justified by showing that $\sum_{i\in S^*\cap\mathcal{U}_{q_t}}\lambda_i \geq \sum_{i\in\hat{U}_{q_t}}\lambda_i$. To this end, it suffices to show that the subset $S^*\cap\mathcal{U}_{q_t}$ is feasible with respect to the min-knapsack instance for which $\hat{U}_q$ is optimal (see Section 3.2). On the account that Property 3 in Claim 1 is satisfied, we have

$$\sum_{i\in S^*\cap\mathcal{U}_q} r_i\lambda_i\hat{\alpha}_q \geq \frac{1}{1-\epsilon}\cdot \sum_{i\in S^*\cap\mathcal{U}_q} r_i\lambda_i\alpha(S^*,q) \geq \sum_{i\in S^*\cap\mathcal{U}_q} r_i\pi(i,S^*) = c_q^* \geq \hat{c}_q \ ,$$

where the second inequality proceeds from Lemma 2, and the third inequality follows from Property 2 of Claim 1. Now, since $\gamma_i = \lfloor \frac{n\lambda_i r_i\hat{\alpha}_q}{\epsilon\hat{c}_q}\rfloor$, we get

$$\sum_{i\in S^*\cap\mathcal{U}_q}\gamma_i \geq \sum_{i\in S^*\cap\mathcal{U}_q}\left(\frac{n\lambda_i r_i\hat{\alpha}_q}{\epsilon\hat{c}_q}-1\right) \geq \frac{n}{\epsilon} - |S^*\cap\mathcal{U}_q| \geq (1-\epsilon)\cdot\frac{n}{\epsilon} \ . \tag{24}$$

Inequality (24) precisely states that $S^*\cap\mathcal{U}_{q_t}$ is feasible with respect to the min-knapsack instance associated with the class of items $\mathcal{U}_{q_t}$, thereby completing the proof of (21).

Inequality (22) proceeds from Lemma 1 and the independence between the random variables $V_{-t}$ and $|C_{S\cap\mathcal{U}_{q_t}}|$, where $S$ designates a random assortment sampled according to the decomposable assortment $\hat{B}_{\text{unlike}}$. Indeed, observe that $|C_{S\cap\mathcal{U}_{q_t}}| = \sum_{i\in\hat{U}_{q_t}}\tilde{X}_i$ where $\{\tilde{X}_i\}_{i\in\hat{U}_{q_t}}$ is a collection of independent Bernoulli random variables with probabilities of success $\tilde{\lambda}_i = (1-\epsilon)\cdot\lambda_i$ for every $i\in\hat{U}_{q_t}$. (Recall from Section 3.2 that $\hat{U}_q$ is the set of items that are picked with positive probability by the decomposable assortment $\hat{B}_{\text{unlike}}$ within $\mathcal{U}_q$.) The second inequality of Lemma 1 implies that $|C_{S\cap\mathcal{U}_{q_t}}| \succeq_{\text{cni}} P((1+\epsilon)\cdot\sum_{i\in\hat{U}_{q_t}}(1-\epsilon)\cdot\lambda_i) \succeq_{\text{cni}} P(\sum_{i\in\hat{U}_{q_t}}\lambda_i)$.

### D.6. Proof of Claim 6

This result is proven by induction over $t\in[0,T]$. The base case $t=0$ is clearly satisfied since $\hat{S}_0=\emptyset$ by definition. For the general case $t\geq 1$, we have $\hat{S}_{t-1}\subseteq S^*\cap\mathcal{N}_{\text{likely}}$ by the induction hypothesis. Recall that the set of items $\hat{S}_t$ was constructed by adding item $i_{t-1}$ to $\hat{S}_{t-1}$ in iteration $t-1$; we let $p\in[P_{\max}]$ and $q\in\mathcal{Q}$ be the unique indices for which $i_{t-1}\in\Lambda_p\cap\mathcal{W}_q$. Now, suppose that $i_{t-1}\notin S^*$. By the revenue-ordered-by-class property stated in Lemma 3, the subset $S^*\cap\Lambda_p\cap\mathcal{W}_q$ is formed by the $k_{p,q}$ highest revenue items in $\Lambda_p\cap\mathcal{W}_q$, where $k_{p,q}=|S^*\cap\Lambda_p\cap\mathcal{W}_q|$. We first observe that $|\hat{S}_{t-1}\cap\Lambda_p\cap\mathcal{W}_q| = k_{p,q}$. Otherwise, if $|\hat{S}_{t-1}\cap\Lambda_p\cap\mathcal{W}_q| < k_{p,q}$, our greedy selection rule ensures that item $i_{t-1}$ is necessarily among the $k_{p,q}$ highest revenue items in $\Lambda_p\cap\mathcal{W}_q$, implying that $i_{t-1}\in S^*\cap\Lambda_p\cap\mathcal{W}_q$, in contradiction to having $i_{t-1}\notin S^*$. Therefore, combining this observation with the induction hypothesis, we infer that $(\hat{S}_{t-1}\cap\Lambda_p\cap\mathcal{W}_q) = (S^*\cap\Lambda_p\cap\mathcal{W}_q)$. Now, just before picking item $i_{t-1}$ at iteration $t-1$ of the greedy algorithm, we had

$$\sum_{i\in\hat{S}_{t-1}\cap\Lambda_p\cap\mathcal{W}_q} r_i\cdot\mathbb{E}_{S\sim\hat{B}_{\text{unlike}}}\left[\pi\left(i,S\cup\hat{S}_{t-1}\right)\right]$$

$$
\begin{aligned}
&\leq \sum_{i \in \hat{S}_{t-1} \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}} \left[ \tilde{\pi}(i, S \cup \hat{S}_{t-1}) \right] \\
&< (1-\epsilon) \cdot \hat{c}_{p,q} \\
&< \hat{c}_{p,q} \ ,
\end{aligned} \tag{25}
$$

where the first inequality holds since the FPTAS of Theorem 2 ensures that $\mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}}[\pi(i, S \cup \hat{S}_{t-1})] \leq \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}}[\tilde{\pi}(i, S \cup \hat{S}_{t-1})]$. The second inequality follows from the termination criterion of our greedy procedure, since inequality (2) is necessarily violated.

On the other hand,

$$
\begin{aligned}
&\sum_{i \in \hat{S}_{t-1} \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}} \left[ \pi(i, S \cup \hat{S}_{t-1}) \right] \\
&\qquad = \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}} \left[ \pi(i, S \cup \hat{S}_{t-1}) \right] \\
&\qquad = \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \lambda_i w_i \cdot \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}} \left[ \frac{1}{1 + w_i + w(C_S) + w(C^{-i}_{\hat{S}_{t-1}})} \right] \\
&\qquad \geq \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \lambda_i w_i \cdot \mathbb{E}_{S \sim \hat{B}_{\mathrm{unlike}}} \left[ \frac{1}{1 + w_i + w(C_S) + w(C^{-i}_{S^* \cap \mathcal{N}_{\mathrm{likely}}})} \right] \\
&\qquad \geq \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \lambda_i w_i \cdot \mathbb{E} \left[ \frac{1}{1 + w_i + w(C_{S^* \cap \mathcal{N}_{\mathrm{unlike}}}) + w(C^{-i}_{S^* \cap \mathcal{N}_{\mathrm{likely}}})} \right] \\
&\qquad = \sum_{i \in S^* \cap \Lambda_p \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) \\
&\qquad = c^*_{p,q} \ .
\end{aligned}
$$

Here, the first equality holds since $(\hat{S}_{t-1} \cap \Lambda_p \cap \mathcal{W}_q) = (S^* \cap \Lambda_p \cap \mathcal{W}_q)$, as shown above. The first inequality holds since $\hat{S}_{t-1} \subseteq S^* \cap \mathcal{N}_{\mathrm{likely}}$ by the induction hypothesis, implying that $w(C^{-i}_{S^* \cap \mathcal{N}_{\mathrm{likely}}}) \succeq_{\mathrm{st}} w(C^{-i}_{\hat{S}_{t-1}})$. In addition, we note that $C_S$ is independent of $C^{-i}_{\hat{S}_{t-1}}$ and $C^{-i}_{S^* \cap \mathcal{N}_{\mathrm{likely}}}$, where $S$ is the random realization of the decomposable assortment $\hat{B}_{\mathrm{unlike}}$. The second inequality proceeds from Claim 5, instantiated with $W = 1 + w_i + w(C^{-i}_{S^* \cap \mathcal{N}_{\mathrm{likely}}})$. Combining the latter inequality with (25), we obtain $c^*_{p,q} < \hat{c}_{p,q}$, which contradicts Property 1 of Claim 1. It follows that $i_{t-1} \in S^* \cap \Lambda_p \cap \mathcal{W}_q$, and therefore $\hat{S}_t \subseteq S^* \cap \mathcal{N}_{\mathrm{likely}}$, as desired.

# Online Companion

## Appendix EC.1: Proof of Lemma 4

Our proof proceeds in two steps, where we separately analyze the effects of two rounding procedures. First, we use the rounded weights $\{\tilde{w}_i\}_{i\in[n]}$ and revenues $\{\tilde{r}_i\}_{i\in[n]}$, but we keep the original consideration probabilities $\{\lambda_i\}_{i\in[n]}$. We use $\mathcal{R}_{\tilde{w}}(\cdot)$ to denote the resulting expected revenue function. The following claim is established in Appendix EC.1.1.

Claim EC.1. *For every $S \subseteq [n]$, we have $(1-\epsilon)\cdot\mathcal{R}(S) \leq \mathcal{R}_{\tilde{w}}(S) \leq (1+2\epsilon)\cdot\mathcal{R}(S)$.*

Next, suppose that in addition to using the rounded weights $\{\tilde{w}_i\}_{i\in[n]}$ and revenues $\{\tilde{r}_i\}_{i\in[n]}$, we also consider the rounded consideration probabilities $\{\tilde{\lambda}_i\}_{i\in[n]}$ in place of $\{\lambda_i\}_{i\in[n]}$. Recall that $\tilde{\mathcal{R}}(\cdot)$ denotes the resulting expected revenue function. We fix an arbitrary assortment $S \subseteq [n]$. Our analysis proceeds from the next two claims.

Claim EC.2. *$\tilde{\mathcal{R}}(S) \geq (1-\epsilon)\cdot\mathcal{R}_{\tilde{w}}(S)$.*

Claim EC.3. *$\mathbb{E}_{S'\sim B^S}[\mathcal{R}_{\tilde{w}}(S')] \geq (1-\epsilon)\cdot\tilde{\mathcal{R}}(S)$.*

To conclude the proof of Lemma 4, note that

$$\mathbb{E}_{S'\sim B^S}[\mathcal{R}(S')] \geq (1-2\epsilon)\cdot\mathbb{E}_{S'\sim B^S}[\mathcal{R}_{\tilde{w}}(S'))] \geq (1-3\epsilon)\cdot\tilde{\mathcal{R}}(S) \ ,$$

where the former inequality is due to Claim EC.1 and the latter follows from Claim EC.3. Reciprocally, we have

$$\tilde{\mathcal{R}}(S) \geq (1-\epsilon)\cdot\mathcal{R}_{\tilde{w}}(S) \geq (1-2\epsilon)\cdot\mathcal{R}(S) \ ,$$

where the inequalities are derived by applying Claims EC.2 and EC.1 in succession.

### EC.1.1. Proof of Claim EC.1

Let $\pi_w(i,S)$ be the choice probability of item $i$ out of the assortment $S$, with respect to the original weight parameters $w$. The analogous choice probability $\pi_{\tilde{w}}(i,S)$ with respect to $\tilde{w}$ is defined similarly, noting that both quantities are independent of the price parameters. In order to establish the desired claim, it suffices to prove that, for every item $i \in S$,

$$(1-\epsilon)\cdot r_i\cdot\pi_w(i,S) \leq \tilde{r}_i\cdot\pi_{\tilde{w}}(i,S) \leq (1+2\epsilon)\cdot r_i\cdot\pi_w(i,S) \ .$$

To obtain the first inequality, rather than exploiting representation (1) of the choice probabilities, we express the latter by conditioning on how the consideration set $C_S$ is realized. Based on this idea, we have

$$\tilde{r}_i\cdot\pi_{\tilde{w}}(i,S) = \tilde{r}_i\cdot\sum_{T\subseteq S:i\in T}\Pr\left[C_S = T\right]\cdot\frac{\tilde{w}_i}{1+\tilde{w}(T)}$$

$$
\begin{aligned}
&\geq r_i \cdot \sum_{T \subseteq S : i \in T} \Pr\left[C_S = T\right] \cdot \frac{w_i}{1+\epsilon+(1+\epsilon)\cdot w(T)} \\
&\geq (1-\epsilon) \cdot r_i \cdot \sum_{T \subseteq S : i \in T} \Pr\left[C_S = T\right] \cdot \frac{w_i}{1+w(T)} \\
&= (1-\epsilon) \cdot r_i \cdot \pi_w(i,S) \ ,
\end{aligned}
$$

where the first inequality holds since $\tilde{r}_i \tilde{w}_i = r_i w_i$ and since

$$
\tilde{w}(T) = \sum_{i \in T \cap \mathcal{W}_0} \tilde{w}_i + \sum_{q \geq 1} \sum_{i \in T \cap \mathcal{W}_q} \tilde{w}_i \leq \frac{|T \cap \mathcal{W}_0|}{n} \cdot \epsilon + (1+\epsilon) \cdot \sum_{q \geq 1} \sum_{i \in T \cap \mathcal{W}_q} w_i \leq \epsilon + (1+\epsilon) \cdot w(T) \ .
$$

Using similar arguments in the opposite direction, we have

$$
\begin{aligned}
\tilde{r}_i \cdot \pi_{\tilde{w}}(i,S) &= \tilde{r}_i \cdot \sum_{T \subseteq S : i \in T} \Pr\left[C_S = T\right] \cdot \frac{\tilde{w}_i}{1+\tilde{w}(T)} \\
&\leq r_i \cdot \sum_{T \subseteq S : i \in T} \Pr\left[C_S = T\right] \cdot \frac{w_i}{1+w(T)-\epsilon} \\
&\leq (1+2\epsilon) \cdot r_i \cdot \sum_{T \subseteq S : i \in T} \Pr\left[C_S = T\right] \cdot \frac{w_i}{1+w(T)} \\
&= (1+2\epsilon) \cdot r_i \cdot \pi_w(i,S) \ ,
\end{aligned}
$$

Here, the first inequality holds since

$$
\tilde{w}(T) \geq \sum_{q \geq 1} \sum_{i \in T \cap \mathcal{W}_q} \tilde{w}_i \geq \sum_{q \geq 1} \sum_{i \in T \cap \mathcal{W}_q} w_i + \left( \sum_{i \in T \cap \mathcal{W}_0} w_i - \epsilon \right) = w(T) - \epsilon \ .
$$

### EC.1.2. Proof of Claim EC.2

Let $\tilde{\pi}(i,S)$ be the choice probability of item $i$ out of the assortment $S$, with respect to the rounded parameters. We simply have to show that $\tilde{\pi}(i,S) \geq (1-\epsilon) \cdot \pi_{\tilde{w}}(i,S)$. For this purpose, we utilize representation (1) of the choice probabilities to obtain

$$
\tilde{\pi}(i,S) = \tilde{\lambda}_i \cdot \mathbb{E}\left[ \frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(\tilde{C}_S^{-i})} \right] \ ,
$$

where $\tilde{C}_S^{-i}$ represents the random consideration set induced by $S \setminus \{i\}$ with respect to the rounded consideration probabilities $\{\tilde{\lambda}_i\}_{i \in [n]}$. Consequently, we have

$$
\begin{aligned}
\tilde{\pi}(i,S) &= \tilde{\lambda}_i \cdot \mathbb{E}\left[ \frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(\tilde{C}_S^{-i})} \right] \\
&\geq (1-\epsilon) \cdot \lambda_i \cdot \mathbb{E}\left[ \frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(\tilde{C}_S^{-i})} \right] \\
&\geq (1-\epsilon) \cdot \lambda_i \cdot \mathbb{E}\left[ \frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(C_S^{-i})} \right] \\
&= (1-\epsilon) \cdot \pi_{\tilde{w}}(i,S) \ ,
\end{aligned}
$$

where the first inequality follows from the construction of the rounded consideration probabilities. The second inequality proceeds from the stochastic orderings $\tilde{w}(\tilde{C}_S^{-i}) \sim \sum_{j\in S\setminus\{i\}} \tilde{w}_j \cdot B(\tilde{\lambda}_j) \preceq \sum_{j\in S\setminus\{i\}} \tilde{w}_j \cdot B(\lambda_j) \sim \tilde{w}(C_S^{-i})$, where the inequality holds since the usual stochastic order is closed under convolution and $\tilde{\lambda}_j \le \lambda_j$ for every $j \in S \setminus \{i\}$.

### EC.1.3. Proof of Claim EC.3

To prove the desired claim, it is sufficient to show that $\mathbb{E}_{S'\sim B^S}[\pi_{\tilde{w}}(i,S')] \ge (1-\epsilon)\cdot\tilde{\pi}(i,S)$. To this end, we note that

$$\begin{aligned}
\mathbb{E}_{S'\sim B^S}\left[\pi_{\tilde{w}}\left(i,S'\right)\right] &= \mathbb{E}_{S'\sim B^S}\left[(1-\epsilon)\cdot\lambda_i\cdot\mathbb{E}\left[\frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(C_{S'}^{-i})}\right]\right]\\
&\ge \lambda_i\cdot\mathbb{E}\left[\frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(\tilde{C}_S^{-i})}\right]\\
&\ge (1-\epsilon)\cdot\tilde{\lambda}_i\cdot\mathbb{E}\left[\frac{\tilde{w}_i}{1+\tilde{w}_i+\tilde{w}(\tilde{C}_S^{-i})}\right]\\
&= (1-\epsilon)\cdot\tilde{\pi}\left(i,S\right)\ ,
\end{aligned}$$

where the equalities follow from representation (1) of the choice probabilities. The first inequality proceeds from the stochastic ordering $\tilde{w}(C_{S'}^{-i}) \sim \sum_{j\in S\setminus\{i\}} \tilde{w}_j \cdot B((1-\epsilon)\cdot\lambda_j) \preceq \sum_{j\in S\setminus\{i\}} \tilde{w}_j \cdot B(\tilde{\lambda}_j) \sim \tilde{w}(\tilde{C}_S^{-i})$ where $S'$ is the random realization of the decomposable assortment $B^S$. Here, the stochastic inequality holds since the usual stochastic order is closed under convolution and $\tilde{\lambda}_j \ge (1-\epsilon)\cdot\lambda_j$ for every $j \in S\setminus\{i\}$.

## Appendix EC.2: Proof of Theorem 4

In what follows, our objective is to establish Theorem 4, which extends our algorithmic results for general instances of the click-based MNL assortment problem. As explained in Section 4 and Appendix EC.1, Assumption 1 can be enforced by slightly altering our input parameters with a negligible loss of optimality, given the sensitivity analysis of Lemma 4. As such, in the remainder of this appendix, we assume that Assumption 1 is satisfied and focus our attention on eliminating Assumption 2. The latter states that $Q_{\max} - Q_{\min} = O(\frac{1}{\epsilon}\cdot\log n)$, which means that the ratio of extremal preference weights is polynomially bounded.

Now, to handle the general case, in which the extremal preference weights of items are not necessarily polynomially related, we will show how to approximately decompose any given instance into a collection of bounded-ratio ones, glued together by means of dynamic programming. Within this framework, our approach directly leverages the approximation scheme provided in Theorem 3 as a subroutine. Our exposition of this approach is organized as follows: First, we establish in Appendix EC.2.1 a constrained version of Theorem 3, showing that our approximation scheme can easily handle an additional knapsack-like constraint on the selected decomposable assortment.

Next, we present in Appendix EC.2.2 a decomposition method that breaks any given instance into subproblems that satisfy the bounded-ratio assumption. In Appendices EC.2.3 and EC.2.4, we formulate a dynamic program that sequentially examines these subproblems; at each state, the dynamic program calls our approximation scheme for the constrained bounded-ratio setting as a subroutine.

### EC.2.1. Approximation scheme for the emptiness-constrained problem

Here, we argue that the approximation scheme of Theorem 3 can be generalized to handle an additional constraint on the decomposable assortment, termed *the emptiness constraint.* This constraint imposes a lower bound on the probability that the induced consideration set is empty. Specifically, given $\Phi \in [0,1]$, the emptiness constraint requires that the decomposable assortment $B$ satisfies $\Pr_{S\sim B}[C_S = \emptyset] \geq \Phi$.

In the next theorem, we state our main algorithmic result in the emptiness-constrained setting. In what follows, we still consider instances that satisfy Assumptions 1 and 2. Let $S^*$ be an optimal (deterministic) assortment for the click-based MNL assortment problem subject to the emptiness constraint with parameter $\Phi$.

THEOREM EC.1. *For any accuracy level $\epsilon > 0$, we can compute a decomposable assortment $B$ such that $\mathbb{E}_{S\sim B}[\mathcal{R}(S)] \geq (1-\epsilon)\cdot\mathcal{R}(S^*)$ and $\Pr_{S\sim B}[C_S = \emptyset] \geq \Phi$. The running time of our algorithm is $O(n^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$.*

We devote the remainder of this section to proving Theorem EC.1. Since our algorithmic approach is nearly identical to that of the unconstrained setting, we only focus on the differences between these approaches. Our analysis directly builds on the one presented in Appendix D for the unconstrained setting.

*Modified approximation scheme.* We begin by describing how our approximation scheme is modified to account for the emptiness constraint. The algorithm is nearly identical to that of the unconstrained setting with the only exception that, as we enumerate over the collection of input parameters $\Omega$, we only consider decomposable assortments the meet the emptiness criterion. Specifically, we run the specialized algorithms with all input parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}) \in \Omega$ to construct $B_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$ and $B_{\text{likely}} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}, B_{\text{unlike}})$. The algorithm eventually returns the resulting decomposable assortment $B = B_{\text{unlike}} \cup B_{\text{likely}}$ of maximum expected revenue out of all those that satisfy the constraint $\Pr_{S\sim B}[C_S = \emptyset] \geq \Phi$. That is, any decomposable assortment $B$ with $\Pr_{S\sim B}[C_S = \emptyset] < \Phi$ is discarded in Step 3 of the algorithm described in Section 3.4. To summarize, our modified algorithm proceeds as follows:

1. Generate the collection of input parameters $\Omega$.
2. For every $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}) \in \Omega$:

(a) Compute $B_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$.

(b) Compute $B_{\text{likely}} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}, B_{\text{unlike}})$.

3. Return the decomposable assortment $B = B_{\text{unlike}} \cup B_{\text{likely}}$ of maximum expected revenue out of those generated in Step 2 that additionally satisfy the requirement $\Pr_{S\sim B}[C_S = \emptyset] \geq \Phi$.

*Analysis.* Similar to the notation of Appendix D, we denote by $B = B_{\text{unlike}} \cup B_{\text{likely}}$ the decomposable assortment returned by the modified approximation scheme. Recall that, for purposes of analysis, we fix a collection of input parameters $(\{\hat{c}_q, \hat{\alpha}_q\}_{\mathcal{Q}}, \{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}) \in \Omega$ that satisfy Properties 1-3 of Claim 1. We denote by $\hat{B} = \hat{B}_{\text{unlike}} \cup \hat{B}_{\text{likely}}$ the resulting decomposable assortment, generated in Step 2 of our approximation scheme, i.e., $\hat{B}_{\text{unlike}} = \texttt{MinKnapsack}(\{(\hat{c}_q, \hat{\alpha}_q)\}_{\mathcal{Q}})$ and $\hat{B}_{\text{likely}} = \texttt{Greedy}(\{\hat{c}_{p,q}\}_{[P_{\max}]\times\mathcal{Q}}, \hat{B}_{\text{unlike}})$. We let $\hat{U}_q$ be the collection of items generated by solving the min-knapsack instance over the class of items $\mathcal{U}_q$, and let $\hat{B}_q$ be the resulting decomposable assortment specified by $\texttt{MinKnapsack}(\cdot)$ within this class. Similarly, we let $\hat{S}_T \subseteq \mathcal{N}_{\text{likely}}$ be the subset of likely items constructed in the last iteration of $\texttt{Greedy}(\cdot)$.

A close inspection of our analysis of the unconstrained setting in Appendix D.4 reveals that an identical performance guarantee holds in the emptiness-constrained setting. That is, using precisely the same line of reasoning as in inequality (19), we obtain $\mathbb{E}_{S\sim\hat{B}}[\mathcal{R}(S)] \geq (1-6\epsilon)\cdot\mathcal{R}(S^*)$. Hence, in order to establish Theorem EC.1, it remains to show that the decomposable assortment $\hat{B}$ indeed satisfies the emptiness criterion $\Pr_{S\sim\hat{B}}[C_S = \emptyset] \geq \Phi$, and thus, it is not discarded in Step 3 of our modified approximation scheme.

To this end, we relate $\Pr_{S\sim\hat{B}_{\text{unlike}}}[C_S = \emptyset]$ to the analogous probability with respect to $S^* \cap \mathcal{N}_{\text{unlike}}$. The lower bound stated in Lemma EC.1 is the first step towards showing that our final assortment $\hat{B}$ meets the emptiness constraint.

Lemma EC.1. $\Pr_{S\sim\hat{B}_{\text{unlike}}}[C_S = \emptyset] \geq \Pr[C_{S^*\cap\mathcal{N}_{\text{unlike}}} = \emptyset]$.

*Proof.* By definition of the decomposable assortment $\hat{B}_{\text{unlike}}$, we have:

$$\begin{aligned}
\Pr_{S\sim\hat{B}_{\text{unlike}}}[C_S = \emptyset] &= \prod_{q\in\mathcal{Q}} \Pr_{S\sim\hat{B}^q}[C_S = \emptyset] \\
&= \prod_{q\in\mathcal{Q}} \prod_{i\in\hat{U}_q} (1-(1-\epsilon)\cdot\lambda_i) \\
&\geq \prod_{q\in\mathcal{Q}} \prod_{i\in\hat{U}_q} e^{-\lambda_i} \\
&= \prod_{q\in\mathcal{Q}} e^{-\sum_{i\in\hat{U}_q}\lambda_i} ,
\end{aligned}$$

where the second equality proceeds from our construction of the decomposable assortment $\hat{B}^q$, where $\Pr[\hat{B}_i^q = 1] = 1-\epsilon$ for every $i \in U_q$. The subsequent inequality holds since $1-(1-\epsilon)\cdot\lambda_i \geq e^{-\lambda_i}$

for $\lambda_i \in [0,\epsilon]$, noting that we indeed have $\lambda_i \leq \epsilon$ since all items $i \in \hat{U}_q$ are unlikely. Recalling that inequality (24) guarantees that $\sum_{i\in \hat{U}_q} \lambda_i \leq \sum_{i\in S^*\cap \mathcal{U}_q} \lambda_i$, we conclude that

$$\begin{aligned}\mathrm{Pr}_{S\sim \hat{B}_{\mathrm{unlike}}}[C_S=\emptyset] &\geq \prod_{q\in\mathcal{Q}} e^{-\sum_{i\in S^*\cap\mathcal{U}_q}\lambda_i}\\ &\geq \prod_{q\in\mathcal{Q}}\prod_{i\in S^*\cap\mathcal{U}_q}(1-\lambda_i)\\ &= \mathrm{Pr}\left[C_{S^*\cap\mathcal{N}_{\mathrm{unlike}}}=\emptyset\right] \ .\end{aligned}$$

□

To complete the proof of Theorem EC.1, we verify that the decomposable assortment $\hat{B}$ meets the emptiness criterion, as shown in the next lemma.

LEMMA EC.2. $\mathrm{Pr}_{S\sim\hat{B}}[C_S=\emptyset]\geq\Phi$.

*Proof.* Observe that

$$\begin{aligned}\mathrm{Pr}_{S\sim\hat{B}}[C_S=\emptyset] &= \mathrm{Pr}_{S\sim\hat{B}_{\mathrm{unlike}}}[C_S=\emptyset]\cdot\mathrm{Pr}_{S\sim\hat{B}_{\mathrm{likely}}}[C_S=\emptyset]\\ &\geq \mathrm{Pr}\left[C_{S^*\cap\mathcal{N}_{\mathrm{unlike}}}=\emptyset\right]\cdot\mathrm{Pr}\left[C_{S^*\cap\mathcal{N}_{\mathrm{likely}}}=\emptyset\right]\\ &= \mathrm{Pr}\left[C_{S^*}=\emptyset\right]\\ &\geq \Phi\ ,\end{aligned}$$

where the first equality holds since the decomposable assortments $\hat{B}_{\mathrm{likely}}$ and $\hat{B}_{\mathrm{unlike}}$ are independent. The first inequality above follow from Lemmas EC.1 and Claim 6. In particular, in order to apply Claim 6, we note the assortment $S$ sampled from $\hat{B}_{\mathrm{likely}}$ is deterministic and equal to $\hat{S}_T$, thus $S\subseteq S^*\cap\mathcal{N}_{\mathrm{likely}}$ with probability 1. □

### EC.2.2. Well-separated weight clusters

We begin by defining a way to partition the collection of weight classes into well-separated clusters, which constitute the elementary units of our decomposition into bounded-ratio instances. The concrete meaning of "well-separated clusters" will become clear once we introduce the necessary definitions and constructions.

*Creating weight clusters.* We begin by reminding the reader that, following Assumption 1, the entire set of items is partitioned into weights classes $\mathcal{W}_0,\mathcal{W}_1,\ldots$, where $\mathcal{W}_q=\{i\in[n]: w_i=(1+\epsilon)^q\cdot\frac{\epsilon}{n}\}$. For ease of notation, let $Q_{\mathrm{skip}}=2\cdot\lceil\log_{1+\epsilon}(\frac{n}{\epsilon})\rceil$ and $Q_{\mathrm{pick}}=\lceil\frac{1}{\epsilon}\rceil\cdot Q_{\mathrm{skip}}$. In addition, fix a non-positive integer $s\in[-Q_{\mathrm{skip}}-Q_{\mathrm{pick}}+1,0]$, which is referred to as the shifting index; the value of this index will be determined shortly.

Given these parameters, we initially define a sequence of pairwise-disjoint subsets $I_1^s, I_2^s,\ldots$, where each $I_\ell^s$ corresponds to a consecutive block of non-negative integers. These subsets are iteratively defined as follows:

- First, we define $I_1^s$ as the interval starting at $s$ and consisting of $Q_{\text{pick}}$ successive indices, namely, $I_1^s = \{s, s+1, \ldots, s+Q_{\text{pick}}-1\}$.
- Next, we skip $Q_{\text{skip}}$ indices, and define $I_2^s$ as the interval formed by the next $Q_{\text{pick}}$ indices. Namely, $I_2^s = \{s+Q_{\text{pick}}+Q_{\text{skip}}, s+Q_{\text{pick}}+Q_{\text{skip}}+1, \ldots, s+2Q_{\text{pick}}+Q_{\text{skip}}-1\}$.
- Similarly, we skip the next $Q_{\text{skip}}$ indices and pick the next $Q_{\text{pick}}$ ones to define $I_3^s = \{s+2Q_{\text{pick}}+2Q_{\text{skip}}, s+2Q_{\text{pick}}+2Q_{\text{skip}}+1, \ldots, s+3Q_{\text{pick}}+2Q_{\text{skip}}\}$. So on and so forth, as we alternate between skipping $Q_{\text{skip}}$ indices and forming intervals comprised of $Q_{\text{pick}}$ consecutive indices.

With respect to this sequence, we introduce its corresponding sequence of clusters $\mathcal{C}_1^s, \ldots, \mathcal{C}_L^s$, where each cluster $\mathcal{C}_\ell^s$ is the union of weight classes over all indices in $I_\ell^s$, i.e., $\mathcal{C}_\ell^s = \bigcup_{q \in I_\ell^s} \mathcal{W}_q$. Here, $L$ stands for the largest integer $\ell$ for which $\mathcal{C}_\ell^s$ is non-empty.

*Cluster properties.* We proceed by highlighting two structural properties of the sequence $\mathcal{C}_1^s, \ldots, \mathcal{C}_L^s$, which come as immediate implications of our construction:

1. For every $\ell \in [L]$, the preference weights of any two items within the cluster $\mathcal{C}_\ell^s$ differ by a factor of at most $(2n/\epsilon)^{2\lceil 1/\epsilon \rceil}$. This property holds since $\mathcal{C}_\ell^s$ is comprised of at most $Q_{\text{pick}} = 2 \cdot \lceil \frac{1}{\epsilon} \rceil \cdot \lceil \log_{1+\epsilon}(\frac{n}{\epsilon}) \rceil$ consecutive weight classes, implying that

$$\frac{\max_{i \in \mathcal{C}_\ell^s} w_i}{\min_{i \in \mathcal{C}_\ell^s} w_i} \leq (1+\epsilon)^{Q_{\text{pick}}} \leq \left( \frac{2n}{\epsilon} \right)^{2\lceil 1/\epsilon \rceil} .$$

Hence, each cluster taken in isolation forms a bounded-ratio instance, where all item weights are polynomially-related.

2. For every $\ell \in [L-1]$, the preference weights of any two items, one in $\mathcal{C}_\ell^s$ and the other in $\mathcal{C}_{\ell+1}^s$, differ by a factor of at least $(n/\epsilon)^2$. This property holds since $\mathcal{C}_\ell^s$ and $\mathcal{C}_{\ell+1}^s$ are separated by $Q_{\text{skip}} = 2 \cdot \lceil \log_{1+\epsilon}(\frac{n}{\epsilon}) \rceil$ consecutive weight classes, implying that

$$\frac{\min_{i \in \mathcal{C}_{\ell+1}^s} w_i}{\max_{i \in \mathcal{C}_\ell^s} w_i} \geq (1+\epsilon)^{Q_{\text{skip}}} \geq \left( \frac{n}{\epsilon} \right)^2 .$$

Consequently, due to properties 1 and 2, we say that the clusters $\mathcal{C}_1^s, \ldots, \mathcal{C}_L^s$ are *well-separated.*

*The effect of eliminating unclustered weight classes.* A careful examination of the sequence of clusters $\mathcal{C}_1^s, \ldots, \mathcal{C}_L^s$ reveals that various weight classes do not appear in any of these clusters, corresponding to sequences of $Q_{\text{skip}}$ indices over which we skip during the construction of $I_1^s, I_2^s, \ldots$. Therefore, since the dynamic programming approach in Appendices EC.2.3 and EC.2.4 will limit assortment decisions to the clustered weight classes $\mathcal{C}^s = \bigcup_{\ell \in [L]} \mathcal{C}_\ell^s$, the crucial question is: Why can we overlook revenue contribution of unclustered ones? Letting $S^* \subseteq [n]$ be an optimal assortment, the next claim shows that the shifting index $s$ can be chosen such that we incur a negligible loss in optimality when restricting $S^*$ to the clustered weight classes $\mathcal{C}^s$.

Lemma EC.3. *There exists a shifting index $s^* \in [-Q_{\text{skip}} - Q_{\text{pick}} + 1, 0]$ for which $\mathcal{R}(S^* \cap \mathcal{C}^{s^*}) \geq (1-\epsilon) \cdot \mathcal{R}(S^*)$.*

*Proof.* Our proof is based on a simple application of the probabilistic method (Alon and Spencer 2016). Specifically, let $\sigma$ be a discrete random variable, uniformly distributed over the set of integers in $[-Q_{\text{skip}} - Q_{\text{pick}} + 1, 0]$. In this case, $S^* \cap \mathcal{C}^\sigma$ is clearly a random assortment, whose expected revenue is given by

$$\begin{aligned}
\mathbb{E}_\sigma\left[\mathcal{R}\left(S^* \cap \mathcal{C}^\sigma\right)\right] &= \sum_{i \in S^*} \mathbb{E}_\sigma\left[\mathbb{I}\left[i \in \mathcal{C}^\sigma\right] \cdot r_i \cdot \pi\left(i, S^* \cap \mathcal{C}^\sigma\right)\right] \\
&= \sum_{i \in S^*} \Pr_\sigma\left[i \in \mathcal{C}^\sigma\right] \cdot r_i \cdot \mathbb{E}_\sigma\left[\pi\left(i, S^* \cap \mathcal{C}^\sigma\right) \middle| i \in \mathcal{C}^\sigma\right] \\
&\geq \sum_{i \in S^*} \Pr_\sigma\left[i \in \mathcal{C}^\sigma\right] \cdot r_i \cdot \pi\left(i, S^*\right) \\
&= \frac{Q_{\text{pick}}}{Q_{\text{pick}} + Q_{\text{skip}}} \cdot \sum_{i \in S^*} r_i \cdot \pi\left(i, S^*\right) \\
&\geq (1-\epsilon) \cdot \mathcal{R}\left(S^*\right) \ .
\end{aligned}$$

Here, the first inequality follows from representation (1) of the choice probabilities, which implies that $\pi(i, S^* \cap \mathcal{C}^\sigma) \geq \pi(i, S^*)$ for any item $i \in S^* \cap \mathcal{C}^\sigma$. The third equality can easily be derived by observing that, due to the uniform choice of $\sigma$, any given weight class is clustered into $\mathcal{C}^\sigma$ with probability $\frac{Q_{\text{pick}}}{Q_{\text{pick}} + Q_{\text{skip}}}$. The last inequality holds since $Q_{\text{pick}} = \lceil \frac{1}{\epsilon} \rceil \cdot Q_{\text{skip}}$. In summary, since we have just shown that $\mathbb{E}_\sigma[\mathcal{R}(S^* \cap \mathcal{C}^\sigma)] \geq (1-\epsilon) \cdot \mathcal{R}(S^*)$, there exists at least one realization $s^*$ of $\sigma$ for which $\mathcal{R}(S^* \cap \mathcal{C}^{s^*}) \geq (1-\epsilon) \cdot \mathcal{R}(S^*)$. $\square$

In the remainder of this section, we assume that the shifting index $s^*$ mentioned in Lemma EC.3 is known in advance. From an algorithmic standpoint, the value of $s^*$ can be guessed by exhaustively enumerating all $O(\frac{1}{\epsilon^2} \log \frac{n}{\epsilon})$ integers in $[-Q_{\text{skip}} - Q_{\text{pick}} + 1, 0]$. Furthermore, out of the clusters $\mathcal{C}_1^{s^*}, \ldots, \mathcal{C}_L^{s^*}$, we focus our attention on the subsequence $\mathcal{C}_{\psi(1)}^{s^*}, \ldots, \mathcal{C}_{\psi(L')}^{s^*}$ of non-empty clusters. In what follows, the reference to $s^*$ and $\psi$ will be implicit, meaning that the clusters $\mathcal{C}_{\psi(1)}^{s^*}, \ldots, \mathcal{C}_{\psi(L')}^{s^*}$ will simply be denoted by $\mathcal{C}_1, \ldots, \mathcal{C}_L$.

### EC.2.3. Value approximation

We proceed by describing an approximate decomposition of the expected revenue $\mathbb{E}_{S \sim B}[\mathcal{R}(S)]$ generated by any decomposable assortment $B$, whose realization is contained within the clusters $\mathcal{C}_1, \ldots, \mathcal{C}_L$, based on the individual revenue contribution of each such cluster. This decomposition will motivate our dynamic programming approach for computing near-optimal assortments, formally presented in Appendix EC.2.4.

To describe our value approximation, for every item $i \in \mathcal{C} = \bigcup_{\ell \in [L]} \mathcal{C}_\ell$, we denote by $\varphi(i)$ the index of the unique cluster that contains item $i$. In addition, let $\mathcal{C}_{>\ell}$ designate the union of all clusters indexed strictly greater than $\ell$, i.e., $\mathcal{C}_{>\ell} = \bigcup_{k \geq \ell+1} \mathcal{C}_k$. We remind the reader that representation (1) of the choice probabilities shows that the expected revenue of any decomposable assortment $B$ can be written as:

$$\mathbb{E}_{S \sim B}\left[\mathcal{R}(S)\right] = \mathbb{E}_{S \sim B}\left[\sum_{i \in S} r_i \cdot \pi(i,S)\right] = \mathbb{E}_{S \sim B}\left[\sum_{i \in S} r_i \lambda_i w_i \cdot \mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i})}\right]\right] . \qquad \text{(EC.1)}$$

In Lemma EC.4 below, we provide a lower bound on the inner expectation term $\mathbb{E}[\frac{1}{1+w_i+w(C_S^{-i})}]$ for any fixed assortment $S$. The important feature of this lower bound is that it can computed knowing only the following ingredients:

1. The distribution of the random subset $S \cap \mathcal{C}_{\varphi(i)}$.
2. The probability that the induced consideration set $C_{S \cap \mathcal{C}_{>\varphi(i)}}$ is empty.

Consequently, by recalling that the preference weights within a single cluster are polynomially-related (see property 1 in Appendix EC.2.2), and by observing that the second ingredient above corresponds to an emptiness constraint, we facilitate an approximate reduction to the bounded-ratio setting studied in Appendix EC.2.1.

Lemma EC.4. *For every assortment $S \subseteq \mathcal{C}$ and item $i \in S$,*

$$\mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i})}\right] \geq (1-\epsilon) \cdot \Pr\left[C_{S \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i} \cap \mathcal{C}_{\varphi(i)})}\right] .$$

*Proof.* The proof proceeds by observing that

$$\begin{aligned}
&\mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i})}\right] \\
&\quad \geq \Pr\left[C_{S \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i})} \middle| C_{S \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \\
&\quad = \Pr\left[C_{S \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1 + w_i + w(C_S \cap (\bigcup_{\ell=1}^{\varphi(i)-1} \mathcal{C}_\ell)) + w(C_S^{-i} \cap \mathcal{C}_{\varphi(i)})}\right] \\
&\quad \geq (1-\epsilon) \cdot \Pr\left[C_{S \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1 + w_i + w(C_S^{-i} \cap \mathcal{C}_{\varphi(i)})}\right] .
\end{aligned}$$

Here, the middle equality follows from the independence of $C_S \cap \mathcal{C}_{>\varphi(i)}$ and $C_S^{-i} \cap (\bigcup_{\ell=1}^{\varphi(i)} \mathcal{C}_\ell)$. The last inequality holds since

$$w\left(C_S \cap \left(\bigcup_{\ell=1}^{\varphi(i)-1} \mathcal{C}_\ell\right)\right) \leq \left|C_S \cap \left(\bigcup_{\ell=1}^{\varphi(i)-1} \mathcal{C}_\ell\right)\right| \cdot \left(\frac{\epsilon}{n}\right)^2 \cdot \min_{j \in \mathcal{C}_{\varphi(i)}} w_j \leq \epsilon w_i ,$$

given that the clusters $\mathcal{C}_1, \ldots, \mathcal{C}_L$ are well-separated (see property 2 in Appendix EC.2.2). □

In the next claim, we argue that the lower bound stated in Lemma EC.4 is useful in obtaining a tight approximation for the optimal expected revenue. To formalize this notion, we overload our previous notation by letting $S^* \subseteq \mathcal{C}$ be an optimal assortment restricted to weight classes clustered into $\mathcal{C}$. As shown in Lemma EC.3, this restriction has negligible losses in optimality in comparison to arbitrary assortments, that may offer items from unclustered weight classes as well.

LEMMA EC.5. *For any* $\epsilon \in (0,1/2)$,

$$\mathcal{R}(S^*) \le (1+2\epsilon) \cdot \sum_{i \in S^*} r_i \lambda_i w_i \cdot \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^*} \cap \mathcal{C}_{\varphi(i)})}\right] .$$

*Proof.* Using arguments analogous to those appearing in the proof of Lemma EC.4, we derive an upper bound on the expectation $\mathbb{E}[\frac{1}{1+w_i+w(C^{-i}_{S^*})}]$, for every item $i \in S^*$, by observing that

$$\begin{aligned}
&\mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^*})}\right] \\
&\quad = \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} \neq \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^*})} \middle| C_{S^* \cap \mathcal{C}_{>\varphi(i)}} \neq \emptyset\right] \\
&\qquad + \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^*})} \middle| C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \\
&\quad \le \frac{1}{1+w_i+(n/\epsilon)^2 \cdot w_i} + \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^* \cap \mathcal{C}_{\varphi(i)}})}\right] \\
&\quad \le \frac{\epsilon}{n} \cdot \frac{1}{1+w_i} + \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^* \cap \mathcal{C}_{\varphi(i)}})}\right] .
\end{aligned} \tag{EC.2}$$

Here, the first inequality holds since the clusters $\mathcal{C}_1, \ldots, \mathcal{C}_L$ are well-separated, implying that conditional on the random consideration set $C_{S^* \cap \mathcal{C}_{>\varphi(i)}}$ being non-empty, its random total weight is at least $(\frac{n}{\epsilon})^2 \cdot w_i$. The second inequality is obtained by observing that $1 + w_i + (\frac{n}{\epsilon})^2 \cdot w_i \ge \frac{n}{\epsilon} \cdot (1+w_i)$ whenever $w_i \ge \frac{\epsilon}{n}$. The property that the preference weight of each item is at least $\frac{\epsilon}{n}$ follows from the weight transformation in Section 4, which we analyzed in Appendix EC.1.

Now, according to the revenue representation (EC.1), the optimal expected revenue $\mathcal{R}(S^*)$ can be upper bounded as follows:

$$\begin{aligned}
\mathcal{R}(S^*) &= \sum_{i \in S^*} r_i \lambda_i w_i \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^*})}\right] \\
&\le \frac{\epsilon}{n} \cdot \sum_{i \in S^*} \frac{r_i \lambda_i w_i}{1+w_i} + \sum_{i \in S^*} r_i \lambda_i w_i \cdot \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^* \cap \mathcal{C}_{\varphi(i)}})}\right] \\
&\le \epsilon \cdot \mathcal{R}(S^*) + \sum_{i \in S^*} r_i \lambda_i w_i \cdot \Pr\left[C_{S^* \cap \mathcal{C}_{>\varphi(i)}} = \emptyset\right] \cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S^* \cap \mathcal{C}_{\varphi(i)}})}\right] ,
\end{aligned}$$

where the first inequality is obtained by plugging the upper bound (EC.2) and the second inequality holds since $\mathcal{R}(S^*) \ge r_i \lambda_i \cdot \frac{w_i}{1+w_i}$ for any item $i \in \mathcal{C}$. To see this, note that the optimality of $S^*$ ensures that its expected revenue $\mathcal{R}(S^*)$ is at least the expected revenue of the singleton assortment $\{i\}$, which is precisely $r_i \lambda_i \cdot \frac{w_i}{1+w_i}$. The desired claim is now obtained by rearranging the above bound.

□

Based on the preceding discussion, Lemmas EC.4 and EC.5 immediately imply that, in order to compute a random assortment whose expected revenue is within factor $1-O(\epsilon)$ of the optimal revenue $\mathcal{R}(S^*)$, it suffices to optimize the value approximation $\tilde{\mathcal{R}}(\cdot)$, defined over the set of decomposable assortments $B$ as follows:

$$\begin{aligned}
\tilde{\mathcal{R}}(B) &= \mathbb{E}_{S\sim B}\left[\sum_{i\in S\cap\mathcal{C}} \Pr\left[C_{S\cap\mathcal{C}_{>\varphi(i)}}=\emptyset\right]\cdot r_i\lambda_i w_i\cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S\cap\mathcal{C}_{\varphi(i)}})}\right]\right] \\
&= \mathbb{E}_{S\sim B}\left[\sum_{\ell=1}^{L}\left(\Pr\left[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset\right]\cdot \sum_{i\in S\cap\mathcal{C}_\ell} r_i\lambda_i w_i\cdot \mathbb{E}\left[\frac{1}{1+w_i+w(C^{-i}_{S\cap\mathcal{C}_\ell})}\right]\right)\right] \\
&= \sum_{\ell=1}^{L}\Pr_{S\sim B}\left[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset\right]\cdot \mathbb{E}_{S\sim B}\left[\mathcal{R}\left(S\cap\mathcal{C}_\ell\right)\right] \ , \qquad \text{(EC.3)}
\end{aligned}$$

where the second equality is simply a decomposition of $S\cap\mathcal{C}$ into the clusters $\mathcal{C}_1,\dots,\mathcal{C}_L$, and the third equality follows by noting that the inner sum is precisely the expected revenue generated by $S\cap\mathcal{C}_\ell$, conditional on $S$ being the realization of the decomposable assortment $B$.

### EC.2.4. Gluing approximate dynamic program

*Continuous dynamic program.* The important observation is that the problem of maximizing the value approximation of Appendix EC.2.3 can be formulated as a dynamic program whose state space is described by the following parameters:

- An index $\ell\in[L]$, corresponding to the cluster $\mathcal{C}_\ell$ that is currently being considered.
- A continuous variable $\theta\in[0,1]$, specifying the probability that none of the items picked from higher-index clusters $\mathcal{C}_{>\ell}$ will appear in our random consideration set.

Given these parameters, we define $F(\ell,\theta)$ as the maximum possible value of $\sum_{k=1}^{\ell}\Pr_{S\sim B}[C_{S\cap\mathcal{C}_{>k}}=\emptyset]\cdot\mathbb{E}_{S\sim B}[\mathcal{R}(S\cap\mathcal{C}_k)]$ over all decomposable assortments $B$, subject to the probabilistic constraint $\Pr_{S\sim B}[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset]=\theta$. An immediate implication of this definition is that, in order to compute an assortment maximizing the value approximation $\tilde{\mathcal{R}}(\cdot)$, we can equivalently compute the one for which $F(L,1)$ is attained. Concurrently, the definition of $F$ can be directly used to express this function in recursive form, by observing that for $\ell\ge 2$,

$$F(\ell,\theta)=\max_{B_\ell\in\mathcal{B}_\ell}\left\{\theta\cdot\mathbb{E}_{S\sim B_\ell}\left[\mathcal{R}(S)\right]+F\left(\ell-1,\theta\cdot\Pr_{S\sim B_\ell}\left[C_S=\emptyset\right]\right)\right\} \ . \qquad \text{(EC.4)}$$

Here, $\mathcal{B}_\ell$ designates the collection of decomposable assortments over the set of items $\mathcal{C}_\ell$. In addition, for $\ell=1$, we have

$$F(\ell,\theta)=\theta\cdot\max_{B_1\in\mathcal{B}_1}\mathbb{E}_{S\sim B_1}\left[\mathcal{R}(S)\right] \ . \qquad \text{(EC.5)}$$

Due to the continuity of $\theta\in[0,1]$, the function $F$ should be viewed as a recursive characterization of optimal random assortments for our value approximation $\tilde{\mathcal{R}}(\cdot)$.

*Discretization.* For this reason, we explain how to approximately solve the recursive equations (EC.4) and (EC.5) by leveraging the bounded-ratio algorithm given in Theorem EC.1 as a subroutine. To simplify the exposition, we assume that $\max_{i\in\mathcal{C}} \lambda_i \le 1-\epsilon$. By the sensitivity analysis established in Lemma 4, this assumption can be made without loss of generality.

Now, in order to discretize the state variable $\theta$, we define the finite set $\Theta = \{(1+\frac{\epsilon}{n})^{-k} : 0 \le k \le n\cdot(1+\lceil \log_{1+\frac{\epsilon}{n}}(\frac{1}{\epsilon})\rceil)\}$. For any $\Phi \in [0,1]$, let $B(\ell,\Phi)\in\mathcal{B}_\ell$ be the decomposable assortment returned by the approximation scheme stated in Theorem EC.1, when the underlying set of items is precisely the cluster $\mathcal{C}_\ell$, subject to the emptiness constraint $\Pr_{S\sim B}[C_S=\emptyset]\ge\Phi$.

With these definitions at hand, we introduce a discrete dynamic program $\tilde{F}$ over the state space $[L]\times\Theta$, i.e., a formulation where our second parameter $\theta$ is restricted to take values only within $\Theta$. This program is formally specified through the following recursive equations, for $\ell\ge 2$,

$$\tilde{F}(\ell,\theta) = \max_{\substack{\tilde{\kappa}\in\Theta:\\ \tilde{\kappa}\le\theta}} \left\{ \theta\cdot\mathbb{E}_{S\sim B(\ell,\tilde{\kappa}/\theta)}\left[\mathcal{R}(S)\right] + \tilde{F}(\ell-1,\tilde{\kappa}) \right\} \ . \tag{EC.6}$$

In addition, for $\ell=1$, we have

$$\tilde{F}(\ell,\theta) = \theta\cdot\mathbb{E}_{S\sim B(1,0)}\left[\mathcal{R}(S)\right] \ . \tag{EC.7}$$

*Analysis.* By property 1 in Appendix EC.2.2, within each cluster $\mathcal{C}_\ell$, the ratio between the extremal preference weights is upper-bounded by $(\frac{2n}{\epsilon})^{2\lceil 1/\epsilon\rceil}$. Hence, it is easy to verify that the instance of the constrained click-based MNL instance corresponding to each cluster satisfies Assumption 2. By leveraging the bounded-ratio algorithm of Theorem EC.1, we infer that our recursion can be solved in time $O((\frac{n}{\epsilon})^{O(\frac{1}{\epsilon^4}\log\frac{1}{\epsilon})})$. To establish the performance guarantee of our algorithm, let $\tilde{B}^*$ be the optimal decomposable assortment resulting from the solution of equations (EC.6) and (EC.7). Namely, $\tilde{B}^*$ is obtained by concatenating the decomposable assortments picked at each step of the recursion, in the optimal sequence of dynamic programming decisions that attain the function value $\tilde{F}(L,1)$. In the remainder of this section, we conclude the analysis by relating $\tilde{\mathcal{R}}(\tilde{B}^*)$ to the optimum value $F(L,1)$ of the continuous dynamic program $F$, which was shown earlier to capture the assortment maximizing our value approximation $\tilde{\mathcal{R}}(\cdot)$.

Lemma EC.6. *$\tilde{\mathcal{R}}(\tilde{B}^*)\ge(1-3\epsilon)\cdot F(L,1)$.*

To prove the desired inequality, we begin by considering the optimal sequence of dynamic programming decisions for equations (EC.6) and (EC.7) that yields the random assortment $\tilde{B}^*$. This sequence is denoted by $(\tilde{\theta}_1^*,\ldots,\tilde{\theta}_L^*)\in\Theta^L$, where $\tilde{\theta}_L^*=1$. Namely, $\tilde{B}^*=\bigcup_{\ell=1}^L B(\ell,\tilde{\Phi}_\ell^*)$ where $\tilde{\Phi}_\ell^* = \tilde{\theta}_{\ell-1}^*/\tilde{\theta}_\ell^*$ for every $\ell\in[1,L]$ and $\tilde{\theta}_0^*=0$. In the next claim, we relate our value approximation for the decomposable assortment $\tilde{B}^*$ to the value computed by the discretized dynamic program (EC.6)-(EC.7) along the sequence of decisions $(\tilde{\theta}_1^*,\ldots,\tilde{\theta}_L^*)$.

Claim EC.4. $\tilde{\mathcal{R}}(\tilde{B}^*) = \sum_{\ell=1}^{L} \tilde{\theta}^*_\ell \cdot \mathbb{E}_{S\sim B(\ell,\tilde{\Phi}^*_\ell)}[\mathcal{R}(S)]$.

*Proof.* By definition of $\tilde{\mathcal{R}}(\cdot)$ in equation (EC.3), we have:

$$\begin{aligned}\tilde{\mathcal{R}}\left(\tilde{B}^*\right) &= \mathbb{E}_{S\sim\tilde{B}^*}\left[\mathcal{R}\left(S\cap\mathcal{C}_L\right)\right] + \sum_{\ell=1}^{L-1}\mathrm{Pr}_{S\sim\tilde{B}^*}\left[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset\right]\cdot\mathbb{E}_{S\sim\tilde{B}^*}\left[\mathcal{R}\left(S\cap\mathcal{C}_\ell\right)\right]\\ &= \mathbb{E}_{S\sim B(L,\tilde{\Phi}^*_L)}\left[\mathcal{R}\left(S\right)\right] + \sum_{\ell=1}^{L-1}\left(\prod_{k=\ell+1}^{L}\mathrm{Pr}_{S\sim B(k,\tilde{\Phi}^*_k)}\left[C_S=\emptyset\right]\right)\cdot\mathbb{E}_{S\sim B(\ell,\tilde{\Phi}^*_\ell)}\left[\mathcal{R}\left(S\right)\right]\\ &= \mathbb{E}_{S\sim B(L,\tilde{\Phi}^*_L)}\left[\mathcal{R}\left(S\right)\right] + \sum_{\ell=1}^{L-1}\left(\prod_{k=\ell+1}^{L}\frac{\tilde{\theta}^*_{k-1}}{\tilde{\theta}^*_k}\right)\cdot\mathbb{E}_{S\sim B(\ell,\tilde{\Phi}^*_\ell)}\left[\mathcal{R}\left(S\right)\right]\\ &= \sum_{\ell=1}^{L}\tilde{\theta}^*_\ell\cdot\mathbb{E}_{S\sim B(\ell,\tilde{\Phi}^*_\ell)}\left[\mathcal{R}\left(S\right)\right]\ ,\end{aligned}$$

where the second equality holds since the random sets $C_{S(k_1)}$ and $C_{S(k_2)}$ are independent for every $k_1 \neq k_2$, where $S(k)$ is the random realization of $B(k,\tilde{\Phi}^*_k)$. The next equality follows from the definition of $B(k,\tilde{\Phi}^*_k)$ as a decomposable assortment that satisfies the emptiness constraint $\mathrm{Pr}_{S\sim B(k,\tilde{\Phi}^*_k)}[\mathcal{C}_S=\emptyset] \geq \tilde{\Phi}^*_k$. $\square$

Conversely, we now consider a decomposable assortment $\mathcal{A}^*$ that maximizes our value approximation $\tilde{\mathcal{R}}(\cdot)$, namely $\tilde{\mathcal{R}}(\mathcal{A}^*) = F(L,1)$. Consequently, we recursively define $\theta^*_L = 1$, and $\theta^*_\ell = \lfloor \theta^*_{\ell+1} \cdot \mathrm{Pr}_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_{\ell+1}}=\emptyset]\rfloor_\Theta$ for every $\ell \in [1, L-1]$, where $\lfloor\cdot\rfloor_\Theta$ is an operator that rounds its argument down to the nearest number in $\Theta$. Next, we define $\Phi^*_\ell = \theta^*_{\ell-1}/\theta^*_\ell$ for every $\ell \in [1,L]$, where $\theta^*_0 = 0$. It is not difficult to verify that the sequence of dynamic programming decisions $(\theta^*_1,\dots,\theta^*_L)$ is feasible with respect to the discretized dynamic program (**??**). To conclude, we relate the value generated by the sequence of dynamic programming decisions $(\theta^*_0,\theta^*_1,\dots,\theta^*_L)$ to $F(L,1)$.

Claim EC.5. $\sum_{\ell=1}^{L} \theta^*_\ell \cdot \mathbb{E}_{S\sim B(\ell,\Phi^*_\ell)}[\mathcal{R}(S)] \geq (1-3\epsilon)\cdot F(L,1)$.

Before presenting the proof of Claim EC.5, we explain how the desired inequality proceeds from the above two claims. By Claim EC.4, we have:

$$\begin{aligned}\tilde{\mathcal{R}}(\tilde{B}^*) &= \sum_{\ell=1}^{L}\tilde{\theta}^*_\ell\cdot\mathbb{E}_{S\sim B(\ell,\tilde{\Phi}^*_\ell)}\left[\mathcal{R}\left(S\right)\right]\\ &\geq \sum_{\ell=1}^{L}\theta^*_\ell\cdot\mathbb{E}_{S\sim B(\ell,\Phi^*_\ell)}\left[\mathcal{R}\left(S\right)\right]\\ &\geq (1-3\epsilon)\cdot F(L,1)\ ,\end{aligned}$$

where the first inequality holds by the optimality of the dynamic programming decisions $(\tilde{\theta}^*_1,\dots,\tilde{\theta}^*_L)$, and the last inequality immediately follows from Claim EC.5.

*Proof of Claim EC.5.* We begin by bounding the approximation error generated by the rounding operator $\lfloor\cdot\rfloor_\Theta$. By the definition of the set $\Theta$, for every $x \geq \epsilon^n \cdot (1+\frac{\epsilon}{n})^{-n}$, we necessarily have

$(1-\frac{\epsilon}{n})\cdot x \le \lfloor x \rfloor_\Theta \le x$. Hence, by our earlier assumption that $1-\lambda_i \ge \epsilon$ for every $i\in[n]$, and by using the relationships $\theta^*_{\ell-1} = \lfloor \theta^*_\ell \cdot \Pr_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_\ell}=\emptyset]\rfloor_\Theta$ and $\Phi^*_\ell = \theta^*_{\ell-1}/\theta^*_\ell$, a straightforward induction over $\ell\in[1,L]$ yields

$$\left(1-\frac{\epsilon}{n}\right)\cdot \Pr_{S\sim\mathcal{A}^*}\left[C_{S\cap\mathcal{C}_\ell}=\emptyset\right] \le \Phi^*_\ell \le \Pr_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_\ell}=\emptyset]\ . \tag{EC.8}$$

Thus, by noting that $\theta^*_\ell = \prod_{k=\ell+1}^L \Phi^*_k$, we obtain

$$\theta^*_\ell \ge \left(1-\frac{\epsilon}{n}\right)^{L-\ell}\cdot \Pr_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset]\ . \tag{EC.9}$$

In addition, for every $\ell\in[1,L]$, by applying Theorem EC.1 to the collection of items contained in $\mathcal{C}_\ell$ with $\Phi=\tilde{\Phi}^*_\ell$, we infer from inequality (EC.8) that

$$\mathbb{E}_{S\sim B(\ell,\Phi^*_\ell)}\left[\mathcal{R}\left(S\right)\right] \ge (1-\epsilon)\cdot \mathbb{E}_{S\sim\mathcal{A}^*}\left[\mathcal{R}\left(S\cap\mathcal{C}_\ell\right)\right]\ . \tag{EC.10}$$

By combining inequalities (EC.9) and (EC.10), we finally obtain

$$\begin{aligned}
\sum_{\ell=1}^L \theta^*_\ell\cdot\mathbb{E}_{S\sim B(\ell,\Phi^*_\ell)}\left[\mathcal{R}\left(S\right)\right] &\ge (1-\epsilon)\cdot\sum_{\ell=1}^L\left(1-\frac{\epsilon}{n}\right)^{L-\ell}\cdot \Pr_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset]\cdot\mathbb{E}_{S\sim\mathcal{A}^*}\left[\mathcal{R}\left(S\cap\mathcal{C}_\ell\right)\right]\\
&\ge (1-\epsilon)\cdot\left(1-\frac{\epsilon}{n}\right)^{n}\cdot\sum_{\ell=1}^L \Pr_{S\sim\mathcal{A}^*}[C_{S\cap\mathcal{C}_{>\ell}}=\emptyset]\cdot\mathbb{E}_{S\sim\mathcal{A}^*}\left[\mathcal{R}\left(S\cap\mathcal{C}_\ell\right)\right]\\
&\ge (1-3\epsilon)\cdot\tilde{\mathcal{R}}\left(\mathcal{A}^*\right)\\
&= (1-3\epsilon)\cdot F\left(L,1\right)\ .
\end{aligned}$$

where the second inequality holds since $L\le n$, by observing that the number of non-empty clusters, as defined in Section EC.2.2, is of at most $n$. □

## Appendix EC.3: Computational Experiment

### EC.3.1. The Coupon Display PTAS

In what follows, we describe how we modify our PTAS to cater to the cardinality constrained instances of Section 5 . First, we compute $B_{\text{likely}}$ via an enumeration approach similar to Step 2(b) of our approximation scheme with respect to the likely items. Next, we re-define the unlikely items to be those with consideration probabilities $\lambda_i\in[\epsilon/n,\epsilon]$. The items with consideration probability $\lambda_i<\epsilon/n$ are termed "rare". For the unlikely items, we first guess the quantity $|S^*\cap\mathcal{U}_q|$, which represents the number of unlikely items from each item class added by the optimal assortment. Then, we carry out Step 2(b) of our approximation, where $\texttt{MinKnapsack}(\{(\hat{c}_q,\hat{\alpha}_q)\}_{\mathcal{Q}})$ is exactly as described in Section 3.2, with the addition of cardinality constraints ensuring that $|S^*\cap\mathcal{U}_q|$ items are added from item class $q\in[Q_{\min},Q_{\max}]$. This step yields the decomposable random assortment $B_{\text{unlike}}$. These updated set of steps for likely and unlikely items are described next in more detail.

*Guessing procedure and running time.* We describe our guessing procedure for $B_{\text{likely}}$ and analyze the corresponding running time. By leveraging the revenue-ordered-by-class property established in Section 3.3, we can guess $B_{\text{likely}}$ by simply enumerating over all ways in which the $C = 6$ products can divided among the item classes $(p, q) \in [P_{\max}] \times \mathcal{Q}$, of which there are at most

$$O\left(6^{(Q_{\max}-Q_{\min})\cdot(P_{\max}-P_{\text{rare}})}\right) = O\left(6^{O((1/\epsilon^2)\cdot\log\frac{w_{\max}}{w_{\min}}\cdot\log\frac{n}{\epsilon})}\right) .$$

To ensure that the ratio $\frac{w_{\max}}{w_{\min}}$ is polynomial in $n$ and $\frac{1}{\epsilon}$, one can exploit dynamic programming ideas similar to those presented in Appendix EC.2. All in all, this approach reduces the overall running time of this step to $O(n^{(\frac{1}{\epsilon})^{O(1)}})$. This guessing scheme is akin to carrying Step 2a of our approach with $\epsilon = 0$, and hence for all future analysis, we assume that both versions of the PTAS return the same set of likely items.

For the unlikely items, a similar analysis shows that we can guess $|S^* \cap \mathcal{U}_q|$ in polynomial time as well. After this guessing step, we carry out Step 2a exactly as detailed in Section 3.2, meaning our overall approach for the unlikely items also runs in polynomial time. It is important to note that in a setting with a cardinality constraint, our original PTAS must also employ this initial guessing step and hence we may again assume that both version of the PTAS return the same set of unlikely items.

Finally, we construct our choice of rare items $B_{\text{rare}}$ by myopically selecting the rare items of largest $\rho_i$-quantities until the cardinality constraint is met, where $\rho_i = \lambda_i r_i w_i$. For $i \in B_{\text{rare}} \cup B_{\text{unlike}}$, we let $B_i$ be a Bernoulli random variable with success probability $1-\epsilon$ (all items in $B_{\text{likely}}$ continue to be selected with probabaility 1). Ultimately, the algorithm returns the assortment $B_{\text{likely}} \cup B_{\text{unlike}} \cup B_{\text{rare}}$, which is shown to be $(1 - O(\epsilon))$-optimal in our next claim.

Claim EC.6. *$\mathcal{R}(B_{\text{likely}} \cup B_{\text{unlike}} \cup B_{\text{rare}}) \geq (1 - O(\epsilon)) \cdot \mathcal{R}(S^*)$.*

*Proof of Claim EC.6.* Given the discussion above, all that is required to prove the claim is to analyze the rare items selected under the two version of the PTAS. For this purpose, let $\hat{B}_{\text{rare}}$ and $B_{\text{rare}}$ denote the set of rare items selected under our original and modified PTAS respectively. To help clarify, $\hat{B}_{\text{rare}}$ is the subset of items returned from Step 2b that satisfy $\lambda_i < \epsilon/n$. Furthermore, let $B_{\text{other}} = B_{\text{likely}} \cup B_{\text{unlike}}$ denote their shared choices for the likely and unlikely items (with $\lambda_i \in [\epsilon/n, \epsilon]$), and let $\hat{B} = B_{\text{other}} \cup \hat{B}_{\text{rare}}$ and $B = B_{\text{other}} \cup B_{\text{rare}}$. To establish the claim, we first show that

$$\sum_{i \in B_{\text{rare}}} r_i \cdot \mathbb{E}_{S \sim B}\left[\pi(i, B_{\text{other}} \cup B_{\text{rare}})\right] \geq (1 - 2\epsilon) \cdot \sum_{i \in \hat{B}_{\text{rare}}} r_i \cdot \mathbb{E}_{S \sim \hat{B}}\left[\pi(i, B_{\text{other}} \cup \hat{B}_{\text{rare}})\right], \tag{EC.11}$$

and then we show

$$\sum_{i \in B_{\text{other}}} r_i \cdot \mathbb{E}_{S \sim B}\left[\pi(i, B_{\text{other}} \cup B_{\text{rare}})\right] \geq (1 - 2\epsilon) \cdot \sum_{i \in \hat{B}_{\text{other}}} r_i \cdot \mathbb{E}_{S \sim \hat{B}}\left[\pi(i, B_{\text{other}} \cup \hat{B}_{\text{rare}})\right]. \tag{EC.12}$$

Combining (EC.11) and (EC.12) with Theorem 4 yields the desired claim.

To prove (EC.11), observe that

$$\begin{aligned}\sum_{i\in B_{\text{rare}}} r_i\cdot\mathbb{E}_{S\sim B}\left[\pi(i,B_{\text{other}}\cup B_{\text{rare}})\right] &\geq \sum_{i\in B_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})+w(C_{B_{\text{rare}}})}\right]\\ &\geq \Pr\left[C_{B_{\text{rare}}}=\emptyset\right]\cdot\sum_{i\in B_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right]\\ &\geq \left(1-\frac{\epsilon}{n}\right)^n\cdot\sum_{i\in B_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right]\\ &\geq (1-2\epsilon)\cdot\sum_{i\in B_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right].\end{aligned}$$

Recalling that $\rho_i=\lambda_i r_i w_i$, we get that

$$\begin{aligned}\sum_{i\in B_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right] &= \sum_{i\in B_{\text{rare}}}\rho_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right]\\ &\geq \sum_{i\in \hat{B}_{\text{rare}}}\rho_i\cdot\mathbb{E}_{S\sim \hat{B}}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right]\\ &= \sum_{i\in \hat{B}_{\text{rare}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim \hat{B}}\left[\frac{1}{1+w_i+w(C_{B_{\text{other}}})}\right]\\ &\geq \sum_{i\in \hat{B}_{\text{rare}}} r_i\cdot\mathbb{E}_{S\sim \hat{B}}\left[\pi(i,B_{\text{other}}\cup\hat{B}_{\text{rare}})\right]\end{aligned}$$

where the first inequality holds since $B_{\text{rare}}$ is the assortment of rare items that maximizes the quantity $\sum_{i\in S}\rho_i$, while the second inequality is due to $B_{\text{other}}\subseteq(B_{\text{other}}\cup\hat{B}_{\text{rare}})\setminus\{i\}$. Inequality (EC.11) immediately follows from the above sequence of inequalities.

We conclude by proving (EC.12), which is the simple case. To this end, we note that

$$\begin{aligned}\sum_{i\in B_{\text{other}}} r_i\cdot\mathbb{E}_{S\sim B}\left[\pi(i,B_{\text{other}}\cup B_{\text{rare}})\right] &\geq \Pr\left[C_{B_{\text{rare}}}=\emptyset\right]\cdot\sum_{i\in B_{\text{other}}}\lambda_i r_i w_i\cdot\mathbb{E}_{S\sim B}\left[\frac{1}{1+w_i+w(C^{-i}_{B_{\text{other}}})}\right]\\ &\geq (1-2\epsilon)\cdot\sum_{i\in B_{\text{other}}} r_i\cdot\mathbb{E}_{S\sim B}\left[\pi(i,B_{\text{other}})\right]\\ &\geq (1-2\epsilon)\cdot\sum_{i\in B_{\text{other}}} r_i\cdot\mathbb{E}_{S\sim \hat{B}}\left[\pi(i,B_{\text{other}})\right]\\ &\geq (1-2\epsilon)\cdot\sum_{i\in \hat{B}_{\text{other}}} r_i\cdot\mathbb{E}_{S\sim \hat{B}}\left[\pi(i,B_{\text{other}}\cup\hat{B}_{\text{rare}})\right],\end{aligned}$$

where the last inequality holds since $B_{\text{other}}\subseteq B_{\text{other}}\cup\hat{B}_{\text{rare}}$. □

### EC.3.2. SAA Heuristic

Using classic linearization tricks, it is possible to reformulate the nonlinear integer program (3) as the following linear integer program:

$$
\begin{aligned}
\max \quad & \sum_{i\in[n]}\sum_{k\in[K]} \lambda_i r_i w_i z_{i,i}^k && \text{(SAA-IP)}\\
\text{s.t.} \quad & x_i \le M z_{i,j}^k, && \forall k\in[K], i,j\in[n]\\
& z_{i,j}^k \le y_0^k, && \forall k\in[K], i,j\in[n]\\
& y_0^k - z_{i,j}^k \le M(1-x_i), && \forall k\in[K], i,j\in[n]\\
& y_0^k + \sum_{j\in C^k_{[n]\setminus\{i\}}} w_j z_{i,j}^k = 1, && \forall k\in[K], i\in[n]\\
& \sum_{i=1}^n x_i = 6 &&\\
& x_i\in\{0,1\}, y_0^k, z_{i,j}^k \ge 0, &&
\end{aligned}
$$

where $M$ is a large constant. The decision variables $y_0^k$ and $z_{i,j}^k$ both capture the no-purchase probability for sample $k$ of the random consideration set $C_{[n]\setminus\{i\}}$.

## Appendix EC.4: Estimation Case Study

### EC.4.1. Catboost Implementation for Estimating Click Probabilities

The following function takes as input a pandas dataframe sales_data, which has a row for each product displayed to each each customer and whose columns give the various feature values. Furthermore, there is an additional binary column (is_click) indicating whether the given product was clicked or not. The second input to this function is feature_list; a list of column names corresponding to the features that will be used to fit the catboost model. Note that in the code below, we optimize over two hyperparameters, depth and l2_leaf_reg, whose exact nature is formalized next. These two hyperparameters were selected after extensive trial and error to find the hyperparameters that had the greatest effect on prediction accuracy. The depth parameter controls the maximum depth allowed for any of the fitted trees and the l2_leaf_reg parameters is the coefficient of the L2 regularization term of the cost function that is minimized.

```
def Get_Click_Probs_Catboost(sales_data, feature_list):

    # Get feature values (X) and response (y) for each customer.
    # Convert both to numpy arrays for catboost
    X = np.array(sales_data.loc[:, feature_list])
    y = np.array(sales_data.is_click)
```

```
# Hyperparameter tuning
params = {'depth': [3,4,5], 'l2_leaf_reg': [1,4,9]}

# Instantiate a catboost model, and fit with X and y.
# We do 5-fold cross validation (cv=5) to get the best hyperparameters
cb = CatBoostClassifier(logging_level = "Silent")
model = GridSearchCV(cb, params, cv = 5).fit(X, y)

return model
```

### EC.4.2. Additional Estimation Experiment with the Mixed-MNL Model

Figure EC.1 displays the results of an additional set of experiments where fit mixed-MNL models with $G \in \{1,2,3,4,5\}$, where we assess the accuracy of each fitted model based its improvement over a traditional MNL model in terms of out-of-sample log-likelihood. It is important to note that these experiments used a new collection of randomly generated train/test splits, and hence the results presented here will differ slightly from those presented in Section 6.5.

### EC.4.3. The Cost of Model Misspecification

In this section, we evaluate the gap between the MNL and click-based MNL choice models from a decision-making standpoint. Since the click-based MNL model generalizes the standard MNL model, and its fit to historical data is significantly more accurate (see Table 5), we wish to estimate the operational cost of a model misspecification. *Does our generalization of the MNL model generate significantly different assortment recommendations?* We quantity the potential loss of revenue in using the standard MNL model in place of the click-based MNL model, when the latter forms the ground truth. For this purpose, we propose a generative model that mirrors the real-world assortment optimization setting faced by Alibaba whenever customers claim a coupon. For the instances that we generate, we take the viewpoint that true purchasing patterns are governed by a click-based MNL model, and we study the cost of making assortment decisions according to an MNL model.

*Experimental set-up* As explained in Section 6.1, when a customer lands on a seller's page and then clicks on a coupon, she is brought to a coupon sub-page where six products are displayed. Since Alibaba's objective in this setting is to maximize the revenue garnered from each customer who claims a coupon, the problem of choosing the six product displays can be formulated as a

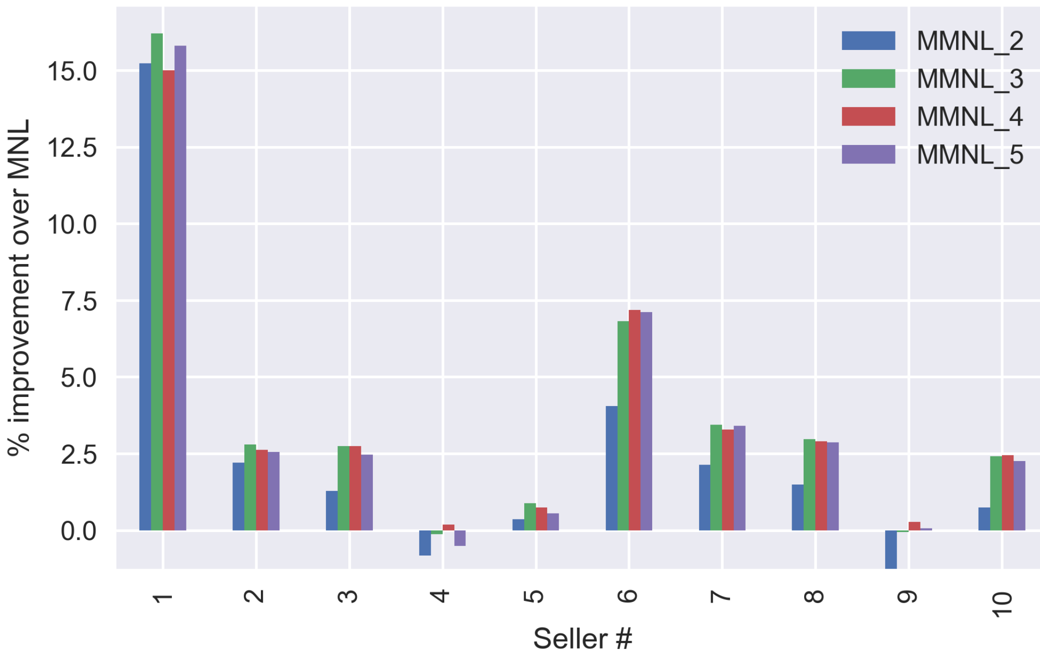


**Figure EC.1 Percentage improvements in out-of-sample log-likelihood of mixed-MNL models with $G \in \{2,3,4,5\}$ over the MNL model.**

cardinality-constrained assortment optimization problem. We focus specifically on replicating the assortment problems faced by Sellers 7 and 10, whose summary statistics are provided in Table 4. For this purpose, we model the collection of products available to each seller as those that have been offered at least once in our data set. The 25-dimensional feature vector characterizing a particular product is randomly generated from the product-specific empirical distributions induced by our historical data set. Combining these feature vectors along with our fitted models (see Sections 6.3 and 6.4), we jointly instantiate a click-based MNL and a standard MNL assortment optimization problems. The optimal assortment associated with the standard MNL model is computed using the linear programming formulation of Gallego et al. (2015) and Sumida et al. (2020). In order to compute assortment recommendations for the click-based MNL model, we employ the updated PTAS detailed in Section 5.2 with $\epsilon \in \{0.05, 0.04, 0.03\}$.

Next, we compare the expected revenues of the assortments recommended by the MNL and click-based MNL models. Assuming that the fitted click-based MNL models prescribe the "ground-truth" purchase probabilities, we compute the revenue loss incurred if Alibaba were to adopt the assortments recommended by the MNL fits instead. The gap between these two recommended assortments can be interpreted as the cost of a model misspecification with respect to these nested families of choice models. That is, this is the loss of revenue from applying a standard MNL choice model in settings where customers follow a more general two-stage choice-making process. Clearly, this performance gap does not realistically measure the extent to which revenue would increase if a platform such as Alibaba utilizes the click-based MNL model vis-a-vis the standard MNL model – this difference can only be measured by conducting a controlled experiment. However, such simulations indicate the degree of dissimilarity between the assortment recommendations generated by these two choice models. Indeed, the gap of predictive performance observed in Section 6.5 does not necessarily imply that the click-based MNL model should be credited with significantly better assortment decisions than the MNL model. At face value, it is entirely plausible that optimal assortment decisions are insensitive to the underlying choice model, given the notorious robustness of the MNL choice model in applications and the similarity with the click-based MNL model.

*Instance generator.* For each of the three sellers, we randomly generate 100 instances, jointly describing the click-based MNL and standard MNL assortment optimization problems, using the following procedure. First, we generate a vector of 25 feature values $(x_i^1, \ldots, x_i^{25})$ for each product $i$ offered by the particular seller, using product-specific empirical distributions derived from the sales data. Specifically, using the notation of Section 6, each feature value $x_i^j$ is sampled uniformly at random from the historical data $\{X_{it}^j : t \in [\tau]\}$, where $X_{it}^j$ is the $j$-th coordinate of the feature vector $X_{it}$ describing item $i \in S_t$ offered to customer $t$. Next, we utilize the resulting product features and the estimates for the model parameters $\beta$ and the function $g(\cdot)$ obtained by solving the MLE problem described in Section 6.3 to specify the MNL weights $w_i$ and the consideration probabilities $\lambda_i$ for the click-based MNL model. Similarly, we utilize these product features and the MLE estimates obtained by solving problem (MLE MMNL) to specify the preference weights of the standard MNL model. Finally, since precisely six products should be chosen, we add a cardinality constraint to the assortment problem, that forces any feasible assortment to contain exactly six products.

*Results.* Let $S^{\mathrm{PTAS}}$ be the assortment returned by algorithm PTAS, and let $S^{\mathrm{MNL}}$ be the assortment recommended by the fitted MNL model. In Table 3, we report the average optimality gap of the assortment $S^{\mathrm{PTAS}}$ over $S^{\mathrm{MNL}}$ for each seller over the 100 generated instances. More formally, we define the optimality gap of assortment $S \in \{S^{\mathrm{PTAS}}, S^{\mathrm{MNL}}\}$ as $\frac{\mathcal{R}(S^*)-\mathcal{R}(S)}{\mathcal{R}(S^*)}$, where the optimal assortment $S^*$ is computed via complete enumeration over all feasible assortments. It is important to note that $\mathcal{R}(S)$ stands for the expected revenue of assortment $S$ under the fitted click-based MNL model, and thus, we compute the optimality gap assuming that the click-based MNL model is the ground-truth.

Table EC.1 reveals that the optimality gaps of the assortments recommended by the click-based MNL model are generally between 6-9% smaller than those observed under the assortments recommended by the MNL fits. This suggests that the two choice models lead to significantly distinct assortment decisions. Interestingly, the optimality gaps observed for these Alibaba-inspired

| | | | Avg. % Opt. Gap | |
|---|---|---|---|---|
| Seller | $n$ | $\epsilon$ | PTAS | MNL |
| 7 | 49 | 0.03 | 0.62 | 7.60 |
| | | 0.04 | 1.04 | 7.60 |
| | | 0.05 | 0.90 | 7.60 |
| 10 | 38 | 0.03 | 0.039 | 9.34 |
| | | 0.04 | 0.061 | 9.34 |
| | | 0.05 | 0.016 | 9.34 |

**Table EC.1 Percent optimality gap of the assortments $S^{\text{PTAS}}$ and $S^{\text{MNL}}$.**

instances are slightly larger than those observed in Section 5, where the performance of PTAS was tested on synthetic instances.