Assortment Optimization under the Multi-Purchase Multinomial Logit Choice Model

Yicheng Bai
School of Operations Research and Information Engineering, Cornell Tech, New York, New York 10044. yb279@.cornell.edu

Jacob Feldman
Olin Business School, Washington University, 1 Brookings Dr., St. Louis, Missouri 63130, USA. jbfeldman@wustl.edu

Danny Segev
Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. segevdanny@tauex.tau.ac.il

Huseyin Topaloglu
School of Operations Research and Information Engineering, Cornell Tech, New York, New York 10044. topaloglu@orie.cornell.edu

Laura Wagner
Universidade Católica Portuguesa, Católica Lisbon School of Business and Economics, Católica Lisbon Research Unit in Business and Economics, Portugal, Palma de Cima, 1649-023 Lisboa, Portugal. lwagner@ucp.pt

In this paper, we introduce the Multi-Purchase Multinomial Logit choice model, which extends the random utility maximization framework of the classical Multinomial Logit model to a multiple-purchase setting. In this model, customers sample random utilities for each offered product as in the Multinomial Logit model. However, rather than focusing on a single product, they concurrently sample a "budget" parameter M, which indicates the maximum number of products that the customer is willing to purchase. Subsequently, the M highest utility products are purchased, out of those whose utilities exceed that of the no-purchase option. When fewer than M products satisfy the latter condition, only these products will be purchased.

Our primary contribution resides in proposing the first multi-purchase choice model that can be fully operationalized. Specifically, we provide a recursive procedure to compute the choice probabilities in this model, which in turn provides a framework to study its resulting assortment problem, where the goal is to select a subset of products to make available for purchase so as to maximize expected revenue. Our main algorithmic results consist of two distinct polynomial time approximation schemes (PTAS); the first, and simpler of the two, caters to a setting where each customer may buy only a constant number of products, whereas the second more nuanced algorithm applies to our multi-purchase model in its general form. Additionally, we study the revenue-potential of making assortment decisions that account for multi-purchase behavior in comparison to those that overlook this phenomenon. In particular, we relate both the structure and revenue performance of the optimal assortment under a traditional single-purchase model to that of the optimal assortment in the multi-purchase setting. Finally, we complement our theoretical work with an extensive set of computational experiments, where the efficacy of our proposed PTAS is tested against natural heuristics. Ultimately, we find that our approximation scheme outperforms these approaches by 1-5% on average.

1. Introduction

The standard assumption underlying the vast majority of classical choice models is that each arriving customer purchases at most a single product. For example, the ever-popular Multinomial Logit (MNL) model (Luce 1959, McFadden 1974, Plackett 1975) and the Nested Logit model (Williams 1977, McFadden 1978) both make this fundamental assumption, as do more recent choice models such as the Markov chain (Blanchet et al. 2016, Feldman and Topaloglu 2017) and the non-parametric ranking-based (Mahajan and Van Ryzin 2001, Farias et al. 2013, Honhon et al. 2012) models. Consequently, the massive body of analytical papers that study operational questions under these models crucially rely on the assumption that customer behavior is exclusively restricted to single-product purchases. In reality, however, there is an abundance of concrete applications where customers often make multiple simultaneous purchases from a single category of substitutable products.

To demonstrate this phenomenon, we observe this trend among the customers of a leading flash-sales e-retailer, who graciously agreed to share with us a small portion of their sales data. This company, whose identity cannot be revealed for confidentiality reasons, runs what they refer to as "campaigns", during which well-known brands are offered from a particular product category at a steep discount for a short period of time. We were given access to sales data from numerous campaigns across multiple product categories. In Table 1, we summarize the frequency with which customers are observed to be making multiple simultaneous purchases from a single product category. For example, 82% of customers who made a purchase from the eyewear category limited themselves to a single product. On the other hand, this percentage drops to about 40% for the underwear category. Another interesting insight from Table 1 is that customers rarely purchased multiple versions or copies of the same product. In fact, out of those who purchased two products in the same campaign, the vast majority – over 89% – purchased two different products.

The above example elucidates the real-life necessity of choice models that can meaningfully capture customers making multiple purchases from a single product category. To the best of our knowledge, the only existing models that fall into this framework are those proposed by Kim et al. (2002), Ferreira and Goh (2021), Fox et al. (2018), and Tulabandhula et al. (2020), which are thoroughly discussed in Section 1.2. In short, these four papers propose distinct random-utility-based choice models that capture customers making multiple purchases. Unfortunately, the structural assumptions underlying each of these models lead to complex expressions for their inherent choice probabilities, and in turn, it is currently unknown whether one can develop tractable algorithms with provable performance guarantees for assortment or pricing problems under these models. In fact, among these four papers, only Tulabandhula et al. (2020) make an explicit effort to tackle the corresponding assortment problem; however, their proposed algorithms are heuristic in nature

Table 1 Distribution of Multi-Purchase Events

Sector	$k=1$	$k=2$	$k=3$	$k\geq 4$	Mean Total purchase	Number of campaigns
Underwear					511	2
Orders of size k	40%	40%	14%	6%		
Orders containing k different products	100%	99%	99%	92%		
Dresses					228	2
Orders of size k	56%	23%	15%	6%		
Orders containing k different products	100%	100%	100%	97%		
Denim & Casualwear					280	9
Orders of size k	66%	21%	8%	5%		
Orders containing k different products	100%	90%	65%	62%		
Shoes					674	32
Orders of size k	66%	24%	6%	4%		
Orders containing k different products	100%	95%	76%	73%		
Bags					122	11
Orders of size k	76%	17%	5%	2%		
Orders containing k different products	100%	93%	80%	72%		
Jewellery & Watches					258	27
Orders of size k	81%	14%	3%	2%		
Orders containing k different products	100%	89%	75%	39%		
Eyewear					414	43
Orders of size k	82%	15%	2%	1%		
Orders containing k different products	100%	93%	85%	79%		

Note: The percentage of orders containing k different products (opposed to orders containing the same kind) are calculated from the number of orders of size k. To address the concern that these products could be complements, campaigns with more than one product description were removed.

and do not come with any theoretical guarantees. Motivated by this apparent void in the revenue management literature, we set-out to address the following research questions:

1. Can we develop a choice model, based on the well-established framework of random utility maximization (RUM), that captures customers making multiple simultaneous purchases? Moreover, can we derive this model within the utility-based guidelines of the MNL model, thus extending the practical scope of this perennial choice model?

2. Given such a multi-purchase model, can we develop tractable algorithms with provable near-optimal guarantees for its corresponding assortment optimization problem, thus overcoming the previously mentioned operational roadblocks that have limited the practical appeal of existing multi-purchase models?

3. Does the structure of optimal assortments change in a meaningful way as we move from single- to multi-purchase settings? Do such differences lead to high-level managerial insights that could guide product offering decisions?

1.1. Modeling approach and technical contributions

In this paper, we introduce the first fully operational extension of the classical MNL model to a setting where customers can simultaneously purchase multiple substitutable products. We aptly refer to the resulting model as the Multi-Purchase Multinomial Logit (MP-MNL) choice model. After formalizing the probabilistic structure of this model, we study its corresponding assortment optimization problem, where the goal is to select the revenue-maximizing subset of products to offer to each arriving customer. We first prove various structural properties related to the make-up of an optimal assortment. Beyond leading to useful managerial insights, these structural properties directly yield elegant algorithms for identifying assortments with provable performance guarantees. While these preliminary algorithms possess the desirable qualities of being simple and easily implementable, their recommended assortments could be rather far from optimal in the worst case. For this reason, as our cornerstone algorithmic results, we provide two distinct polynomial time approximation schemes (PTAS) that compete against the optimal revenue to within any level of precision. The first, and simpler of the two, considers an intermediate setting in which each customer is assumed to purchase a constant number of products, while the second more nuanced PTAS applies to the MP-MNL model in its most general form. We proceed by summarizing these contributions in greater detail.

A choice model for multiple purchases. In what follows, we present a high-level overview of the MP-MNL model to set the stage for its formal derivation in Section 2. In a nutshell, the MP-MNL model is perhaps the most natural extension of the traditional MNL model to a multi-purchase setting. Under this model, there are two sources of randomness that guide the purchasing process of each customer. First, each arriving customer samples a random "budget" M, which represents the maximum number of products the customer would ever be willing to buy from the category in question. One natural interpretation of M is as a spending budget: A customer arriving to GAP, for example, has a budget of \$30 to spend on t-shirts, and since each t-shirt costs approximately \$10, this equates to sampling $M = 3$. Next, mirroring the RUM framework of the standard MNL model, the customer associates Gumbel-distributed random utilities with each offered product. Finally, the customer purchases the M highest utility products, out of those whose utilities exceed that of the no-purchase option. When fewer than M products have utilities that satisfy the latter condition, only these products will be purchased.

Interestingly, the marketing literature offers a great deal of evidence in support of the MP-MNL model dynamics, and in particular, its core feature where customers are "variety-seeking", i.e., prefer adding a new item to their basket rather than doubling-up on an existing item. Indeed, Kahn (1995) presents an exhaustive review of the "explosion of research in the marketing literature" on the topic of variety-seeking customers. Of particular relevance is the line of work that aims to characterize and explain variety-seeking purchasing behavior when customers anticipate the temporal separation between purchase and consumption. In this case, there is overwhelming experimental (Simonson 1990), empirical (Harlam and Lodish 1995) and theoretical (Walsh 1995) evidence suggesting that customers prefer, and are generally better off from utility perspectives, purchasing baskets of distinct items rather than multiple units of the same item, in an effort to gain flexibility and hedge against uncertain future preferences.

First and foremost, we show in Section 2.2 that the MP-MNL model admits an intuitive recursive expression for its choice probabilities. That said, for the most general form of the MP-MNL model, a direct implementation of this recursion is not computationally efficient, which leads us to develop a sampling-based alternative, exploited as a black-box within the approximation scheme developed in Section 5. Furthermore, as a sanity check, in the specialized setting of $M = 1$ purchases, we observe that the choice probabilities under the MP-MNL model exactly match those of the standard MNL model. On the other hand, when M exceeds the number of offered products, the choice probabilities become assortment-independent. In other words, product cannibalization/substitution patterns vanish since each offered product will be purchased when its utility exceeds that of the no-purchase option. Hence, the assortment problem is easily solvable in either of these two corner cases. However, as we explain in the sequel, there are many algorithmic hurdles that need to be overcome when the maximum number of purchases M is arbitrarily distributed.

The assortment optimization problem. In Section 2.3, we formalize the assortment optimization problem under the MP-MNL model, where the goal is to select a subset of products to offer so as to maximize expected revenue. Unfortunately, from an intractability perspective, despite our best efforts, we have not been successful at establishing hardness results for the unconstrained assortment problem, whose computational complexity remains an intriguing open question. That said, we show that assortment optimization under the MP-MNL model is $\Omega(n^{1-\epsilon})$-hard to approximate under Totally Unimodular (TU) constraint structures, where n stands for the number of available products. Additionally, we prove that this problem remains NP-hard even for extremely simple versions, where M takes only two values. These hardness results establish a strong separation between the MP-MNL model and the traditional MNL model, where TU-constrained assortment optimization admits a linear programming-based polynomial time algorithm (Sumida et al. 2021).

Revenue-ordered assortments. In Section 3, we consider the efficacy of offering revenue-ordered assortments, consisting of the highest revenue products up to a given threshold. We begin by considering an optimal assortment under the false assumption that customers are willing to purchase at most one product, which reduces the MP-MNL model to a standard MNL model. The optimal assortment in the latter setting is known to be revenue-ordered (Talluri and Van Ryzin 2004). We first show that this assortment is necessarily contained in the true optimal assortment of the multi-purchase setting, indicating that, as we move from a single- to a multi-purchase setting, larger and more diverse assortments are needed to extract optimal revenues. Additionally, this structural property leads to the natural question: Can optimal or near-optimal revenues be attained by revenue-ordered assortments? On the negative side, we construct an exemplary instance of our MP-MNL assortment problem in which revenue-ordered assortments garner at most half of the optimal expected revenue. On the positive side, we derive two distinct performance guarantees.

First, we show that the optimal (single-purchase) MNL-based assortment garners at least a $\frac{1}{\mathbb{E}[M]}$-fraction of the optimal expected revenue. Next, we show that the most profitable revenue-ordered assortment under the MP-MNL model earns an $\Omega(\frac{1}{\log n})$-fraction of the optimal expected revenue.

Approximation schemes. Our primary algorithmic results come in Sections 4 and 5, where we present two distinct approximation schemes for the assortment optimization problem, varying in sophistication and applicability. More formally, for any accuracy level $\epsilon > 0$, we provide two algorithmic approaches for computing an assortment whose expected revenue is within a $(1 - \epsilon)$-factor of the optimal expected revenue. In Section 4, our approximation scheme pertains to a simplified setting, where each arriving customer is willing to purchase at most $m_{\max} = O(1)$ products. Under this assumption, we observe that our recursive approach for computing the choice probabilities runs in polynomial time, and hence our algorithmic focus can turn to directly tackling the assortment problem. Along this line, our key algorithmic insight in this limited purchase setting is that, with knowledge of the $O(\frac{m_{\max}^2}{\epsilon})$ products with the highest MNL-based preference weight offered in the optimal assortment, the problem can be approximately reduced to a knapsack-like problem. Next, in Section 5, we present a more nuanced PTAS for the assortment problem in its utmost generality, where the number of products that each customer is willing to purchase is arbitrarily distributed. The high-level idea of this PTAS is to decompose the original problem into a collection of sub-problems based on partitioning the products by their respective weights. We then develop a PTAS for each of the resulting sub-problems, whose solutions are "stitched" together via a carefully crafted dynamic program that accounts for the subtle interactions between these sub-problems. Overall, our PTAS for the general setting can be viewed as an intriguing theoretical finding. While this approach may not be directly implementable for large-scale instances, the algorithmic insights we uncover in its development certainly have the potential to lay the groundwork for future research, aimed at tackling this general setting in a more efficient way.

Computational experiments. In Appendices G and H, we present two distinct sets of computational experiments. The first is aimed at establishing that traditional single-purchase models are generally inadequate in settings where multi-purchase behavior is prevalent. For this purpose, we generate an extensive collection of sales data, assuming that arriving customers make purchasing decisions according to an MP-MNL model. We then fit the ever-popular mixed-MNL model to the resulting data set via maximum likelihood estimation, and assess its accuracy through the ability to recapture estimates of the underlying choice probabilities as dictated by the ground truth MP-MNL model. More specifically, for each mixed-MNL fit, we compute its relative absolute percent error in predicted choice probability averaged over all assortments, and find that this metric exceeds 23% for all test cases considered. We further show that these prediction errors have downstream consequences, leading to highly sub-optimal assortment recommendations that leave up to 20% of revenue on the table.

Our second set of experiments is devoted to measuring the efficacy of our approximation scheme under limited purchases against two heuristic approaches on a wide variety of randomly generated instances, with up to 60 products and with a cardinality constraint, which encodes the notion that exactly C products must be offered. As such, we discuss how the PTAS presented in Section 4 can easily be adapted to this cardinality constrained setting. One of the tested heuristics assumes that customers are willing to purchase at most one product, meaning that this heuristic solves an assortment problem where customer choice is governed by a traditional MNL model. This comparison allows us to gauge the potential revenue-loss that ensues by ignoring multi-purchase behavior. Overall, we find that our PTAS consistently outperforms the other two approaches by approximately 1-5% across all test cases. We find that this revenue improvement can be potentially attributed to fundamental differences in the make-up of the assortments recommended by our PTAS. Specifically, the latter generally computes assortments that contain a few high weight products that are missing from the MNL-based assortment recommendations, ultimately leading to the superior performance of the former assortments for customers who are willing to purchase a relatively large number of products.

1.2. Related literature

In what follows, we begin by briefly reviewing past work that considers assortment optimization problems under logit-based choice models. We follow up with a survey of existing results concerning choice models in which customers may simultaneously purchase multiple products.

Logit-based models. The seminal work of Talluri and Van Ryzin (2004) was the first to consider the assortment optimization problem under MNL preferences. They show that an optimal assortment in this setting is revenue-ordered, meaning that it consists of some subset of the highest revenue products. Subsequently, Rusmevichientong et al. (2010) considered the cardinality-constrained variant of this problem, in which there is an upper bound on the total number of products that the retailer can feasibly offer. They provide a purely combinatorial polynomial time algorithm for this constrained variant. Finally, Sumida et al. (2021) study an MNL-based assortment problem in which the offered assortment is restricted by any set of constraints that can be encoded via a Totally Unimodular (TU) matrix. Numerous variants of the basic assortment optimization problem fall under this framework. In this TU-constrained setting, Sumida et al. (2021) show that the assortment problem can be efficiently solved via a carefully crafted linear program.

When customer choice is governed by the Nested Logit model, Davis et al. (2014) show that the unconstrained assortment problem can be solved in polynomial time when the no-purchase option is not available within nests. If, on the other hand, customers can choose the no-purchase option after having selected a nest, the authors show that the problem becomes NP-Hard, and provide a 2-approximation for this case. Gallego and Topaloglu (2014) and Feldman and Topaloglu (2015) provide various constant-factor approximations for cardinality and space constrained variants of the assortment problem under the Nested Logit model, while the very recent work of Segev (2020) proposed approximation schemes in this context. Moreover, Li et al. (2015) devise an exact polynomial time algorithm for the unconstrained assortment problem under the d-Level Nested Logit model, where the nesting structure of products is encoded as a d-level binary tree, instead of the traditional 2-level tree, as in the standard Nested Logit model.

The final logit-based model for which we summarize past work is the Mixed-Logit model, first studied in the assortment optimization context by Bront et al. (2009). In this setting, Rusmevichientong et al. (2014) prove that the unconstrained assortment problem is NP-Hard even when there are only two customer classes. They also establish several conditions on the parameter space under which revenue-ordered assortments are optimal. Concurrently, Désir et al. (2022) show that the unconstrained assortment problem under a Mixed-Logit model is in fact NP-Hard to approximate within a factor of $O(n^{1-\epsilon})$, for any fixed $\epsilon > 0$. In addition, they provide an approximation scheme for both the unconstrained and constrained problem, whose running time scales exponentially in the number of customer classes. Bront et al. (2009) provide an exact integer-programming approach to tackle the assortment problem under a Mixed-Logit model, showing that this approach scales to reasonably sized instances with hundreds of products and customer segments.

Multi-purchase models. Next, we focus on summarizing earlier work on choice models that capture customers making multiple purchases. Kim et al. (2002) and Ferreira and Goh (2021) both propose RUM-based models for multi-purchase settings based on the assumption that customers receive diminishing gains in marginal utility when they purchase multiple products. More specifically, Kim et al. (2002) propose a model in which customers sample log-normal utilities, and then fill their basket with items so as to maximize their total utility while staying within a given budget. The authors show how to estimate the resulting model parameters from sales data; however, pricing or assortment problems are not considered in this paper. In the model developed by Ferreira and Goh (2021), after sampling the idiosyncratic portion of their utility, customers are assumed to purchase all items that yield a net positive utility. Under this model, the authors weigh the benefits of two assortment strategies: either displaying all products as a single assortment, or sequentially revealing products one after the other. Consequently, their main focus is on uncovering high-level insights regarding the impact and profitability of frequent assortment rotations, rather than on developing tractable assortment policies with provable performance guarantees.

Fox et al. (2018) propose a choice model that captures the dynamics of customers forming so-called "n-packs", which are collections of n substitutable products to be purchased and then consumed at a later point in time. A 4-pack of yogurt, for example, could consist of two strawberry and two blueberry flavored yogurts. For a given n-pack, the authors derive the sequential consumption process that will yield the highest utility payoff. Given this analysis, they show how to determine an optimal n-pack, i.e., the collection of n products that, when consumed in the optimal order, will yield the highest expected utility. That said, the focus of this paper is on characterizing a consumer's optimal buying process, rather than on characterizing optimal operational decisions from a retailer's perspective. Finally, Tulabandhula et al. (2020) propose an extension of the MNL model in which customers purchase bundles of products. For the resulting assortment problem, they prove NP-Hardness, even when customers purchase at most two products, and propose various heuristics that do not come with theoretical performance guarantees.

On top of this literature, it is worth highlighting the rank-ordered Logit model developed by Punj and Staelin (1978) and Beggs et al. (1981), whose underlying structure resembles the MP-MNL model in several aspects. Specifically, the rank-ordered Logit model assumes the same random utility specification as that of the MP-MNL. However, it makes no attempts to explicitly characterize purchasing behavior, and instead, specifies the probability of observing any particular ranking of the available alternatives' utilities. This model has been utilized in many applications, including voter preferences (Koop and Poirier 1994), automobile demand (Dagsvik and Liu 2006), and graduate school rankings (Mark et al. 2004), to mention a few. The MP-MNL model extends this framework to a retailing context by assuming that the different alternatives are products made available for purchase, and that customers will purchase the M highest utility products whose utilities exceed that of the no-purchase option.

2. The Multi-Purchase MNL Choice Model

In this section, we formally describe the MP-MNL choice model. To this end, we first provide a detailed description of its choice dynamics, which are then translated into a recursive expression for computing the underlying choice probabilities. In addition, we formulate the assortment optimization problem under the MP-MNL model and shed light on its computational complexity.

2.1. Model specification

For the remainder of this paper, we consider a retailer who has access to n products, which will be referred to as $1, \dots, n$, alongside the ever-present no-purchase option, which will be designated as a 0-indexed product, for convenience of notation. Under the Multi-Purchase MNL choice model, each customer arrives with the intention of purchasing multiple products. We use the random variable

M to denote the maximum number of products that the arriving customer is willing to purchase, and note that its underlying distribution over $[n]_0 = \{0, 1, \dots, n\}$ is assumed to be a known input.

After sampling their maximum number of purchases, M, customers then sample random utilities for all products. Building upon the RUM framework of the MNL model, we assume that the random utility associated with product $i \in [n]_0$ is given by $U_i = v_i + e_i$, where v_i represents the observable deterministic component of the utility and e_i is a standard Gumbel random variable with location-scale parameters $(0, 1)$. Here, the random variables $e_0, e_1, \dots, e_n$ are independently sampled, and since for all choice-related purposes, random utilities of different products will only be compared, we assume without loss of generality that $v_0 = 0$. Given these settings, an arriving customer with $M = m$, will purchase the m highest utility products whose utilities exceed that of the no-purchase option. When fewer than m products have utilities that satisfy the latter condition, only these products will be purchased.

To formally define the choice probabilities that result from the dynamics described above, we introduce some additional notation. For any product i in an assortment $S \subseteq [n]$, let $\text{rank}(i, S)$ be the random rank of product i's utility in relation to the utilities of all products in S, namely, $\text{rank}(i, S) = |\{j \in S : U_j \geq U_i\}|$. For example, in the event that $U_i = \max_{j \in S} U_j$, then $\text{rank}(i, S) = 1$. In addition, for any $m \in [n]$, let $\pi_m(i, S)$ be the probability that product $i \in S$ is purchased given that $M = m$, i.e., when the arriving customer is willing to purchase up to m products. According to the purchasing dynamics described above, we clearly have

$$\pi_m(i, S) = \Pr\left[[\text{rank}(i, S) \leq m] \wedge [U_i > U_0]\right] \ , \tag{1}$$

since product $i \in S$ is purchased when its rank among S is at most m, and concurrently, it is preferable to the no-purchase option. With this notation in-hand, the choice probability of any product $i \in S$ can be expressed as $\pi(i, S) = \sum_{m \in [n]} \Pr[M = m] \cdot \pi_m(i, S)$.

2.2. Computing the choice probabilities

In what follows, we present a recursive approach to compute the conditional choice probabilities $\pi_m(i, S)$, formally defined in (1). The resulting expression will critically rely on the following well-known properties of Gumbel distributed utilities (Luce 1959, Beggs et al. 1981).

PROPERTY 2.1 **(Top Rank Probability)**. *For any assortment $S \subseteq [n]$ and product $i \in S$,*

$$\Pr\left[U_i = \max_{j \in S} U_j\right] = \frac{e^{v_i}}{\sum_{j \in S} e^{v_j}} \ .$$

PROPERTY 2.2 **(Conditional Independence of Top Ranked Products)**. *Let S_1 and $S_2 = \{j_1, \dots, j_p\}$ be a partition of an assortment S. Then,*

$$\Pr\left[U_{j_1} > \cdots > U_{j_p} \middle| \min_{i \in S_1} U_i > \max_{j \in S_2} U_j\right] = \Pr\left[U_{j_1} > \cdots > U_{j_p}\right] \ .$$

Property 2.1 is simply a reminder of the choice probabilities structure under a standard MNL model. Property 2.2 states that, conditional on $\min_{i\in S_1} U_i > \max_{j\in S_2} U_j$, the relative order between utilities of products in S_2 is independent of the utilities of products in S_1. One important implication of the latter property is that, for any product $i \in S_2$ and integer m,

$$\Pr\left[\mathrm{rank}(i,S) \le m \,\middle|\, \min_{i\in S_1} U_i > \max_{j\in S_2} U_j\right] = \Pr\left[\mathrm{rank}(i, S\setminus S_1) \le m - |S_1|\right] \ .$$

Moving forward, we will use $w_i = e^{v_i}$ as the "preference weight" associated with each product $i \in [n]_0$, where by convention $w(S) = \sum_{j\in S} w_j$ will stand for the total weight of the assortment S. For ease of notation, we will utilize the shorthands $S_{-i} = S \setminus \{i\}$ and $S_{+i} = S \cup \{i\}$.

The recursive relation. Next, we show how to exploit these two properties to derive a recursive expression for the choice probabilities under the MP-MNL model. The following lemma provides a recursive relationship through which $\pi_m(i,S)$ can be computed.

LEMMA 1. *For any assortment $S \subseteq [n]$, product $i \in S$, and $m \in [n]$, we have*

$$\pi_m(i,S) = \frac{w_i}{1+w(S)} + \sum_{j\in S_{-i}} \frac{w_j}{1+w(S)} \cdot \pi_{m-1}(i, S_{-j}) \ .$$

Proof. For the purpose of computing the choice probability $\pi_m(i,S)$, we combine representation (1) along with conditioning on the top ranked product, to obtain

$$\begin{aligned}
\pi_m(i,S) &= \Pr\left[[\mathrm{rank}(i,S)\le m] \wedge [U_i > U_0]\right] \\
&= \sum_{j\in S} \Pr\left[\mathrm{rank}(j, S\cup\{0\}) = 1\right] \cdot \Pr\left[[\mathrm{rank}(i,S)\le m]\wedge[U_i>U_0] \,\middle|\, \mathrm{rank}(j,S\cup\{0\})=1\right] \\
&= \sum_{j\in S} \frac{w_j}{1+w(S)} \cdot \Pr\left[[\mathrm{rank}(i,S)\le m]\wedge[U_i>U_0] \,\middle|\, \mathrm{rank}(j,S\cup\{0\})=1\right] \\
&= \frac{w_i}{1+w(S)} + \sum_{j\in S_{-i}} \frac{w_j}{1+w(S)} \cdot \Pr\left[[\mathrm{rank}(i,S_{-j}) \le m-1] \wedge [U_i > U_0]\right] \\
&= \frac{w_i}{1+w(S)} + \sum_{j\in S_{-i}} \frac{w_j}{1+w(S)} \cdot \pi_{m-1}(i, S_{-j}) \ .
\end{aligned}$$

Here, the third equality holds since $\Pr[\mathrm{rank}(j, S\cup\{0\}) = 1] = \frac{w_j}{1+w(S)}$, by Property 2.1. The fourth equality directly follows from Property 2.2, along with the observation that, when $\mathrm{rank}(i, S\cup\{0\}) = 1$, product i will be purchased with certainty as long as $m > 0$. $\square$

It is not difficult to verify that computing choice probabilities of the form $\pi(i,S) = \sum_{m\in[n]} \Pr[M = m] \cdot \pi_m(i,S)$ via the recursive expression of Lemma 1 requires evaluating $\pi_m(\hat{S})$ for all assortments $\hat{S} \subseteq S$ with $i \in \hat{S}$ and for all $m \le |\hat{S}|$. As such, when the maximum number of purchases M is arbitrarily distributed, the overall running time to compute $\pi(i,S)$ via this recursion is $O(2^n \cdot n^2)$. Additionally, we observe that $\pi_1(i,S) = \frac{w_i}{1+w(S)}$, that is, when the customer is willing to purchase a single product, customers choose according to a standard MNL model. On the other extreme, where $m \geq |S|$, the customer is willing to purchase at least as many products as are offered, and therefore the choice probability in (1) reduces to $\pi_m(i,S) = \Pr[U_i > U_0] = \frac{w_i}{1+w_i}$, reflecting the notion that each offered product will be purchased if its utility exceeds that of the no-purchase option. In this case, customer choice is governed by an independent demand model, where the choice probabilities are assortment-independent. These two corner cases serve as sanity checks for our modeling assumptions.

2.3. The assortment optimization problem

In what follows, we formulate the assortment optimization problem under the MP-MNL model and shed light on its computational complexity. To this end, for each product $i \in [n]$, let $r_i > 0$ denote its exogenously determined revenue. Then, the expected revenue earned from offering the assortment $S \subseteq [n]$ can be expressed as $\mathcal{R}(S) = \sum_{i \in S} r_i \cdot \pi(i,S)$. In the assortment optimization problem, we wish to compute an assortment whose expected revenue is maximized, meaning that this computational question can be formally stated as

$$\max_{S \subseteq [n]} \sum_{i \in S} r_i \cdot \pi(i,S) \ . \tag{2}$$

Finally, it is worth noting that we may assume without loss of generality that $\Pr[M > 0] = 1$, i.e., each customer arrives with the intention of purchasing at least one product. To verify this claim, let $\mathcal{R}_m(S) = \mathbb{E}[\mathcal{R}(S) | M = m]$ be the expected revenue earned, given that customers are willing to purchase up to m products. Clearly, we have $\mathcal{R}(S) = \sum_{m \in [n]} \Pr[M = m] \cdot \mathcal{R}_m(S)$, meaning that the assortment problem (2) remains unchanged when the random number of purchases M is replaced by $[M | M \geq 1]$.

Hardness: open question and results. It is worth pointing out that the concise nature of the assortment problem described above belies its true complexities, which stem from the intricacy of the purchasing dynamics involved and their resulting choice probabilities. However, despite our best efforts, we were unable to prove that unconstrained assortment optimization in this context is NP-Hard. Consequently, whether computing an optimal assortment without any side constraints is intractable or not remains an intriguing open question for future work. Still, in Appendix A, we present two hardness results related to TU-constrained assortment optimization. Surprisingly, when choice is governed by a standard MNL model, Sumida et al. (2021) show that this constrained variant can be solved in polynomial time via linear programming methods. In sharp contrast, for the MP-MNL model in its general form, we show that it is NP-hard to approximate TU-constrained assortment optimization within factor $O(n^{1-\epsilon})$, for any fixed $\epsilon > 0$. This finding strongly separates the inherent computational complexity of our multi-purchase assortment setting from that of the standard MNL model. In addition, we prove that TU-constrained assortment optimization remains NP-Hard, even when the support of M consists of only two values.

3. The Efficacy of Revenue-ordered Assortments

One of the most well-known results in the assortment optimization literature is that of Talluri and Van Ryzin (2004), who showed that the optimal assortment under the standard MNL model is revenue-ordered, meaning that it is made-up of the highest revenue products up to some threshold. Given that the MP-MNL model is a direct relative of the tradition MNL model, it is natural to ask whether revenue-ordered assortments can perform well in our multi-purchase setting. The entirety of this section is devoted to addressing this question.

3.1. Suboptimality of revenue-ordered assortments

As can only be expected, it is not difficult to construct elementary examples where, under the MP-MNL model, no revenue-ordered assortment is optimal. In fact, the following example shows that revenue-ordered assortments cannot generally guarantee more than half of the optimal expected revenue.

Example 1. Let us consider the following assortment optimization instance, which is parameterized by $\epsilon \in (0, \frac{1}{2})$, with the additional assumption that $\frac{1}{\epsilon}$ is an integer. For ease of notation, we also define $\alpha = \frac{1}{\epsilon} + 1$ and $r = \frac{1}{\epsilon}$.

- The underlying collection of products is $1, \ldots, \alpha r + 1$.
- Product 1 has a revenue of $r_1 = r$ and a preference weight of $w_1 = 1$. For product 2, we set $r_2 = 1 + \epsilon$ and $w_2 = \infty$. For any other product $3 \leq i \leq \alpha r + 1$, we set $r_i = 1$ and $w_i = \frac{\epsilon}{\alpha r - 1}$. The no-purchase option has $w_0 = 0$.
- The maximum number of purchases M takes the values 1 and $\alpha r + 1$, such that $\Pr[M = 1] = \frac{\alpha}{1+\alpha}$ and $\Pr[M = \alpha r + 1] = \frac{1}{1+\alpha}$.

Clearly, since $r_1 > r_2 > r_3 = \cdots = r_{\alpha r+1}$, any revenue-ordered assortment falls into one of the next two cases:

- $S_1 = \{1\}$, which earns an expected revenue of $\mathcal{R}(S_1) = r = \frac{1}{\epsilon}$.
- $S_t = \{1, \ldots, t\}$ for some $2 \leq t \leq \alpha r + 1$, with an expected revenue of

$$
\begin{aligned}
\mathcal{R}(S_t) &= \Pr[M = 1] \cdot (1 + \epsilon) + \Pr[M = \alpha r + 1] \cdot (r + t - 1 + \epsilon) \\
&\leq \frac{\alpha}{1+\alpha} \cdot (1 + \epsilon) + \frac{1}{1+\alpha} \cdot ((1 + \alpha) \cdot r + \epsilon) \\
&\leq r + 1 + \epsilon \\
&\leq \frac{1}{\epsilon} + 2 \ .
\end{aligned}
$$

However, focusing on the non-revenue-ordered assortment $S = \{1, 3, 4, \ldots, \alpha r + 1\}$, we obtain an expected revenue of

$$
\mathcal{R}(S) = \Pr[M = 1] \cdot \frac{r + \epsilon}{1 + \epsilon} + \Pr[M = \alpha r + 1] \cdot (r + \alpha r - 1)
$$

$$
\begin{aligned}
&= \frac{\alpha}{1+\alpha} \cdot \frac{r+\epsilon}{1+\epsilon} + \frac{1}{1+\alpha} \cdot ((1+\alpha) \cdot r - 1) \\
&= \Bigg(1 + \underbrace{\frac{\alpha}{1+\alpha} \cdot \frac{1}{1+\epsilon}}_{=\frac{1}{1+2\epsilon}}\Bigg) \cdot r + \underbrace{\frac{\alpha}{1+\alpha} \cdot \frac{\epsilon}{1+\epsilon} - \frac{1}{1+\alpha}}_{=0} \\
&= \left(1 - \frac{\epsilon}{1+2\epsilon}\right) \cdot \frac{2}{\epsilon} \ .
\end{aligned}
$$

Consequently, we have just shown that the ratio between the best expected revenues achievable by revenue-ordered assortments and arbitrarily-structured ones can be upper-bounded by

$$
\frac{\max_{t \in [\alpha r+1]} \mathcal{R}(S_t)}{\mathcal{R}(S)} \leq \frac{\frac{1}{\epsilon} + 2}{(1 - \frac{\epsilon}{1+2\epsilon}) \cdot \frac{2}{\epsilon}} \ ,
$$

which tends to $\frac{1}{2}$ as ϵ tends to 0.

3.2. MNL-optimal assortments

While Example 1 cautions against blindly using revenue-ordered assortments in our multi-purchase setting, we are nonetheless able to derive several practical insights and results in this context. To start, we examine the effects of making assortment decisions that do not account for arriving customers who may purchase multiple products. More specifically, we study the efficacy of offering an assortment that forms an optimal solution to

$$
\max_{S \subseteq [n]} \left\{ \frac{1}{1+w(S)} \cdot \sum_{i \in S} r_i w_i \right\} \ , \tag{3}
$$

corresponding to the simplistic setting where we ignore purchases beyond each customer's most preferred product. Equivalently, such an assortment is optimal with respect to the standard MNL model, with preference weights $w_1, \dots, w_n$. The first result we establish reveals that there exists an optimal assortment under the MP-MNL model that includes each and every product offered under the standard MNL model. We then show that the expected revenue earned by an MNL-optimal assortment in our multi-purchase setting is within factor $\frac{1}{\mathbb{E}[M]}$ of the optimal expected revenue; this guarantee will be demonstrated to be tight.

The following property, which we exploit numerous times in our future analysis, formalizes the result of Talluri and Van Ryzin (2004), as it relates to the optimality of revenue-ordered assortments when choice is governed by a standard MNL model.

Property 3.1. *Let* $z^* = \max_{S \subseteq [n]} \{ \frac{1}{1+w(S)} \cdot \sum_{i \in S} r_i w_i \}$ *be the optimal objective value of problem* (3)*. Then,* $\tilde{S} = \{ i \in [n] : r_i \geq z^* \}$ *is an optimal assortment for this problem.*

Nesting structure. Lemma 2 below establishes an interesting structural property that relates the make-up of $\tilde{S}$ to that of an optimal assortment in our multi-purchase setting. Specifically, we show that there exists an optimal assortment to the latter problem that includes all products offered by $\tilde{S}$. This property, whose proof is provided in Appendix B.1, shows that in the presence of multi-purchase behavior, a larger and more diverse assortment is needed to capture optimal expected revenues.

LEMMA 2. *Let S^* be a maximum-cardinality optimal assortment to problem* (2)*. Then, $\tilde{S} \subseteq S^*$.*

Given this structural result, a very natural question is whether generalized "nestedness" properties can be established beyond the one related to single-purchase settings. More specifically, suppose that for $\bar{m} \in [n]$, we ignore purchases beyond each customer's $\bar{m}$-th most preferred product. In this case, the random demand M is replaced by $\min\{M, \bar{m}\}$, meaning that an optimal assortment would be

$$S^*_{\bar{m}} = \arg\max_{S \subseteq [n]} \sum_{m \in [\bar{m}]} \Pr\left[\min\{M, \bar{m}\} = m\right] \cdot \sum_{i \in S} r_i \cdot \pi_m(i, S) \ .$$

With this notation, Lemma 2 shows that $S^*_1 \subseteq S^*$, and it is natural to wonder whether $S^*_1 \subseteq S^*_2 \subseteq \cdots \subseteq S^*_n \subseteq S^*$. However, the following example shows that this nesting structure fails to hold in general.

EXAMPLE 2. Consider an assortment optimization instance consisting of five products, whose revenues and weights are given by the next table:

product	price	weight
1	115	12
2	153	2
3	155	2
4	157	0.5
5	374	4

In addition, customers are willing to purchase up to three products (i.e., $\Pr[M = 3] = 1$). By enumerating over all 2^5 possible assortments, it is easy to verify that $S^*_1 = \{5\}$ and $S^*_2 = \{1, 5\}$, whereas $S^* = \{2, 3, 4, 5\}$. In this case, we observe that $S^*_2 \not\subseteq S^*$, and hence the nesting property breaks for $\bar{m} = 2$.

Revenue guarantee. Interestingly, a close inspection of the proof we provide for Lemma 2 reveals that augmenting any assortment S with a product from $\tilde{S} \setminus S$ can only increase its expected revenue in our multi-purchase setting, hinting that offering $\tilde{S}$ by itself might be profitable in certain settings. Our next result formalizes this notion, showing that the assortment $\tilde{S}$ is guaranteed to garner at least a $\frac{1}{\mathbb{E}[M]}$-fraction of the optimal expected revenue. Consequently, in scenarios where the "average" customer is only willing to purchase very few products, revenue-ordered assortments have a small optimality gap under the MP-MNL model. In Lemma 3, whose proof appears in Appendix B.1, we use S^* to denote an arbitrary optimal assortment to problem (2).

Lemma 3. $\mathcal{R}(\tilde{S}) \geq \frac{1}{\mathbb{E}[M]} \cdot \mathcal{R}(S^*)$.

To establish the tightness this bound, the next example describes an instance of our multi-purchase assortment problem for which $\mathcal{R}(\tilde{S}) = O(\frac{1}{\mathbb{E}[M]}) \cdot \mathcal{R}(S^*)$.

Example 3. Suppose we have at our possession n products, all associated with a preference weight of 1, whereas the no-purchase option has weight $w_0 = 0$. In addition, the revenue of each product is $1 - \frac{1}{n}$, except for that of product 1, which is set to 1. When customers are always willing to purchase up to n products (i.e., $\Pr[M = n] = 1$), the assortment $S^* = [n]$ that includes all products is clearly optimal, garnering an expected revenue of $\mathcal{R}(S^*) = n$. In contrast, it is easy to verify that $\tilde{S} = \{1\}$ is the unique MNL-based optimal assortment, with $\mathcal{R}(\tilde{S}) = 1$, and we thus have $\mathcal{R}(\tilde{S}) = \frac{1}{n} \cdot \mathcal{R}(S^*) = \frac{1}{\mathbb{E}[M]} \cdot \mathcal{R}(S^*)$.

3.3. The Best Revenue-Ordered Assortment

Finally, we step outside of the traditional MNL framework, and analyze the performance of the most profitable revenue-ordered assortment, showing that its expected revenue is within factor $O(\log n)$ of the optimal expected revenue. Specifically, assuming that products are indexed in non-increasing revenue order, i.e, $r_1 \geq \cdots \geq r_n$, let $S^*_{\mathrm{ro}} = \max_{i \in [n]} \mathcal{R}([i])$ be the most profitable revenue-ordered assortment, whereas S^* will denote an arbitrary optimal assortment to problem (2). The following lemma, whose proof is presented in Appendix B.2, formalizes the performance guarantee attained by S^*_{ro}.

Lemma 4. $\mathcal{R}(S^*_{\mathrm{ro}}) = \Omega(\frac{1}{\log n}) \cdot \mathcal{R}(S^*)$.

It is worth pointing out that, since we clearly have $\mathcal{R}(S^*_{\mathrm{ro}}) \geq \mathcal{R}(\tilde{S})$, the performance guarantee reported in Lemma 3 also holds for the assortment S^*_{ro}. That said, Lemma 4 provides a stronger guarantee than that of Lemma 3 when the average customer wishes to purchase a large number of products.

Open Question. Whether or not revenue-ordered assortments can be shown to garner a constant fraction of the optimal expected revenue remains open.

4. An Approximation Scheme under Limited Purchases

In this section, we present an approximation scheme for the assortment optimization problem, which constitutes a PTAS when the maximum number of purchases M is supported over $\{0, 1, \ldots, m_{\max}\}$, with $m_{\max} = O(1)$. The precise performance guarantees of this algorithm can be formally stated as follows.

Theorem 1. *For any $\epsilon > 0$, there is a deterministic $O(n^{O(m_{\max}^2/\epsilon)})$-time algorithm for computing an assortment whose expected revenue is within factor $1 - \epsilon$ of optimal.*

Technical overview. At a high-level, our approximation scheme unfolds in two sequential steps. First, in Section 4.1, we recover the $O(\frac{m_{\max}^2}{\epsilon})$ highest weight products offered in the optimal assortment S^*. We refer to these products as "heavy" ones, while the remaining products are referred to as "light". For the latter, we guess the total weight of light products offered by S^* within a sufficiently small factor. Subsequently, in Section 4.2, we propose an approximate way of reducing our assortment problem of interest to a knapsack-like problem. To this end, we prove that one can utilize efficiently-computable coefficients to approximate the choice probabilities of any assortment that: (1) Includes all heavy products, and (2) Satisfies our total weight bounds for light products. These coefficients will be obtained through a carefully crafted adaptation of the recursive procedure prescribed by Lemma 1, originally employed to compute the choice probabilities in an exact way.

4.1. Step 1: Guessing

In what follows, we employ efficient enumeration ideas to recover the set of so-called heavy products. In parallel, we explain how to approximately estimate the combined weight of all light products offered in the optimal assortment S^*. To avoid cumbersome notation, we assume without loss of generality that $\frac{1}{\epsilon}$ takes an integer value.

Guessing heavy products. First, we guess the $\mathcal{N} = \frac{6m_{\max}^2}{\epsilon}$ largest-weight products offered in S^*, and refer to this collection as the set of heavy products, $\mathcal{H}^*$. Clearly, $\mathcal{H}^*$ can be exactly recovered by enumerating over all possible $O(n^{\mathcal{N}}) = O(n^{O(m_{\max}^2/\epsilon)})$ subsets of cardinality $\mathcal{N}$. When $|S^*| \leq \mathcal{N}$, this initial guessing step will recover the optimal assortment. To handle the non-trivial case where $|S^*| > \mathcal{N}$, we use $\mathcal{L} = \{i \in [n] \setminus \mathcal{H}^* : w_i \leq \min_{j \in \mathcal{H}^*} w_j\}$ to designate the remaining set of light products, whose weights do not exceed that of any heavy product, and discard all products outside of $\mathcal{H}^* \cup \mathcal{L}$. Finally, we let $\mathcal{L}^* = S^* \cap \mathcal{L}$ denote the set of light products included in the optimal assortment.

Guessing the total weight of light products. To begin, we guess $\bar{w} = \max\{w_i : i \in \mathcal{L}^*\}$, which is the heaviest weight of a light product offered in the optimal assortment S^*, for which there are only $O(n)$ options. Next, let k^* be the smallest non-negative integer satisfying $\bar{w} \cdot (1+\delta)^k \geq w(\mathcal{L}^*)$, where $\delta = \frac{\epsilon}{6m_{\max}}$. Given that $w(\mathcal{L}^*) \in [\bar{w}, n\bar{w}]$, it is easy to verify that $k^* = O(\frac{m_{\max}}{\epsilon} \cdot \log n)$, and we can therefore efficiently enumerate over all possible values for this parameter. We conclude our guessing step by defining $W = \bar{w} \cdot (1+\delta)^{k^*}$; clearly, $\frac{W}{1+\delta} \leq w(\mathcal{L}^*) \leq W$, by the choice of k^*.

4.2. Step 2: Reduction to a knapsack-like problem

Here, we first show how to efficiently compute assortment-independent coefficients $\{\tilde{\pi}_i\}_{i \in [n]}$ that will be plugged-in as estimates for the choice probabilities. Next, we replace the choice probabilities within the expected revenue function (2) with these coefficients, thereby transforming our assortment problem of interest into a knapsack-like problem. Our final piece of analysis concludes that an approximate solution to the latter problem directly yields an assortment whose expected revenue is within factor $1-\epsilon$ of optimal.

Computing the choice probability estimators. Consider the following dynamic program, where each state (i,m,H) corresponds to a product $i \in [n]$, a number of purchases $m \in [m_{\max}]_0$, and a subset of heavy products $H \subseteq \mathcal{H}^*$. Our value function is defined through the recursive expression

$$V(i,m,H) = \underbrace{\frac{1}{1+(1+\delta)\cdot W + w(H)}}_{\text{part (i)}} \cdot \left(\underbrace{w_i + \left[\frac{W}{(1+\delta)^2} - \delta \cdot w(H_{-i}) \right]^+ \cdot V(i,m-1,H)}_{\text{part (ii)}} + \underbrace{\sum_{j \in H_{-i}} w_j \cdot V(i,m-1,H_{-j})}_{\text{part (iii)}} \right) , \tag{4}$$

with base cases of $V(\cdot,0,\cdot)=0$. It is not difficult to verify that, given $\mathcal{H}^*$, the total running time required to compute $V(i,m,\mathcal{H})$ over all possible states is only $O(n^2 m_{\max} \cdot 2^{\mathcal{N}}) = O(n^2 \cdot 2^{O(m_{\max}^2/\epsilon)})$.

Lemma 5 below, whose proof is deferred to Section 4.4, relates this value function to choice probabilities with respect to assortments that "resemble" the optimal one, S^*. Specifically, let

$$\mathcal{F} = \left\{ S \subseteq (\mathcal{L} \cup \mathcal{H}^*) : w(S \cap \mathcal{L}) \in \left[\frac{W}{(1+\delta)^2}, (1+\delta)\cdot W \right], S \supseteq \mathcal{H}^* \right\}$$

be the family of assortments that contain all heavy products, in which the total weight of all light products nearly matches $w(\mathcal{L}^*)$, as defined above. Then, the next result reveals that our value function can be used to ϵ-estimate the choice probability of any product offered by such an assortment.

Lemma 5. *For any assortment $S \in \mathcal{F}$ and product $i \in S$, we have*

$$(1-2\epsilon)\cdot \pi(i,S) \leq \mathbb{E}[V(i,M,\mathcal{H}^*)] \leq \pi(i,S) ,$$

where the expectation above is taken over the randomness in M.

Constructing the knapsack problem. We proceed by presenting a knapsack-based approximation of our original assortment optimization problem, where one wishes to compute an assortment $S \subseteq [n]$ whose expected revenue $\mathcal{R}(S) = \sum_{i \in S} r_i \cdot \pi(i,S)$ is maximized. For this purpose, we replace each choice probability $\pi(i,S)$ with light-product-independent proxies $\tilde{\pi}_i = \mathbb{E}[V(i,M,\mathcal{H}^*)]$. By plugging these coefficients into the expected revenue function, we introduce an extension of the classic knapsack problem, where our objective is to maximize $\sum_{i \in S} r_i \tilde{\pi}_i$, subject to picking all heavy products and offering light products whose total weight is within $[\frac{W}{1+\delta}, W]$. This problem can be compactly written as

$$\mathrm{OPT}_{\text{knapsack}} = \sum_{i \in \mathcal{H}^*} r_i \tilde{\pi}_i + \left[\begin{array}{ll} \max\limits_{S \subseteq \mathcal{L}} & \sum\limits_{i \in S} r_i \tilde{\pi}_i \\ \text{s.t.} & w(S) \in [\frac{W}{1+\delta}, W] \end{array} \right] \tag{5}$$

The inner problem above can be viewed as a variant of the classic binary knapsack problem, where a lower bound on the total weight is enforced in conjunction with the traditional upper bound. By rounding down product weights to the nearest multiple of $\frac{\delta W}{n}$, standard dynamic programming ideas (see, e.g., (Vazirani 2013, Chap. 8)) can be employed to compute, in $O(\frac{n^2}{\delta})$ time, a subset of products $\tilde{S} \in \mathcal{F}$ with

$$\sum_{i \in \tilde{S}} r_i \tilde{\pi}_i \geq \mathrm{OPT}_{\mathrm{knapsack}} \geq \sum_{i \in S^*} r_i \tilde{\pi}_i \ , \tag{6}$$

where the second inequality follows since the optimal assortment $S^* = \mathcal{H}^* \cup \mathcal{L}^*$ constitutes a feasible solution to problem (5), as $\frac{W}{1+\delta} \leq w(\mathcal{L}^*) \leq W$. In essence, rounding down weights to multiples of $\frac{\delta W}{n}$ allows us to compute a super-optimal knapsack solution, at the expense of violating the bounds $[\frac{W}{1+\delta}, W]$ by a factor of $1 \pm \delta$, thereby constructing an assortment in $\mathcal{F}$. We conclude our approximation scheme by returning the set of products $\tilde{S}$ as its final assortment.

4.3. Performance guarantee and running time

Approximation ratio. We proceed by arguing that the assortment $\tilde{S}$ is indeed near-optimal. To this end, one should observe that the expected revenue attained by $\tilde{S}$ is

$$\begin{aligned}
\mathcal{R}(\tilde{S}) &= \sum_{i \in \tilde{S}} r_i \cdot \pi(i, \tilde{S}) \\
&\geq \sum_{i \in \tilde{S}} r_i \tilde{\pi}_i \\
&\geq \sum_{i \in S^*} r_i \tilde{\pi}_i \\
&\geq (1 - 2\epsilon) \cdot \sum_{i \in S^*} r_i \cdot \pi(i, S^*) \\
&= (1 - 2\epsilon) \cdot \mathcal{R}(S^*).
\end{aligned}$$

Here, the first and last inequalities hold since $\pi(i, \tilde{S}) \geq \tilde{\pi}_i$ and $\tilde{\pi}_i \geq (1 - 2\epsilon) \cdot \pi(i, S^*)$, respectively, by Lemma 5. The middle inequality is precisely the relation between $\tilde{S}$ and S^*, given by (6).

Running time. The specific points that should be taken into account when deriving the overall running time of our approach can be briefly summarized as follows:

- In Step 1, there are $O(n^{O(m_{\max}^2/\epsilon)})$ guesses for the set of heavy products $\mathcal{H}^*$. In addition, to obtain the approximate total weight W of light products, only $O(n \cdot \frac{m_{\max}}{\epsilon} \cdot \log n)$ guesses are required.
- For each such set of guesses, we proceed in Step 2 to compute the choice probability estimators $V(\cdot, \cdot, \cdot)$ in $O(n^{O(1)} \cdot 2^{O(\frac{m_{\max}}{\epsilon})})$ time. Then, their related knapsack problem is approximated in $O(\frac{n^2 m_{\max}}{\epsilon})$ time.

All together, we arrive at an overall running time of $O(n^{O(m_{\max}^2/\epsilon)})$, as stated in Theorem 1.

4.4. Proof of Lemma 5

Auxiliary claims. We begin by stating two basic results, whose proofs are provided in Appendices C.1 and C.2, respectively. The first claim shows that the choice probability of any particular product within a given assortment increases more by the removal of a more attractive product than by the removal of a less attractive product. The second claim places a lower bound on the total weight of particular collections of light products.

Claim 1. *Let $S \subseteq [n]$ be an assortment, and let $j, \ell \in S$ be two distinct products with $w_j \geq w_\ell$. Then, for every product $i \in S \setminus \{j, \ell\}$, we have $\pi_m(i, S_{-j}) \geq \pi_m(i, S_{-\ell})$ for all $m \in [n]$.*

Claim 2. *Let $S \in \mathcal{F}$ be an assortment and let $Q \subseteq S$ be a subset of products with $|S \setminus Q| \leq m_{\max}$. Then, $w(Q \cap \mathcal{L}) \geq \frac{W}{(1+\delta)^2} - \delta \cdot w(Q \cap \mathcal{H}^*)$.*

The generalized claim. In what follows, we prove Lemma 5 by establishing a somewhat more general claim. For this purpose, consider some assortment $S \in \mathcal{F}$ and a product $i \in S$ within this assortment. The next claim provides tight bounds, in terms of the value function V, on the probability that product i is purchased out of certain sub-assortments of S.

Claim 3. *Let $Q \subseteq S$ be a subset of products with $i \in Q$, and suppose that $|S \setminus Q| \leq m_{\max} - m$ for some $m \in [m_{\max}]$. Then,*

$$(1-\delta)^{6m} \cdot \pi_m(i, Q) \leq V(i, m, \mathcal{H}^* \cap Q) \leq \pi_m(i, Q) \ .$$

To understand how the above claim leads to Lemma 5, note that by choosing $Q = S$, we have $V(i, m, \mathcal{H}^*) \leq \pi_m(i, S)$ for all $m \in [m_{\max}]$, and therefore, $\mathbb{E}[V(i, M, \mathcal{H}^*)] \leq \pi(i, S)$. In the opposite direction, since we simultaneously get $V(i, m, \mathcal{H}^*) \geq (1-\delta)^{6m} \cdot \pi_m(i, S)$ for all $m \in [m_{\max}]$, it follows that

$$\begin{aligned}
\mathbb{E}\left[V(i, M, \mathcal{H}^*)\right] &\geq \sum_{m \in [m_{\max}]} \Pr\left[M = m\right] \cdot (1-\delta)^{6m} \cdot \pi_m(i, S) \\
&\geq \left(1 - \frac{\epsilon}{6m_{\max}}\right)^{6m_{\max}} \cdot \pi(i, S) \\
&\geq e^{-2\epsilon} \cdot \pi(i, S) \\
&\geq (1-2\epsilon) \cdot \pi(i, S) \ .
\end{aligned}$$

We prove both inequalities of Claim 3 via induction over m. For this purpose, consider the following representation of the choice probability $\pi_m(i, Q)$, which is derived from Lemma 1:

$$\begin{aligned}
\pi_m(i, Q) &= \frac{w_i}{1 + w(Q)} + \sum_{\ell \in Q_{-i}} \frac{w_\ell}{1 + w(Q)} \cdot \pi_{m-1}(i, Q_{-\ell}) \\
&= \underbrace{\frac{1}{1 + w(Q)}}_{\text{part (i)}} \cdot \left(\underbrace{w_i + \sum_{\ell \in \mathcal{L} \cap Q_{-i}} w_\ell \cdot \pi_{m-1}(i, Q_{-\ell})}_{\text{part (ii)}} + \underbrace{\sum_{j \in \mathcal{H}^* \cap Q_{-i}} w_j \cdot \pi_{m-1}(i, Q_{-j})}_{\text{part (iii)}} \right). \qquad (7)
\end{aligned}$$

Our induction argument will compare parts (i)-(iii) above to their counterparts in the recursive equation (4) by which the value function V is instantiated with $H = \mathcal{H}^* \cap Q$.

The upper bound proof: $\pi_m(i,Q) \geq V(i,m,H)$. Since the base case of $m = 0$ trivially holds, we move on to proving the required bound for general $m \geq 1$. To this end, we show that each part in (7) exceeds its counterpart in (4):

- Part (i): To verify this comparison, note that

$$\begin{aligned} w(Q) &= w(\mathcal{L} \cap Q) + w(\mathcal{H}^* \cap Q) \\ &\leq w(\mathcal{L} \cap S) + w(H) \\ &\leq (1+\delta) \cdot W + w(H) \ , \end{aligned}$$

where the first inequality holds since $Q \subseteq S$, and the second inequality follows by recalling that $w(S \cap \mathcal{L}) \leq (1+\delta) \cdot W$, as $S \in \mathcal{F}$.

- Part (ii): To make this comparison, we observe that

$$\begin{aligned} \sum_{\ell \in \mathcal{L} \cap Q_{-i}} w_\ell \cdot \pi_{m-1}(i, Q_{-\ell}) &\geq \pi_{m-1}(i,Q) \cdot w(\mathcal{L} \cap Q_{-i}) \\ &\geq \pi_{m-1}(i,Q) \cdot \left[\frac{W}{(1+\delta)^2} - \delta \cdot w(H_{-i}) \right]^+ \\ &\geq V(i, m-1, H) \cdot \left[\frac{W}{(1+\delta)^2} - \delta \cdot w(H_{-i}) \right]^+ \ . \end{aligned}$$

Here, the first inequality holds since adding any product $\ell \in \mathcal{L} \cap Q_{-i}$ to Q can only reduce the choice probability of product i. The second inequality is obtained by applying Claim 2, noting that $\mathcal{H}^* \cap Q_{-i} = H_{-i}$, by definition of H, regardless of whether product i is light or heavy. The last inequality follows from the induction hypothesis, which can be applied since $|S \setminus Q| \leq m_{\max} - m < m_{\max} - (m-1)$.

- Part (iii): In this case,

$$\begin{aligned} \sum_{j \in \mathcal{H}^* \cap Q_{-i}} w_j \cdot \pi_{m-1}(i, Q_{-j}) &= \sum_{j \in H_{-i}} w_j \cdot \pi_{m-1}(i, Q_{-j}) \, . \\ &\geq \sum_{j \in H_{-i}} w_j \cdot V(i, m-1, H_{-j}) \end{aligned}$$

The equality above holds since $\mathcal{H}^* \cap Q_{-i} = H_{-i}$, as explained in the previous item. The next inequality follows by invoking the induction hypothesis, which is indeed applicable since $|S \setminus Q_{-j}| = |S \setminus Q| + 1 \leq m_{\max} - (m-1)$, noting that $\mathcal{H}^* \cap Q_{-j} = H_{-j}$, as $j \in \mathcal{H}^*$.

The lower bound proof: $V(i,m,H) \geq (1-\delta)^{6m} \cdot \pi_m(i,Q)$. While the current claim follows from arguments similar to those of our upper bound proof, its specifics are somewhat more involved, and we therefore present these details in Appendix C.3.

5. A PTAS for the General Problem

This section is devoted to presenting a high-level overview of our approximation scheme for the assortment optimization problem in its most general form, without any structural assumptions whatsoever. The specifics of this result are formally stated in the next theorem, noting that its success probability can be amplified to $1-\delta$, for any $\delta \in (0,1)$, with $O(\log \frac{1}{\delta})$ independent repetitions.

Theorem 2. *For any $\epsilon > 0$, there is a randomized algorithm for computing, with probability at least $1/2$, an assortment $S \subseteq [n]$ whose expected revenue is within factor $1-\epsilon$ of optimal. The running time of this algorithm is $O(n^{O(1/\epsilon^3)})$.*

The upcoming overview provides a detailed account of our algorithm and its analysis, while concurrently highlighting the technical hurdles that arise along the way. It turns out that avoiding an $O(1)$-bound on the number of maximum purchases requires a new set of tools and ideas, which are substantially more intricate than those presented in Section 4. As such, for readability purposes, we focus on the bigger picture and defer most technical details to subsequent appendices. It is important to emphasize that we do not attempt to present the most efficient implementation by any means, and leave for future work the task of distilling our approach into a more practical variant.

5.1. The bounded-ratio setting

For ease of exposition, in the remainder of this section, we derive a weaker version of Theorem 2, where the running time of our approach includes a dependency on the extremal product weights. As formally stated below, we devise a randomized approximation scheme whose running time is polynomial in the input size as well as in the ratio between extremal weights, $\frac{w_{\max}}{w_{\min}}$, where $w_{\max} = \max_{i\in[n]} w_i$ and $w_{\min} = \min_{i\in[n]} w_i$. Here, one can assume without loss of generality that $w_{\min} > 0$, as all products with zero preference weights can be discarded in advance.

Theorem 3. *For any $\epsilon > 0$, there is a randomized algorithm for computing, with probability at least $1/2$, an assortment $S \subseteq [n]$ whose expected revenue is within factor $1-\epsilon$ of optimal. The running time of this algorithm is $O((n \cdot \frac{w_{\max}}{w_{\min}})^{O(1/\epsilon^2)})$.*

Addressing the general setting. It is important to point out that, for arbitrarily-structured instances, the ratio $\frac{w_{\max}}{w_{\min}}$ could be exponential in the input size, implying that Theorem 3 should not be viewed as a PTAS by itself. That said, this result will be employed as a subroutine in Appendix E, upon addressing the assortment optimization problem in its utmost generality. At a high level, our general approximation scheme will be obtained by decomposing any given instance into a sequence of sub-instances, each with a polynomially-bounded $\frac{w_{\max}}{w_{\min}}$ ratio; in this case, the approximation scheme of Theorem 3 actually constitutes a true PTAS. Here, the technical crux resides in bounding the interaction between these sub-instances through appropriately-defined cardinality constraints. To this end, we mention that our bounded-ratio PTAS works in its current form even subject to a cardinality constraint on the offered assortment, and briefly discuss the intuition needed to confirm the validity of this extension at the end of Section 5.5.

5.2. The stability of choice probabilities under rounding

To simplify the analysis of our algorithmic ideas, we first examine the following question: What is the extent to which choice probabilities are affected by small perturbations to the preference weights of the underlying products?

The effects of single product rounding. To derive sufficiently strong bounds, let us denote by $\pi^{\langle w \rangle}(i, S)$ the choice probability of product i within the assortment S, assuming that the underlying preference weights are given by $w = (w_0, w_1, \ldots, w_n)$. Now suppose we define a new vector $w^{\downarrow i}$, whose coordinates are identical to those of w, except for the preference weight of product i, which is rounded down to $w_i^{\downarrow i} \in [(1-\epsilon) \cdot w_i, w_i]$. The choice probabilities with respect to this vector will be denoted by $\pi^{\langle w^{\downarrow i} \rangle}(\cdot, \cdot)$. The next claim, whose proof appears in Appendix D.1, argues that with respect to any assortment, the choice probability of product i will remain within factor $1-\epsilon$ of its original value, whereas the choice probability of any other product can only increase.

LEMMA 6. *For any assortment $S \subseteq [n]$, we have:*
- $\pi^{\langle w^{\downarrow i} \rangle}(i, S) \geq (1-\epsilon) \cdot \pi^{\langle w \rangle}(i, S)$.
- $\pi^{\langle w^{\downarrow i} \rangle}(j, S) \geq \pi^{\langle w \rangle}(j, S)$, *for every product $j \neq i$.*

Fully rounded instances. Now, in order to inject additional structure into a given instance $\mathcal{I}$ of the assortment optimization problem, we define its rounded-down counterpart $\mathcal{I}^{\downarrow}$. In the latter instance, the preference weight of each product is rounded down to the nearest power of $1+\epsilon$, thereby obtaining the weight vector $w^{\downarrow}$. In contrast, the no-purchase option retains its weight, i.e., $w_0^{\downarrow} = w_0 = 1$. In the next lemma, we exploit the single-product rounding bound of Lemma 6 to show that the expected revenue $\mathcal{R}^{\langle w^{\downarrow} \rangle}(S)$ of any assortment S with respect to $w^{\downarrow}$ differs by a multiplicative factor of $1 \pm 2\epsilon$ from its original value $\mathcal{R}^{\langle w \rangle}(S)$ with respect to w. The specifics of this proof are provided in Appendix D.2.

LEMMA 7. $\mathcal{R}^{\langle w^{\downarrow} \rangle}(S) \in (1 \pm 2\epsilon) \cdot \mathcal{R}^{\langle w \rangle}(S)$, *for any assortment $S \subseteq [n]$.*

5.3. The within-class revenue-orderedness of optimal assortments

In light of the preceding discussion, it remains to propose an approximation scheme for the rounded instance $\mathcal{I}^{\downarrow}$, as Lemma 7 argues that we incur only $O(\epsilon)$-losses in optimality when moving between $\mathcal{I}$ and $\mathcal{I}^{\downarrow}$ in either direction. For this purpose, letting $\mathcal{Q} = \lceil \log_{1+\epsilon}(\frac{w_{\max}}{w_{\min}}) \rceil$, we focus our attention on assortment optimization instances where each preference weight is of the form $w_{\min} \cdot (1+\epsilon)^q$, for some integer $q \in [\mathcal{Q}]_0$. We refer to the collection of products $i \in [n]$ with $w_i = w_{\min} \cdot (1+\epsilon)^q$ as the weight class $\mathcal{W}_q$, and for an integer $k \in [|\mathcal{W}_q|]_0$, we use $\mathcal{W}_q[k]$ to denote the set of k most expensive products in $\mathcal{W}_q$, breaking ties arbitrarily. The next claim, whose proof appears in Appendix D.3, shows that within each weight class, had we known the number of products picked by some fixed optimal assortment, it is optimal to pick the most expensive ones.

LEMMA 8. *There exists an optimal assortment S^* where, for every $q \in [\mathcal{Q}]_0$,*

$$S^* \cap \mathcal{W}_q = \mathcal{W}_q\left[|S^* \cap \mathcal{W}_q|\right] \ .$$

5.4. The enumeration-based approximation scheme

Letting S^* be an optimal assortment satisfying the structural condition of Lemma 8, it suffices to guess the number of products $k_q^* = |S^* \cap \mathcal{W}_q|$ offered out of each class $\mathcal{W}_q$. A naive implementation would require enumerating over $O(n^{O(\mathcal{Q})}) = O(n^{O(\frac{1}{\epsilon}\log\frac{w_{\max}}{w_{\min}})})$ joint configurations for the values $\{k_q^*\}_{q\in[\mathcal{Q}]_0}$. However, to eventually obtain a polynomial-time approximation scheme along the lines of Theorem 2, it is crucial that the running time dependency of our current enumeration procedure will be polynomial in $\frac{w_{\max}}{w_{\min}}$. For this purpose, our refined procedure operates as follows:

- **Step 1: Guessing choice probabilities.**

—For each weight class $\mathcal{W}_q$, due to having identical preference weights, all products offered by S^* from this class have precisely the same choice probability, which will be referred to as π_q^*. One can easily verify that $\pi_0^* \leq \cdots \leq \pi_{\mathcal{Q}}^*$.

—We first guess the identity of $q_{\min}$, which is the smallest index $q \in [\mathcal{Q}]_0$ for which $\pi_q^* \geq \frac{\epsilon}{n} \cdot \Pr[M>0] \cdot \frac{w_{\min}}{1+w_{\min}}$, with the convention that $q_{\min} = \infty$ when no such index exists. We also note that $\pi_q^* \leq \Pr[M>0] \cdot \frac{w_{\min}\cdot(1+\epsilon)^{\mathcal{Q}}}{1+w_{\min}\cdot(1+\epsilon)^{\mathcal{Q}}}$ for every $q \in [\mathcal{Q}]_0$; the latter inequality holds since, for any product $i \in S^* \cap \mathcal{W}_q$, we have

$$\begin{aligned}
\pi_q^* &= \pi\left(i, S^*\right) \\
&\leq \pi(i, \{i\}) \\
&= \Pr\left[M>0\right] \cdot \frac{w_{\min} \cdot (1+\epsilon)^q}{1+w_{\min} \cdot (1+\epsilon)^q} \\
&\leq \Pr\left[M>0\right] \cdot \frac{w_{\min} \cdot (1+\epsilon)^{\mathcal{Q}}}{1+w_{\min} \cdot (1+\epsilon)^{\mathcal{Q}}} \ .
\end{aligned}$$

It follows that the subsequence $\pi_{q_{\min}}^*, \ldots, \pi_{\mathcal{Q}}^*$ is bounded within the interval

$$\left[\frac{\epsilon}{n} \cdot \Pr[M>0] \cdot \frac{w_{\min}}{1+w_{\min}}, \Pr[M>0] \cdot \frac{w_{\min} \cdot (1+\epsilon)^{\mathcal{Q}}}{1+w_{\min} \cdot (1+\epsilon)^{\mathcal{Q}}}\right] \ ,$$

whose endpoints differ by a factor of at most $(1+\epsilon)^{\mathcal{Q}} \cdot \frac{n}{\epsilon}$.

—Therefore, by enumerating over all non-decreasing sequences of powers-of-$(1+\epsilon)$ in this interval, we can obtain over-estimates $\tilde{\pi}_{q_{\min}} \leq \cdots \leq \tilde{\pi}_{\mathcal{Q}}$ such that $\tilde{\pi}_q \in [\pi^*_q, (1+\epsilon)\cdot\pi^*_q]$ for every $q_{\min} \leq q \leq \mathcal{Q}$. Since any such sequence can be written as $(1+\epsilon)^{x_{q_{\min}}}, \ldots, (1+\epsilon)^{x_{\mathcal{Q}}}$ for a sequence of integers $x_{q_{\min}} \leq \cdots \leq x_{\mathcal{Q}}$ with $x_{\mathcal{Q}} - x_{q_{\min}} \leq \log_{1+\epsilon}((1+\epsilon)^{\mathcal{Q}}\cdot\frac{n}{\epsilon})$, elementary identical-balls-into-distinct-bins counting shows that the number of options to be examined is

$$\begin{aligned}\binom{\mathcal{Q} + \lceil \log_{1+\epsilon}((1+\epsilon)^{\mathcal{Q}}\cdot\frac{n}{\epsilon})\rceil}{\mathcal{Q}} &= O\left(2^{O(\mathcal{Q}+\log_{1+\epsilon}((1+\epsilon)^{\mathcal{Q}}\cdot\frac{n}{\epsilon}))}\right)\\ &= O\left(2^{O(\frac{1}{\epsilon}\log\frac{w_{\max}}{w_{\min}} + \frac{1}{\epsilon}\log\frac{n}{\epsilon}}\right)\\ &= O\left(\left(n\cdot\frac{w_{\max}}{w_{\min}}\right)^{\tilde{O}(1/\epsilon)}\right) .\end{aligned}$$

- **Step 2: Guessing revenue contributions.**

—We first obtain an under-estimate $\widetilde{\mathrm{OPT}}$ of the optimal expected revenue $\mathcal{R}(S^*)$ that satisfies $\widetilde{\mathrm{OPT}} \in [(1-\epsilon)\cdot\mathcal{R}(S^*), \mathcal{R}(S^*)]$. To identify an ample interval over which we will enumerate, let i^* be the product maximizing $\mathcal{R}(\{i\})$ over all $i \in [n]$. Since it is easy to verify that $\mathcal{R}(S^*) \in [\mathcal{R}(\{i^*\}), n\cdot\mathcal{R}(\{i^*\})]$, our estimate $\widetilde{\mathrm{OPT}}$ can be efficiently guessed by enumerating over all powers of $1+\epsilon$ within this interval. Clearly, the number of candidate estimates is only $O(\frac{1}{\epsilon}\log n)$, as the interval endpoints differ by a factor of n.

—For each weight class $\mathcal{W}_q$, we define its contribution toward the overall expected revenue as $\mathcal{R}_q(S^*) = \sum_{i\in S^*\cap\mathcal{W}_q} r_i\cdot\pi(i,S^*) = \pi^*_q\cdot\sum_{i\in S^*\cap\mathcal{W}_q} r_i$. Now, for every $q \geq q_{\min}$, we guess an over-estimate $\tilde{\mathcal{R}}_q$ for the revenue contribution $\mathcal{R}_q(S^*)$ up to an additive factor of $\frac{\epsilon}{\mathcal{Q}+1}\cdot\widetilde{\mathrm{OPT}}$. In other words, $\mathcal{R}_q(S^*) - \frac{\epsilon}{\mathcal{Q}+1}\cdot\widetilde{\mathrm{OPT}} \leq \tilde{\mathcal{R}}_q \leq \mathcal{R}_q(S^*)$. To enumerate over all possible joint configurations of $\{\tilde{\mathcal{R}}_q\}_{q\geq q_{\min}}$, identical-balls-into-distinct-bins counting implies once again that the number of guesses to consider is only $O(2^{O(\mathcal{Q}/\epsilon)}) = O((\frac{w_{\max}}{w_{\min}})^{O(1/\epsilon^2)})$.

- **Step 3: Inferring a near-optimal assortment.** Finally, out of each weight class $\mathcal{W}_q$ with $q \geq q_{\min}$, we offer the $\tilde{k}_q$ most expensive products, where $\tilde{k}_q$ is the minimal integer k for which $\sum_{i\in\mathcal{W}_q[k]} r_i \geq \frac{\tilde{\mathcal{R}}_q}{\tilde{\pi}_q}$. Aggregating these decisions over all weight classes in question, we use $\tilde{S}$ to designate the resulting assortment, i.e., $\tilde{S} = \bigcup_{q\geq q_{\min}} \mathcal{W}_q[\tilde{k}_q]$.

5.5. Analysis

Approximation guarantee. In the next claim, whose proof is provided in Appendix D.4, we argue that the resulting assortment $\tilde{S}$ is indeed near-optimal, by exploiting the specific structure of our estimates for per-class choice probabilities and revenue contributions.

LEMMA 9. *$\mathcal{R}(\tilde{S}) \geq (1-3\epsilon)\cdot\mathcal{R}(S^*)$.*

Running time and revenue estimation. A close inspection of our enumeration method reveals that the overall number of possible guesses for the choice probability estimates $\tilde{\pi}_{q_{\min}} \leq \cdots \leq \tilde{\pi}_{\mathcal{Q}}$, the class index $q_{\min}$, the optimal revenue estimate $\widetilde{\mathrm{OPT}}$, and the revenue contribution estimates $\{\tilde{\mathcal{R}}_q\}_{q \geq q_{\min}}$ is only $O((n \cdot \frac{w_{\max}}{w_{\min}})^{O(1/\epsilon^2)})$. Each such guess uniquely translates to a candidate assortment, and therefore, it remains to compute the expected revenue of these assortments and to pick the most profitable one. That said, as outlined in Section 2.2, evaluating our recursive expression for the choice probabilities within a given assortment is generally exponential in the number of underlying products.

To bypass this obstacle, we describe an approximate way to efficiently estimate the expected revenue of a given assortment S. Specifically, for an error parameter $\epsilon > 0$ and a confidence level $\delta > 0$, our objective is to compute an estimator $\tilde{\mathcal{R}}(S)$ for the expected revenue $\mathcal{R}(S)$ that satisfies

$$\Pr\left[\left|\tilde{\mathcal{R}}(S) - \mathcal{R}(S)\right| \leq \epsilon \cdot \mathcal{R}(S^*)\right] \geq 1 - \delta \ .$$

It is important to point out that this guarantee is not multiplicative in its current form, since the error term $\epsilon \cdot \mathcal{R}(S^*)$ depends on the optimal assortment S^* rather than on S; we will explain how to go around this issue in the next paragraph. Specifically, our estimator will be derived as follows:

1. Letting $N = \lceil \frac{n^2}{2\epsilon^2} \cdot \ln(\frac{2n^2}{\delta}) \rceil$, we first draw N independent samples $\rho_1, \ldots, \rho_N$ from the distribution $[\text{rank}(i,S) | U_i > U_0]$.
2. For every product $i \in S$, we define an auxiliary estimator $\tilde{\pi}(i,S)$ for the choice probability $\pi(i,S)$, by setting $\tilde{\pi}(i,S) = \sum_{m \in [n]} \Pr[M = m] \cdot \tilde{\pi}_m(i,S)$, where $\tilde{\pi}_m(i,S) = \frac{w_i}{1+w_i} \cdot \frac{1}{N} \cdot \sum_{\nu \in [N]} \mathbb{1}[\rho_\nu \leq m]$.
3. Our final estimator is given by $\tilde{\mathcal{R}}(S) = \sum_{i \in S} r_i \cdot \tilde{\pi}(i,S)$.

The next claim, whose proof is given in Appendix D.5, shows that $\tilde{\mathcal{R}}(S)$ indeed forms an additive $\pm\epsilon \cdot \mathcal{R}(S^*)$ estimate for the expected revenue $\mathcal{R}(S)$.

Lemma 10. $\Pr[|\tilde{\mathcal{R}}(S) - \mathcal{R}(S)| \leq \epsilon \cdot \mathcal{R}(S^*)] \geq 1 - \delta$.

Picking a near-optimal assortment. In conclusion, recalling that our enumeration method is required to compare $O((n \cdot \frac{w_{\max}}{w_{\min}})^{O(1/\epsilon^2)})$ assortments, we employ the estimation procedure described above with a confidence level of $\delta = \Theta(1/(n \cdot \frac{w_{\max}}{w_{\min}})^{O(1/\epsilon^2)})$. By choosing an appropriate constant within the latter term, a simple application of the union bound guarantees that, with probability say $1/2$, we simultaneously get the next two guarantees:

1. In conjunction with Lemma 9, the estimated revenue $\tilde{\mathcal{R}}(\tilde{S})$ of the assortment $\tilde{S}$ evaluates to at least $(1 - 4\epsilon) \cdot \mathcal{R}(S^*)$.
2. The estimated revenue $\tilde{\mathcal{R}}(S)$ of any assortment S with $\mathcal{R}(S) \leq (1 - 6\epsilon) \cdot \mathcal{R}(S^*)$ evaluates to at most $(1 - 5\epsilon) \cdot \mathcal{R}(S^*)$.

Therefore, with probability at least 1/2, by picking the most profitable assortment with respect to the estimated revenue $\tilde{\mathcal{R}}(\cdot)$, we indeed identify one whose true expected revenue is at least $(1-6\epsilon)\cdot\mathcal{R}(S^*)$.

Incorporating a cardinality constraint. As mentioned at the beginning of this section, our algorithmic approach for the bounded-ratio setting works even subject to a cardinality constraint on the offered assortment. This extension turns out to be very useful once we focus on general problem instances in Appendix E. To better understand how one can seamlessly handle a cardinality constraint, the key ideas to take note of are:

- Lemma 8 proves that the optimal assortment S^* for the rounded instance satisfies $S^* \cap \mathcal{W}_q = \mathcal{W}_q[|S^* \cap \mathcal{W}_q|]$ for each weight class $\mathcal{W}_q$, i.e., S^* is revenue-ordered within each weight class.
- With this structure in-mind, we guess in Step 2 an under-estimate $\tilde{R}_q$ for the revenue contribution $\mathcal{R}_q(S^*)$ of each weight class $\mathcal{W}_q$. From each class, products are added in Step 3 by decreasing order of price until their respective contribution exceeds $\tilde{R}_q$.
- However, since $\tilde{R}_q \le \mathcal{R}_q(S^*)$, the number of products we are adding from each class $\mathcal{W}_q$ is upper-bounded by the number of products picked by S^*, and therefore, the resulting assortment must be cardinality-feasible.

6. Concluding Remarks

We conclude this paper with a number of challenging directions for future work, related to the MP-MNL choice model by itself as well as to multi-purchase settings in general.

Theory question 1: Hardness results. As mentioned in Section 2.3, establishing hardness results for unconstrained assortment optimization under the MP-MNL model remains an intriguing open question. We believe that our inapproximability bounds under totally unimodular constraints (Theorems 4 and 5) could play an important role in attaining results of this nature, even though the unconstrained setting has fewer structural restrictions to exploit.

Theory question 2: Choice probabilities. Yet another seemingly-difficult question for future research is that of proposing a polynomial-time algorithm to compute (or ϵ-estimate) choice probabilities under the MP-MNL model. We remind the reader that the exact approach we outlined in Section 2.2 is generally exponential in the number of products. Quite surprisingly, this obstacle has not prevented us from proposing a PTAS for arbitrarily-structured instances (Theorem 2), showing that additive estimation errors are sufficiently strong for this purpose.

Theory question 3: Revenue-ordered assortments. As mentioned in Section 3.3, we still do not know whether the highest earning revenue-ordered assortment garners some constant factor of the optimal expected revenue, although we were able to show that this assortment is within an $\Omega(\frac{1}{\log n})$-fraction of optimal. As such, closing this gap is an important step towards fully characterizing the proximity of the traditional MNL and MP-MNL models.

Estimating the MP-MNL model. From an empirical perspective, an intriguing direction for future work is that of finding an effective estimation procedure through which the MP-MNL model could be fit to sales data. The numerical experiments presented in Appendix G.1 depict the downsides of utilizing traditional single-purchase models in settings where multi-purchase behavior is present, yet further experimentation is needed to cement these insights. In particular, an estimation case study in which the MP-MNL model is fit to real sales data, and then benchmarked against traditional choice models, would be a natural step to consider.

Complementary multi-purchase events. Generally speaking, incorporating correlations between product purchases within the MP-MNL framework, or within any other multi-purchase model in this spirit, represents a potential direction for future work. These correlations would enable retailers to capture customers purchasing multiple complementary products. For example, in a grocery setting, it would be interesting to examine whether one can meaningfully augment the MP-MNL model to reflect the notion that shoppers who purchase pasta, for example, are then more likely to purchase some sort of accompanying sauce.

Extensions to additional choice models. Finally, a natural line for future investigation would be to tackle the question of whether alternative well-known choice models, such as the Nested Logit or Markov Chain models, can be adapted to multi-purchase settings. Clearly, any random utility maximization choice model can be extended to a multi-purchase setting, along the ranking-based instructions we consider as part of our model specification (Section 2.1). However, for any such adaptation, the obvious question is whether its resulting assortment optimization problem remains tractable, or alternatively, can be well-approximated.

Author Bios. Yicheng Bai is a PhD student in the School of Operations Research and Information Engineering at Cornell University. His research focuses on revenue management, inventory control, assortment optimization, and approximate dynamic programming. Jacob Feldman is an associate professor in the supply chain, operations and technology group at Olin Business School. Danny Segev is a senior faculty member at the School of Mathematical Sciences, Tel-Aviv University. Huseyin Topaloglu is a professor in the School of Operations Research and Information Engineering at Cornell University. His research focuses on revenue management, assortment optimization, transportation logistics, and approximate dynamic programming. . Laura Wagner is an Assistent Professor of Operations Management and Information Systems at Catolica Lisbon School of Business and Economics, Catholic University Portugal.

Acknowledgments. The work of Danny Segev on this project is supported by Israel Science Foundation grant 1407/20. The work of the fourth author was supported by a seed grant from Urban Tech Hub at Cornell Tech and National Foundation Grant CMMI-1825406.

References

Beggs S, Cardell S, Hausman J (1981) Assessing the potential demand for electric cars. *Journal of Econometrics* 17(1):1–19.

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

Bront JJM, Méndez-Díaz I, Vulcano GJ (2009) A column generation algorithm for choice-based network revenue management. *Operations Research* 57(3):769–784.

Dagsvik JK, Liu G (2006) A framework for analyzing rank ordered panel data with application to automobile demand. Technical Report, Discussion Papers 480.

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

Désir A, Goyal V, Zhang J (2022) Technical note – Capacitated assortment optimization: Hardness and approximation. *Operations Research* 70(2):893–904.

Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Science* 59(2):305–322.

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

Feldman JB, Topaloglu H (2017) Revenue management under the Markov chain choice model. *Operations Research* 65(5):1322–1342.

Ferreira K, Goh J (2021) Assortment rotation and the value of concealment. *Management Science* 67(3):1489–1507.

Fox E, Norman L, Semple J (2018) Choosing an n-pack of substitutable products. *Management Science* 64(5):2366–2379.

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

Harlam BA, Lodish LM (1995) Modeling consumers' choices of multiple items. *Journal of Marketing Research* 32(4):404–418.

Håstad J (1999) Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica* 182(1):105–142.

Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.

Hoffman AJ, Kruskal JB (1956) Integral boundary points of convex polyhedra. Kuhn HW, Tucker AW, eds., *Linear Inequalities and Related Systems*, 223–246 (Princeton University Press).

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.

Kahn BE (1995) Consumer variety-seeking among goods and services: An integrative review. *Journal of Retailing and Consumer Services* 2(3):139–148.

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

Kim J, Allenby GM, Rossi PE (2002) Modeling consumer demand for variety. *Marketing Science* 21(3):229–250.

Koop G, Poirier DJ (1994) Rank-ordered logit models: An empirical analysis of Ontario voter preferences. *Journal of Applied Econometrics* 9(4):369–388.

Li G, Rusmevichientong P, Topaloglu H (2015) The d-level Nested Logit model: Assortment and price optimization problems. *Operations Research* 63(2):325–342.

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

Mahajan S, Van Ryzin G (2001) Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49(3):334–351.

Mark DR, Lusk JL, Daniel MS (2004) Recruiting agricultural economics graduate students: Student demand for program attributes. *American Journal of Agricultural Economics* 86(1):175–184.

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.

McFadden D (1978) Modeling the choice of residential location. *Transportation Research Record* 673:72–77.

McFadden D, Train K (2000) Mixed mnl models for discrete response. *Journal of applied Econometrics* 15(5):447–470.

Plackett RL (1975) The analysis of permutations. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 24(2):193–202.

Punj GN, Staelin R (1978) The choice process for graduate business schools. *Journal of Marketing Research* 15(4):588–598.

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.

Segev D (2020) Approximation schemes for capacity-constrained assortment optimization under the Nested Logit model. In submission. Available online as SSRN report #3553264.

Simonson I (1990) The effect of purchase quantity and timing on variety-seeking behavior. *Journal of Marketing Research* 27(2):150–162.

Sumida M, Gallego G, Rusmevichientong P, Topaloglu H, Davis J (2021) Revenue-utility tradeoff in assortment optimization under the Multinomial Logit model with totally unimodular constraints. *Management Science* 67(5):2845–2869.

Talluri K, Van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1):15–33.

Tulabandhula T, Sinha D, Patidar P (2020) Multi-purchase behavior: Modeling and optimization. Available online as SSRN report #3626788.

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

Walsh JW (1995) Flexibility in consumer purchasing for uncertain future tastes. *Marketing Science* 14(2):148–165.

Williams HC (1977) On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and Planning A* 9(3):285–344.

Appendix A: Hardness under totally unimodular constraints.

An alternative way to encode the offered assortment $S \subseteq [n]$ is through its characteristic vector $x \in \{0,1\}^n$, where $x_i = 1$ if $i \in S$ and $x_i = 0$ otherwise. For an arbitrary TU-matrix A and for an integer-valued vector b, let $\mathcal{F} = \{x \in \{0,1\}^n : Ax \leq b\}$ denote the set of feasible assortment decisions. In what follows, we examine the complexity of the TU-constrained assortment optimization problem, $\max_{x \in \mathcal{F}} \mathcal{R}(x)$, whose broad applications were previously discussed in Section 1.2.

Theorem 4. *TU-constrained assortment optimization under the MP-MNL model is NP-Hard to approximate within factor $O(n^{1-\epsilon})$, for any fixed $\epsilon > 0$.*

We establish Theorem 4 by providing an approximation preserving reduction from the maximum independent set problem. The specifics of this proof can be found in Appendix A.1. Additionally, we show that TU-constrained assortment optimization remains NP-Hard even for extremely restricted settings, as argued in the following theorem.

Theorem 5. *TU-constrained assortment optimization under the MP-MNL model is NP-Hard, even when the support of M consists of only two values.*

The proof of Theorem 5 involves a reduction from the 2-partition problem, which is well known to be NP-Hard (Karp 1972). The details of this proof are provided in Appendix A.2. The key observation is that, when customers either exclusively choose according to an MNL model ($M = 1$) or an independent demand model ($M = n$), the TU-constrained assortment optimization problem can be solved in polynomial time (Hoffman and Kruskal 1956, Sumida et al. 2021). Surprisingly, the proof of Theorem 5 reveals that when we consider the "right" mixture of these two extreme cases within the MP-MNL model, its TU-constrained assortment problem is rendered NP-Hard.

A.1. Proof of Theorem 4

An instance of the independent set problem consists of an undirected graph $G=(V,E)$. A vertex set $U \subseteq V$ is called independent when, for any pair of vertices in U, there is no edge connecting the two. Our objective is to identify an independent set of maximum cardinality. The hardness result we exploit is due to Håstad (1999), who proved that, for any fixed $\epsilon > 0$, the independent set problem cannot be approximated in polynomial time on n-vertex graphs within factor $O(n^{1-\epsilon})$, unless $\mathrm{P}=\mathrm{NP}$.

Construction. Letting $V=\{v_1,\dots,v_n\}$ be the underlying set of vertices, we create an instance of the TU-constrained assortment optimization problem as follows:

- *Products*: For every $i \in [n]$, we create a block $\mathcal{B}_i = \{a_i, b_i, c_i, d_i\}$ of four distinct products, whose prices and preferences weights are listed in the next table:

product	price	weight
a_i	0	$n^{2n\cdot(4(n-i)+4)}$
b_i	n^i	$n^{2n\cdot(4(n-i)+3)}$
c_i	0	$n^{2n\cdot(4(n-i)+2)}$
d_i	0	$n^{2n\cdot(4(n-i)+1)}$

The instance we construct does not include the no-purchase option. For convenience, we denote the entire collection of products by $\mathcal{P}$.

- *Demand*: The maximum number of products M that the customer is willing to purchase is distributed over the support $0,1,3,5,\dots,2n-1$. Specifically, the probabilities of odd values are given by $\Pr[M=2i-1] = \frac{1}{n^i}$ for $i \in [n]$, whereas the residual probability mass is allocated to 0, namely, $\Pr[M=0] = 1-\sum_{i\in[n]} \frac{1}{n^i}$.
- *Totally unimodular constraints*: For every $i \in [n]$, we enforce three types of constraints:

1. When product b_i is offered, we also have to offer every product a_j with $(v_i, v_j) \in E$.
2. Exactly one of the products a_i and c_i should be offered.
3. Exactly one of the products b_i and d_i should be offered.

One can easily verify that the columns of the resulting constraint matrix can be partitioned into two sets, one corresponding to a- and b-products and the other to c- and d-products, such that: (1) Each row contains only two non-zero elements, being either 1 or -1; (2) In each row, non-zero elements of similar signs appear in different sets, whereas non-zero elements of opposite signs appear in the same set. Such matrices are well-known to be totally unimodular, by the Hoffman-Kruskal sufficient condition (1956).

Analysis. Let us first define the "nice" event $\mathcal{N}$, where product ranks satisfy

$$\begin{aligned}\mathrm{rank}(a_1,\mathcal{P}) &< \mathrm{rank}(b_1,\mathcal{P}) < \mathrm{rank}(c_1,\mathcal{P}) < \mathrm{rank}(d_1,\mathcal{P}) \\ &< \cdots < \mathrm{rank}(a_n,\mathcal{P}) < \mathrm{rank}(b_n,\mathcal{P}) < \mathrm{rank}(c_n,\mathcal{P}) < \mathrm{rank}(d_n,\mathcal{P}) \ .\end{aligned}$$

In other words, within each block $\mathcal{B}_i$, the preference order realizes as a_i, b_i, c_i, d_i, and moreover, products in lower-indexed blocks realize to be more preferable than those in higher-indexed blocks. The next claim argues that this event occurs with high probability.

LEMMA 11. $\Pr[\mathcal{N}] \ge 1 - \frac{1}{n^{n+1}}$.

Proof. The important observation is that, for any successive pair of products in the sequence $a_1, b_1, c_1, d_1, \ldots, a_n, b_n, c_n, d_n$, their preference weights differ by a multiplicative factor of precisely n^{2n}. Therefore, focusing for instance on a_1 and b_1, we have

$$\Pr[\mathrm{rank}(a_1,\mathcal{P}) > \mathrm{rank}(b_1,\mathcal{P})] = \Pr[U_{b_1} > U_{a_1}] = \frac{w_{b_1}}{w_{b_1}+w_{a_1}} = \frac{1}{1+n^{2n}} \ ,$$

and this property applies to any of these $4n-1$ successive pairs. Now, by a simple application of the union bound, $\Pr[\bar{\mathcal{N}}] \leq \frac{4n-1}{1+n^{2n}} \leq \frac{1}{n^{n+1}}$, where the last inequality can easily be verified for $n \geq 4$. $\square$

The next two claims explain how large independent sets in the graph G can be efficiently mapped to high-revenue sets of products in the resulting assortment optimization problem and vice versa.

Lemma 12. *Let $U \subseteq V$ be an independent set in G. Then, there exists a feasible assortment S with $\mathcal{R}(S) \geq (1-\frac{1}{n^{n+1}}) \cdot |U|$.*

Proof. With respect to the set of vertices U, we construct an assortment S that picks both b_i and c_i for every vertex $v_i \in U$, and concurrently, picks both a_i and d_i for $v_i \notin U$, meaning that

$$S = \{b_i : v_i \in U\} \cup \{c_i : v_i \in U\} \cup \{a_i : v_i \notin U\} \cup \{d_i : v_i \notin U\} \ .$$

This choice clearly satisfies constraints 2 and 3, for any vertex set U. To verify that constraint 1 is met as well, consider some product $b_i \in S$, which occurs only when $v_i \in U$. Since U is an independent set, for every edge $(v_i, v_j) \in E$ we must have $v_j \notin U$ and therefore $a_j \in S$, as required.

Having shown that the assortment S is feasible, we proceed to account for its expected revenue. To this end, note that the conjunction of constraints 2 and 3 guarantees that S picks exactly two products out of each block $\mathcal{B}_i$. Consequently, conditional on the nice event $\mathcal{N}$, for any offered product $b_i \in S$, since $a_i \notin S$ we know that $\mathrm{rank}(b_i, S) = 2i-1$ with probability 1. Given this observation, it follows that the expected revenue of S can be bounded as follows:

$$\begin{aligned}
\mathcal{R}(S) &= \sum_{i:b_i \in S} r_{b_i} \cdot \pi(b_i, S) \\
&= \sum_{i:b_i \in S} r_{b_i} \cdot \Pr[\mathrm{rank}(b_i, S) \leq M] \\
&\geq \Pr[\mathcal{N}] \cdot \sum_{i:b_i \in S} r_{b_i} \cdot \Pr[M \geq \mathrm{rank}(b_i, S) | \mathcal{N}] \\
&= \Pr[\mathcal{N}] \cdot \sum_{i:b_i \in S} r_{b_i} \cdot \Pr[M \geq 2i-1] \\
&\geq \Pr[\mathcal{N}] \cdot |\{i : b_i \in S\}| \\
&\geq \left(1 - \frac{1}{n^{n+1}}\right) \cdot |U| \ .
\end{aligned}$$

Here, the next-to-last inequality holds since $r_{b_i} = n^i$ and since $\Pr[M \geq 2i-1] \geq \Pr[M = 2i-1] = \frac{1}{n^i}$. The last inequality follows from Lemma 11 and from the definition of S. $\square$

Lemma 13. *Let S be a feasible assortment. Then, G has an independent set U of cardinality $|U| \geq (1-\frac{3}{n}) \cdot (\mathcal{R}(S) - 3)$.*

Proof. We begin by establishing upper bounds on the choice probability of each product $b_i \in S$, depending on two scenarios:

- *Case I: $a_i \notin S$.* Following the arguments discussed within the proof of Lemma 12, we know that $[\mathrm{rank}(b_i, S) \mid \mathcal{N}] = 2i - 1$ and therefore

$$\Pr[M \geq \mathrm{rank}(b_i, S) | \mathcal{N}] = \Pr[M \geq 2i - 1] = \sum_{k=i}^{n} \frac{1}{n^k} \leq \frac{1}{n^i} \cdot \left(1 + \frac{2}{n}\right) \ .$$

Consequently, by utilizing Lemma 11,

$$\begin{aligned} \pi(b_i, S) &= \Pr[\mathcal{N}] \cdot \Pr[M \geq \mathrm{rank}(b_i, S) | \mathcal{N}] + \Pr[\bar{\mathcal{N}}] \cdot \Pr[M \geq \mathrm{rank}(b_i, S) | \bar{\mathcal{N}}] \\ &\leq \frac{1}{n^i} \cdot \left(1 + \frac{2}{n}\right) + \frac{1}{n^{n+1}} \\ &\leq \frac{1}{n^i} \cdot \left(1 + \frac{3}{n}\right) \ . \end{aligned}$$

- *Case II: $a_i \in S$.* In this setting, since the assortment S picks exactly two products out of each block, conditional on the nice event $\mathcal{N}$, we know that $\mathrm{rank}(b_i, S) = 2i$ with probability 1, since $a_i \in S$ by the case hypothesis. As a result,

$$\Pr[M \geq \mathrm{rank}(b_i, S) | \mathcal{N}] = \Pr[M \geq 2i] = \sum_{k=i+1}^{n} \frac{1}{n^k} \leq \frac{2}{n^{i+1}} \ .$$

Therefore, again by Lemma 11,

$$\begin{aligned} \pi(b_i, S) &= \Pr[\mathcal{N}] \cdot \Pr[M \geq \mathrm{rank}(b_i, S) | \mathcal{N}] + \Pr[\bar{\mathcal{N}}] \cdot \Pr[M \geq \mathrm{rank}(b_i, S) | \bar{\mathcal{N}}] \\ &\leq \frac{2}{n^{i+1}} + \frac{1}{n^{n+1}} \\ &\leq \frac{3}{n^{i+1}} \ . \end{aligned}$$

Now, since a-, c-, and d-products are not contributing any revenue, the expected revenue of S can be upper bounded by

$$\begin{aligned} \mathcal{R}(S) &= \sum_{i: b_i \in S} r_{b_i} \cdot \pi(b_i, S) \\ &\leq |\{i : [b_i \in S] \wedge [a_i \notin S]\}| \cdot \left(1 + \frac{3}{n}\right) + |\{i : [b_i \in S] \wedge [a_i \in S]\}| \cdot \frac{3}{n} \\ &\leq |\{i : [b_i \in S] \wedge [a_i \notin S]\}| \cdot \left(1 + \frac{3}{n}\right) + 3 \ , \end{aligned} \tag{8}$$

where the first inequality follows from cases I and II. The crucial observation is that $U = \{v_i \in V : [b_i \in S] \wedge [a_i \notin S]\}$ is necessarily an independent set in G. To verify this claim, note that had there been an edge $(v_i, v_j) \in E$ for some pair of vertices $v_i \neq v_j \in U$, constraint 1 will stipulate in particular that $a_j \in S$ (due to having $b_i \in S$), contradicting the fact that $v_j \in U$. To conclude the proof, note that by rearranging inequality (8), we have $|U| \geq (1 - \frac{3}{n}) \cdot (\mathcal{R}(S) - 3)$. $\square$

Summary. Given these relationships, we argue that an α_n-approximation for the TU-constrained assortment optimization problem readily translates to an $(\alpha_n - O(1))$-approximation for the independent set problem on n-vertex graphs. For this purpose, let U^* be a maximum-cardinality independent set in G. By Lemma 12, there exists a feasible assortment S_{U^*} with $\mathcal{R}(S_{U^*}) \geq (1 - \frac{1}{n^{n+1}}) \cdot |U^*|$, meaning that an α_n-approximation would identify an assortment S with $\mathcal{R}(S) \geq \alpha_n \cdot \mathcal{R}(S_{U^*}) \geq \alpha_n \cdot (1 - \frac{1}{n^{n+1}}) \cdot |U^*|$. Translating the latter back to an independent set U via Lemma 13, we obtain an independent set U of cardinality

$$\begin{aligned}|U| &\geq \left(1 - \frac{3}{n}\right) \cdot (\mathcal{R}(S) - 3) \\ &\geq \left(1 - \frac{3}{n}\right) \cdot \left(\alpha_n \cdot \left(1 - \frac{1}{n^{n+1}}\right) \cdot |U^*| - 3\right) \\ &= (\alpha_n - O(1)) \cdot |U^*| \ .\end{aligned}$$

A.2. Proof of Theorem 5

Our proof relies on a reduction from the 2-partition problem, whose input consists of a collection of n integers $a_1, \ldots, a_n$; our goal is to decide whether these numbers can be partitioned into two sets of equal sum. Namely, letting $\sum_{i \in [n]} a_i = L$, the decision problem in question is whether there exists a subset $I \subseteq [n]$ such that $\sum_{i \in I} a_i = L/2$. Prior to moving forward, we assume without loss of generality that $L \geq 4$, and introduce two additional pieces of notation. We let $a_{\max} = \max_{i \in [n]} a_i$ and $\hat{a}_i = a_i / L$. It is easy to verify that $\sum_{i \in I} a_i = L/2$ is equivalent to $\sum_{i \in I} \hat{a}_i = 1/2$.

Given such an instance, we create a corresponding instance of TU-constrained assortment optimization under the MP-MNL model. To this end, we will not be making use of a no-purchase option, and the set of $4n + 1$ underlying products will be comprised of five types:

- *Product type 1:* For each $i \in [n]$, we introduce a single product, with weight $w_i^1 = \frac{3\hat{a}_i}{L - \hat{a}_i}$ and price $r_i^1 = L - \hat{a}_i$.
- *Product type 2:* For each $i \in [n]$, we introduce a single product, with weight $w_i^2 = \hat{a}_i - \frac{3\hat{a}_i}{L - \hat{a}_i}$ and price $r_i^2 = 0$. Here, it is important to note that w_i^2 is indeed non-negative, since $L \geq 4$ and $\hat{a}_i \leq 1$.
- *Product type 3:* For each $i \in [n]$, we introduce a single product, with weight $w_i^3 = \frac{\hat{a}_i}{L}$ and price $r_i^3 = L$.
- *Product type 4:* For each $i \in [n]$, we introduce a single product, with weight $w_i^4 = \hat{a}_i + \frac{\hat{a}_i}{L}$ and price $r_i^4 = 0$.
- *Product type 5:* In this case, we introduce only one product, with weight $w^5 = 2 - \frac{1}{L}$ and price $r^5 = 0$.

From a demand standpoint, the number of products a customer is willing to purchase will be limited to taking only the values 1 and $4n + 1$, with probabilities $\Pr[M = 1] = 9/11$ and $\Pr[M = 4n + 1] = 2/11$. These two extremes correspond to customers who either choose according to a traditional MNL model or purchase all offered products with certainty.

To express our TU constraints, we use superscripts to index across product types. As such, any assortment decision can be captured through the binary vector $x = (x_i^t)$, where t stands for the product type and i is its index for that particular type. With this notation, we enforce the following set of TU constraints:

- $x_i^1 = x_i^2 = x_i^4$ and $x_i^3 = 1 - x_i^1$, for every $i \in [n]$.
- $x^5 = 1$.

One can easily verify that the columns of the resulting constraint matrix can be partitioned into two sets, one corresponding to product types 1, 2, and 4 and the other to product types 3 and 5, such that: (1) Each row contains at most two non-zero elements, being either 1 or -1; (2) In each row, non-zero elements of similar signs appear in different sets, whereas non-zero elements of opposite signs appear in the same set. Such matrices are well-known to be totally unimodular, by the Hoffman-Kruskal sufficient condition (1956).

Given the TU-constraint structure defined above, we can fully specify any feasible assortment by simply specifying the products selected among product type 1. As such, we drop superscripts and simply denote feasible assortment via the characteristic vector of type 1 products, $x \in \{0,1\}^n$. In this case, the expected revenue of any feasible assortment is

$$\mathcal{R}(x) = \frac{9}{11} \cdot \frac{1 + 2\sum_{i\in[n]} \hat{a}_i x_i}{2 + 2\sum_{i\in[n]} \hat{a}_i x_i} + \frac{2}{11} \cdot \left(nL - \sum_{i\in[n]} \hat{a}_i x_i \right) . \tag{9}$$

We argue that there exists a subset $I \subseteq [n]$ for which $\sum_{i\in I} \hat{a}_i = 1/2$ if and only if there is a vector $x \in \{0,1\}^n$ for which $\mathcal{R}(x) \geq \frac{2}{11} \cdot nL + \frac{5}{11}$. For this purpose, note that:

- Suppose $I \subseteq [n]$ is a subset with $\sum_{i\in I} \hat{a}_i = 1/2$. Then, for the characteristic vector x^I, we have $\mathcal{R}(x^I) = \frac{2}{11} \cdot nL + \frac{5}{11}$.
- Conversely, suppose that $\sum_{i\in I} \hat{a}_i \neq 1/2$ for every subset $I \subseteq [n]$. In the case, let us define the auxiliary function $f : [0,1] \to \mathbb{R}$, given by $f(z) = \frac{9}{11} \cdot \frac{1+2z}{2+2z} - \frac{2}{11} z$. Based on the revenue representation (9), one way of writing the expected revenue of any vector $x \in \{0,1\}^n$ is by observing that $\mathcal{R}(x) = \frac{2}{11} \cdot nL + f(\sum_{i\in[n]} \hat{a}_i x_i)$. However, basic calculus reveals that $z^* = 1/2$ is the unique maximizer of $f(z)$ over $[0,1]$, where $f(z^*) = \frac{5}{11}$. Since $\sum_{i\in[n]} \hat{a}_i x_i \neq 1/2$, it follows that $\mathcal{R}(x) < \frac{2}{11} \cdot nL + \frac{5}{11}$.

Appendix B: Proofs from Section 3

B.1. Proofs of Lemmas 2 and 3

Intermediate claims. Prior to diving into the detailed proofs of Lemmas 2 and 3, we establish two intermediate claims. Recall that, for $m \in [n]_0$, we have defined $\mathcal{R}_m(S)$ to be the expected revenue earned, given that customers are willing to purchase up to m products; in other words, $\mathcal{R}_m(S)$ is the conditional expected revenue $[\mathcal{R}(S)|M = m]$. The following claim, whose proof appears in Appendix B.3, provides a recursive representation for $\mathcal{R}_m(S)$.

Claim 4. *For any assortment $S \subseteq [n]$ and $m \in [n]$, we have*

$$\mathcal{R}_m(S) = \sum_{i\in S} \frac{w_i}{1 + w(S)} \cdot (r_i + \mathcal{R}_{m-1}(S_{-i})) \ .$$

The second claim we establish provides an upper bound on the marginal gain in expected revenue that can result when the number of products a customer is willing to purchase is augmented by one. More specifically, this upper bound is z^*, the optimal MNL-based expected revenue as defined in Property 3.1. The proof of Claim 5 is presented in Appendix B.4.

Claim 5. *For any assortment $S \subseteq [n]$ and $m \in [n]$, we have $\mathcal{R}_m(S) - \mathcal{R}_{m-1}(S) \leq z^*$.*

Proof of Lemma 2. We establish the claim by showing that for any assortment $S \subseteq [n]$, product $\ell \in \tilde{S} \setminus S$, and $m \in [n]_0$, we have $\mathcal{R}_m(S_{+\ell}) \geq \mathcal{R}_m(S)$. Consequently, the expected revenue $\mathcal{R}(S)$ of any assortment S can only improve by adding any product from $\tilde{S} \setminus S$, meaning in particular that any maximum-cardinality optimal assortment necessarily contains $\tilde{S}$.

We prove this result via induction over m, noting that the base case of $m = 0$ trivially holds. In the general case of $m \geq 1$, by applying Claim 4 to express $\mathcal{R}_m(S_{+\ell})$, we get

$$\begin{aligned}
\mathcal{R}_m(S_{+\ell}) &= \sum_{i \in S} \frac{w_i}{1 + w(S_{+\ell})} \cdot (r_i + \mathcal{R}_{m-1}(S_{+\ell-i})) + \frac{w_\ell}{1 + w(S_{+\ell})} \cdot (r_\ell + \mathcal{R}_{m-1}(S)) \\
&\geq \sum_{i \in S} \frac{w_i}{1 + w(S_{+\ell})} \cdot (r_i + \mathcal{R}_{m-1}(S_{-i})) + \frac{w_\ell}{1 + w(S_{+\ell})} \cdot (r_\ell + \mathcal{R}_{m-1}(S)) \\
&= \left(1 - \frac{w_\ell}{1 + w(S_{+\ell})}\right) \cdot \sum_{i \in S} \left(\frac{w_i}{1 + w(S)} \cdot (r_i + \mathcal{R}_{m-1}(S_{-i}))\right) \\
&\quad + \frac{w_\ell}{1 + w(S_{+\ell})} \cdot (r_\ell + \mathcal{R}_{m-1}(S)) \\
&= \left(1 - \frac{w_\ell}{1 + w(S_{+\ell})}\right) \cdot \mathcal{R}_m(S) + \frac{w_\ell}{1 + w(S_{+\ell})} \cdot (r_\ell + \mathcal{R}_{m-1}(S)) \\
&= \mathcal{R}_m(S) + \frac{w_\ell}{1 + w(S_{+\ell})} \cdot (r_\ell + \mathcal{R}_{m-1}(S) - \mathcal{R}_m(S)) \\
&\geq \mathcal{R}_m(S) \ .
\end{aligned}$$

Here, the first inequality holds since $\mathcal{R}_{m-1}(S_{+\ell-i}) \geq \mathcal{R}_{m-1}(S_{-i})$ by the induction hypothesis, while the second equality follows by noting that $\frac{w_i}{1+w(S_{+\ell})} = \frac{w_i}{1+w(S)} \cdot (1 - \frac{w_\ell}{1+w(S_{+\ell})})$. The third equality is again implied by Claim 4. The last inequality follows by observing that $\mathcal{R}_m(S) - \mathcal{R}_{m-1}(S) \leq z^*$, due to Claim 5, and by noting that $r_\ell \geq z^*$, since $\ell \in \tilde{S} = \{i \in [n] : r_i \geq z^*\}$.

Proof of Lemma 3. We prove the claim by showing that $\mathcal{R}_m(S) \leq m z^*$ for any assortment $S \subseteq [n]$ and $m \in [n]_0$, via induction over m. Given this bound, we immediately get

$$\begin{aligned}
\mathcal{R}(S^*) &= \sum_{m \in [n]} \Pr[M = m] \cdot \mathcal{R}_m(S^*) \\
&\leq z^* \cdot \sum_{m \in [n]} \Pr[M = m] \cdot m \\
&\leq \mathbb{E}[M] \cdot \mathcal{R}(\tilde{S}) \ ,
\end{aligned}$$

where the last inequality holds since $\mathcal{R}_m(S)$ is decreasing in m, and therefore

$$\mathcal{R}(\tilde{S}) = \sum_{m \in [n]} \Pr[M = m] \cdot \mathcal{R}_m(\tilde{S}) \geq \mathcal{R}_1(\tilde{S}) = z^* \ .$$

To argue that $\mathcal{R}_m(S) \leq m z^*$ for any assortment $S \subseteq [n]$, first note that the base case of $m = 0$ trivially holds. For the general case of $m \geq 1$, by Claim 4 we have

$$\begin{aligned}
\mathcal{R}_m(S) &= \sum_{i \in S} \frac{w_i}{1 + w(S)} \cdot (r_i + \mathcal{R}_{m-1}(S_{-i})) \\
&\leq \mathcal{R}_1(S) + (m-1) \cdot z^* \cdot \sum_{i \in S} \frac{w_i}{1 + w(S)} \\
&\leq z^* + (m-1) \cdot z^* \cdot \sum_{i \in S} \frac{w_i}{1 + w(S)} \\
&\leq m z^* \ ,
\end{aligned}$$

where the first inequality holds since $\mathcal{R}_{m-1}(S_{-i}) \leq (m-1) \cdot z^*$ by the induction hypothesis, and the second inequality follows by noting that $\mathcal{R}_1(S) \leq \mathcal{R}_1(\tilde{S}) = z^*$.

B.2. Proof of Lemma 4

At a high level, our proof begins by describing a way to partition the products offered by the optimal assortment S^* into $O(\log n)$ groups, based on their contribution to the overall expected number of products purchased under S^*. For each such group, we will argue that there exists a revenue-ordered assortment that garners at least half of the group's revenue, thus directly yielding the desired performance guarantee.

Notation. We assume throughout this section that the underlying products are indexed in non-increasing order of revenue, i.e, $r_1 \geq \cdots \geq r_n$. For each product $i \in S^*$, let $S^*_{\geq i} = \{\ell \in S^* : \ell \geq i\}$ be the set of products included in the optimal assortment whose index is lower-bounded by i, noting that their revenue is upper-bounded by r_i. Moreover, for any assortment $S \subseteq [n]$, let $\mathcal{D}(S) = \sum_{i \in S} \pi(i, S)$ denote its expected total demand, or put differently, the expected number of products purchased when offering the assortment S.

Partitioning of S^*. Our method for partitioning S^* into $O(\log n)$ groups is carried out via the following iterative procedure. Let $i_1 = \max\{i \in S^* : \mathcal{D}(S^*_{\geq i}) \geq \frac{1}{2} \cdot \mathcal{D}(S^*)\}$, and define the first group to be $G_1 = S^*_{\geq i_1}$. From here, for $q \geq 2$, we let $i_q = \max\{i \in S^* : \mathcal{D}(S^*_{\geq i_q} \setminus S^*_{\geq i_{q-1}}) \geq \frac{1}{2} \cdot \mathcal{D}(S^* \setminus S^*_{\geq i_{q-1}})\}$, and define the q-th group as $G_q = S^*_{\geq i_q} \setminus S^*_{\geq i_{q-1}}$; these definitions are applied as long as $\mathcal{D}(S^* \setminus S^*_{\geq i_{q-1}}) \geq \frac{1}{2n}$. Upon reaching a stage q where $\mathcal{D}(S^* \setminus S^*_{\geq i_{q-1}}) < \frac{1}{2n}$, we conclude our partition by setting $G_q = S^* \setminus S^*_{\geq i_{q-1}}$. Assuming that the resulting groups are $G_1, \dots, G_Q$, it is easy to see that they form a partition of the optimal assortment S^*. Furthermore, we have $Q = O(\log n)$, since initially $\mathcal{D}(S^*) \leq n$, and in each stage we halve the remaining expected demand up until the remaining products in S^* yield fewer than $\frac{1}{2n}$ expected purchases.

Completing the proof. Since our expected revenue function $\mathcal{R}(\cdot)$ is subadditive, we have

$$\mathcal{R}(S^*) \leq \sum_{q \in [Q]} \mathcal{R}(G_q) \leq Q \cdot \max_{q \in [Q]} \mathcal{R}(G_q) = O(\log n) \cdot \max_{q \in [Q]} \mathcal{R}(G_q) \ .$$

Consequently, to show that revenue-ordered assortments can garner an $\Omega(\frac{1}{\log n})$-factor of the optimal expected revenue $\mathcal{R}(S^*)$, it suffices to prove the following claim, showing that for any group G_q, the most profitable revenue-ordered assortment S^*_{ro} attains an expected revenue of $\Omega(1) \cdot \mathcal{R}(G_q)$.

Claim 6. *$\mathcal{R}(S^*_{\text{ro}}) \geq \frac{1}{2} \cdot \mathcal{R}(G_q)$, for any $q \in [Q]$.*

Proof. For any given $q \in [Q]$, we construct a revenue-ordered assortment S_q, whose expected revenue satisfies $\mathcal{R}(S_q) \geq \frac{1}{2} \cdot \mathcal{R}(G_q)$. Since $\mathcal{R}(S^*_{\text{ro}}) \geq \max_{q \in [Q]} \mathcal{R}(S_q)$ by definition, this completes the proof. We proceed by considering two cases, depending on the value of q.

- *Case 1:* $q \in [Q-1]$. In this case, we set $S_q = [i_q]$, which is clearly a revenue-ordered assortment, due to our initial assumption that $r_1 \geq \cdots \geq r_n$. The expected revenue of this assortment can be related to that of G_q by observing that

$$\begin{aligned}
\mathcal{R}(S_q) &\geq r_{i_q} \cdot \mathcal{D}(S_q) \\
&\geq \frac{1}{2} \cdot r_{i_q} \cdot \mathcal{D}(S^* \setminus S^*_{\geq i_{q-1}}) \\
&\geq \frac{1}{2} \cdot r_{i_q} \cdot \mathcal{D}(G_q) \\
&\geq \frac{1}{2} \cdot \mathcal{R}(G_q).
\end{aligned}$$

Here, the first and last inequalities follow by noting that product i_q has the smallest revenue within S_q and the largest revenue within G_q. The second inequality holds since $\mathcal{D}(S_q) \geq \frac{1}{2} \cdot \mathcal{D}(S^* \setminus S^*_{\geq i_{q-1}})$, by definition of i_q, whereas the third inequality is obtained by noting that $G_q \subseteq S^* \setminus S^*_{\geq i_{q-1}}$.

- *Case 2:* $q = Q$. In this case, we set $S_Q = \tilde{S}$, which is precisely the optimal (single-purchase) MNL-based assortment, as defined in Property 3.1. In order to relate between the expected revenues of S_Q and G_Q, we begin by establishing the following upper bound on $\mathcal{R}(G_Q)$:

$$
\begin{aligned}
\mathcal{R}(G_Q) &= \sum_{m \in [n]} \Pr[M = m] \cdot \mathcal{R}_m(G_Q) \\
&\leq \mathcal{R}_n(G_Q) \\
&\leq \sum_{i \in G_Q} \mathcal{R}_n(\{i\}) \\
&= \sum_{i \in G_Q} \frac{r_i w_i}{1 + w_i} \\
&\leq \sum_{i \in G_Q} r_i w_i,
\end{aligned} \tag{10}
$$

where the first inequality holds since $\mathcal{R}_n(S) \geq \mathcal{R}_m(S)$ for any assortment S and $m \leq n$. On the other hand, we observe that

$$
\begin{aligned}
\mathcal{R}(\tilde{S}) &\geq \mathcal{R}_1(\tilde{S}) \\
&\geq \mathcal{R}_1(G_Q) \\
&= \sum_{i \in G_Q} \frac{r_i w_i}{1 + w(G_Q)} \\
&\geq \frac{1}{2} \cdot \sum_{i \in G_Q} r_i w_i \\
&\geq \frac{1}{2} \cdot \mathcal{R}(G_Q).
\end{aligned}
$$

Here, the second inequality follows by definition of $\tilde{S}$, while the last inequality is exactly (10). To obtain the third inequality, note that we must have $w(G_Q) \leq \frac{1}{n}$, since otherwise,

$$
\mathcal{D}(G_Q) \geq \sum_{i \in G_Q} \pi_1(i, G_Q) = \frac{w(G_Q)}{1 + w(G_Q)} > \frac{1/n}{1 + 1/n} \geq \frac{1}{2n},
$$

which contradicts the definition of the final group G_Q.

B.3. Proof of Claim 4

To establish this claim, we observe that

$$
\begin{aligned}
\mathcal{R}_m(S) &= \sum_{i \in S} r_i \cdot \pi_m(i, S) \\
&= \sum_{i \in S} r_i \cdot \left(\frac{w_i}{1 + w(S)} + \sum_{j \in S_{-i}} \frac{w_j}{1 + w(S)} \cdot \pi_{m-1}(i, S_{-j}) \right) \\
&= \sum_{i \in S} r_i \cdot \frac{w_i}{1 + w(S)} + \sum_{j \in S} \frac{w_j}{1 + w(S)} \cdot \left(\sum_{i \in S_{-j}} r_i \cdot \pi_{m-1}(i, S_{-j}) \right)
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_{i \in S} r_i \cdot \frac{w_i}{1+w(S)} + \sum_{j \in S} \frac{w_j}{1+w(S)} \cdot \mathcal{R}_{m-1}(S_{-j}) \\
&= \sum_{i \in S} \frac{w_i}{1+w(S)} \cdot \left(r_i + \mathcal{R}_{m-1}\left(S_{-i}\right)\right),
\end{aligned}
$$

where the second equality follows from Lemma 1.

B.4. Proof of Claim 5

We prove the required bound by induction over m. For $m=1$ and for an arbitrary assortment $S \subseteq [n]$, we have $\mathcal{R}_1(S) - \mathcal{R}_0(S) = \mathcal{R}_1(S) \leq z^*$, where the last inequality follows since z^*, as defined in Property 3.1, can be equivalently represented as $\max_{S \subseteq [n]} \mathcal{R}_1(S)$. Next, for $m \geq 2$ and for an arbitrary assortment $S \subseteq [n]$,

$$
\begin{aligned}
\mathcal{R}_m(S) - \mathcal{R}_{m-1}(S) &= \sum_{i \in S} \frac{w_i}{1+w(S)} \cdot \left(\mathcal{R}_{m-1}(S_{-i}) - \mathcal{R}_{m-2}(S_{-i})\right) \\
&\leq \sum_{i \in S} \frac{w_i}{1+w(S)} \cdot z^* \\
&\leq z^*,
\end{aligned}
$$

where the equality above is obtained by applying Claim 4, and the first inequality follows from the induction hypothesis.

Appendix C: Proofs from Section 4

C.1. Proof of Claim 1

To obtain a lower bound on $\pi_m(i, S_{-j})$, note that, by conditioning on the random utility of product ℓ, we have

$$
\begin{aligned}
\pi_m(i, S_{-j}) &= \int_{\mathbb{R}} \Pr\left[[\mathrm{rank}(i, S_{-j}) \leq m] \wedge [U_i > U_0] \mid U_\ell = x\right] dF_{U_\ell}(x) \\
&= \int_{\mathbb{R}} \Pr\left[[\mathrm{rank}(i, S_{-\ell}) \leq m] \wedge [U_i > U_0] \mid U_j = x\right] dF_{U_\ell}(x) \\
&\geq \int_{\mathbb{R}} \Pr\left[[\mathrm{rank}(i, S_{-\ell}) \leq m] \wedge [U_i > U_0] \mid U_j = x\right] dF_{U_j}(x) \\
&= \pi_m(i, S_{-\ell}),
\end{aligned}
$$

where the inequality hold since $U_j \succeq_{\mathrm{st}} U_\ell$, along with the fact that $\Pr[[\mathrm{rank}(i, S_{-\ell}) \leq m] \wedge [U_i > U_0] \mid U_j = x]$ is decreasing in x.

C.2. Proof of Claim 2

To prove the desired lower bound, we observe that

$$
\begin{aligned}
w(Q \cap \mathcal{L}) &= w(S \cap \mathcal{L}) - w((S \setminus Q) \cap \mathcal{L}) \\
&\geq \frac{W}{(1+\delta)^2} - |(S \setminus Q) \cap \mathcal{L}| \cdot \max_{i \in \mathcal{L}} w_i \\
&\geq \frac{W}{(1+\delta)^2} - (|S \setminus Q| - |\mathcal{H}^* \setminus Q|) \cdot \frac{w(\mathcal{H}^* \cap Q)}{|\mathcal{H}^* \cap Q|} \\
&\geq \frac{W}{(1+\delta)^2} - (m_{\max} - |\mathcal{H}^* \setminus Q|) \cdot \frac{w(\mathcal{H}^* \cap Q)}{|\mathcal{H}^* \cap Q|} \\
&= \frac{W}{(1+\delta)^2} - \frac{m_{\max} - |\mathcal{H}^* \setminus Q|}{|\mathcal{H}^*| - |\mathcal{H}^* \setminus Q|} \cdot w(\mathcal{H}^* \cap Q)
\end{aligned}
$$

$$\begin{aligned} &\geq \frac{W}{(1+\delta)^2} - \frac{m_{\max}}{|\mathcal{H}^*|} \cdot w(\mathcal{H}^* \cap Q) \\ &= \frac{W}{(1+\delta)^2} - \delta \cdot w(\mathcal{H}^* \cap Q) \ . \end{aligned}$$

Here, the first inequality holds since $w(S \cap \mathcal{L}) \geq \frac{W}{(1+\delta)^2}$, as $S \in \mathcal{F}$. The second inequality is obtained by noting that the weight of any light product is upper bounded by the weight of any product in $\mathcal{H}^* \setminus Q$, and therefore, also by the average weight of these products. The third inequality follows by recalling that $|S \setminus Q| \leq m_{\max}$. Finally, the last equality holds since $|\mathcal{H}^*| = \mathcal{N} = \frac{6m_{\max}^2}{\epsilon}$ and $\delta = \frac{\epsilon}{6m_{\max}}$.

C.3. Proof of Claim 3: Lower bound

Again, the base case of $m = 0$ trivially holds, and we proceed to consider the general case of $m \geq 1$ by relating the three parts of (7) and (4).

Part (i). First, note that

$$\begin{aligned} w(Q) &= w(\mathcal{L} \cap Q) + w(\mathcal{H}^* \cap Q) \\ &\geq \frac{W}{(1+\delta)^2} + (1-\delta) \cdot w(H) \\ &\geq \frac{1}{(1+\delta)^3} \cdot ((1+\delta) \cdot W + w(H)) \ , \end{aligned}$$

where the first inequality follows from Claim 2. Therefore, we have

$$\underbrace{\frac{1}{1 + (1+\delta) \cdot W + w(H)}}_{\text{part (i) of (4)}} \geq (1-\delta)^3 \cdot \frac{1}{1 + w(Q)} \ . \tag{11}$$

Part (ii). In this case, a direct application of the induction hypothesis yields $V(i, m-1, H) \geq (1-\delta)^{6(m-1)} \cdot \pi_{m-1}(i, Q_{-\ell})$ for any product $\ell \in \mathcal{L} \cap Q_{-i}$. As such, we have

$$\begin{aligned} &\underbrace{w_i + \left[\frac{W}{(1+\delta)^2} - \delta \cdot w(H_{-i})\right]^+ \cdot V(i, m-1, H)}_{\text{part (ii) of (4)}} \\ &\quad \geq w_i + (1-\delta)^{6(m-1)} \cdot \left[\frac{W}{(1+\delta)^2} - \delta \cdot w(H_{-i})\right]^+ \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell}) \\ &\quad \geq w_i + (1-\delta)^{6(m-1)} \cdot \left(\frac{W}{(1+\delta)^2} \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell}) - \delta \cdot w(H_{-i}) \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell})\right) \\ &\quad \geq w_i + (1-\delta)^{6(m-1)} \cdot \left(\frac{1}{(1+\delta)^3} \cdot \sum_{\ell \in \mathcal{L} \cap Q_{-i}} w_\ell \cdot \pi_{m-1}(i, Q_{-\ell}) - \delta \cdot w(H_{-i}) \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell})\right) \\ &\quad \geq (1-\delta)^{6(m-1)} \cdot \left(w_i + \frac{1}{(1+\delta)^3} \cdot \sum_{\ell \in \mathcal{L} \cap Q_{-i}} w_\ell \cdot \pi_{m-1}(i, Q_{-\ell}) - \delta \cdot \sum_{j \in \mathcal{H}^* \cap Q_{-i}} w_j \cdot \pi_{m-1}(i, Q_{-j})\right) \ . \end{aligned} \tag{12}$$

Here, the third inequality holds since

$$\begin{aligned} \sum_{\ell \in \mathcal{L} \cap Q_{-i}} w_\ell \cdot \pi_{m-1}(i, Q_{-\ell}) &\leq w(\mathcal{L} \cap Q_{-i}) \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell}) \\ &\leq w(\mathcal{L} \cap S) \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell}) \\ &\leq (1+\delta) \cdot W \cdot \max_{\ell \in \mathcal{L} \cap Q_{-i}} \pi_{m-1}(i, Q_{-\ell}) \ . \end{aligned}$$

On the other hand, the fourth inequality is obtained by noting that

$$w\left(H_{-i}\right)\cdot\max_{\ell\in\mathcal{L}\cap Q_{-i}}\pi_{m-1}\left(i,Q_{-\ell}\right) = \left(\sum_{j\in\mathcal{H}^*\cap Q_{-i}} w_j\right)\cdot\max_{\ell\in\mathcal{L}\cap Q_{-i}}\pi_{m-1}\left(i,Q_{-\ell}\right)$$
$$\leq \sum_{j\in\mathcal{H}^*\cap Q_{-i}} w_j\cdot\pi_{m-1}(i,Q_{-j}) \ ,$$

where the latter inequality holds since $w_j \geq w_\ell$ for any pair of products $j\in\mathcal{H}$ and $\ell\in\mathcal{L}$, implying that $\pi_{m-1}(i,Q_{-j}) \geq \max_{\ell\in\mathcal{L}\cap Q_{-i}}\pi_{m-1}(i,Q_{-\ell})$ by Claim 1.

Part (iii). Again, a direct application of the induction hypothesis yields $V(i,m-1,H_{-j}) \geq (1-\delta)^{6(m-1)}\cdot\pi_{m-1}(i,Q_{-j})$ for any product $j\in H_{-i}$. Consequently,

$$\underbrace{\sum_{j\in H_{-i}} w_j\cdot V\left(i,m-1,H_{-j}\right)}_{\text{part (iii) of (4)}} \geq (1-\delta)^{6(m-1)}\cdot\sum_{j\in\mathcal{H}^*\cap Q_{-i}} w_j\cdot\pi_{m-1}\left(i,Q_{-j}\right) \ . \tag{13}$$

Putting things together. We are now ready to conclude our analysis by combining parts (i), (ii), and (iii). Specifically, taking the lower bounds provided in (11), (12), and (13) yields

$$\begin{aligned} V(i,m,H) &\geq (1-\delta)^{6m-3}\cdot\frac{1}{1+w(Q)} \\ &\quad\cdot\left(w_i+\frac{1}{(1+\delta)^3}\cdot\sum_{\ell\in\mathcal{L}\cap Q_{-i}} w_\ell\cdot\pi_{m-1}(i,Q_{-\ell})+(1-\delta)\cdot\sum_{j\in\mathcal{H}^*\cap Q_{-i}} w_j\cdot\pi_{m-1}(i,Q_{-j})\right) \\ &\geq (1-\delta)^{6m}\cdot\pi(i,Q) \ . \end{aligned}$$

Appendix D: Proofs from Section 5

D.1. Proof of Lemma 6

In order to prove that $\pi^{\langle w^{\downarrow i}\rangle}(i,S) \geq (1-\epsilon)\cdot\pi^{\langle w\rangle}(i,S)$, we show that $\pi_m^{\langle w^{\downarrow i}\rangle}(i,S) \geq (1-\epsilon)\cdot\pi_m^{\langle w\rangle}(i,S)$ for all m by induction. The base case of $m=0$ trivially holds with equality. Now, for $m\geq 1$, by Claim 1 we have

$$\begin{aligned} \pi_m^{\langle w^{\downarrow i}\rangle}(i,S) &= \frac{w_i^{\downarrow i}}{1+w_i^{\downarrow i}+w(S_{-i})}+\sum_{j\in S_{-i}}\frac{w_j}{1+w_i^{\downarrow i}+w(S_{-i})}\cdot\pi_{m-1}^{\langle w^{\downarrow i}\rangle}(i,S_{-j}) \\ &\geq (1-\epsilon)\cdot\frac{w_i}{1+w_i+w(S_{-i})}+\sum_{j\in S_{-i}}\frac{w_j}{1+w_i+w(S_{-i})}\cdot\pi_{m-1}^{\langle w^{\downarrow i}\rangle}(i,S_{-j}) \\ &\geq (1-\epsilon)\cdot\frac{w_i}{1+w_i+w(S_{-i})}+(1-\epsilon)\cdot\sum_{j\in S_{-i}}\frac{w_j}{1+w_i+w(S_{-i})}\cdot\pi_{m-1}^{\langle w\rangle}(i,S_{-j}) \\ &= (1-\epsilon)\cdot\pi_m^{\langle w\rangle}(i,S) \ , \end{aligned}$$

where the first inequality holds since $w_i^{\downarrow i}\in[(1-\epsilon)\cdot w_i, w_i]$ and the second inequality follows from the induction assumption.

To prove that $\pi^{\langle w^{\downarrow i}\rangle}(j,S) \geq \pi^{\langle w\rangle}(j,S)$, for every product $j\neq i$, it suffices to show that $\pi_m^{\langle w^{\downarrow i}\rangle}(j,S) \geq \pi_m^{\langle w\rangle}(j,S)$ for all m. To this end, let us consider the assortment $S^+ = S\cup\{i^+\}$, where i^+ is an auxiliary product with preference weight $w_i^{\downarrow i}$. Given this definition, for every product $j\neq i$ we have

$$\begin{aligned} \pi_m^{\langle w\rangle}(j,S) &= \pi_m^{\langle w\rangle}\left(j,S^+_{-i^+}\right) \\ &\leq \pi_m^{\langle w\rangle}\left(j,S^+_{-i}\right) \\ &= \pi_m^{\langle w^{\downarrow i}\rangle}(j,S) \ , \end{aligned}$$

where the inequality above follows from instantiating Claim 1 with respect to the assortment S^+ and the pair of products $i \neq i^+$, noting that we indeed have $w_i \geq w_i^{\downarrow i}$.

D.2. Proof of Lemma 7

First, we show that $\mathcal{R}^{\langle w^{\downarrow} \rangle}(S) \geq (1-\epsilon) \cdot \mathcal{R}^{\langle w \rangle}(S)$. To this end, the important observation is that the initial weight vector w can be transformed into $w^{\downarrow}$ in n sequential steps, where in each step i, the preference weight of w_i is rounded down to $w_i^{\downarrow}$, the nearest power of $1+\epsilon$. With this view, repeated applications of Lemma 6 imply that $\pi^{\langle w^{\downarrow} \rangle}(i,S) \geq (1-\epsilon) \cdot \pi^{\langle w \rangle}(i,S)$ for every product $i \in S$. This lower bound on the choice probabilities yields the desired inequality, $\mathcal{R}^{\langle w^{\downarrow} \rangle}(S) \geq (1-\epsilon) \cdot \mathcal{R}^{\langle w \rangle}(S)$.

Next, we show that $\mathcal{R}^{\langle w^{\downarrow} \rangle}(S) \leq (1+2\epsilon) \cdot \mathcal{R}^{\langle w \rangle}(S)$. For this purpose, let $w^{\downarrow\downarrow}$ be the weight vector obtained from $w^{\downarrow}$ by rounding down the no-purchase weight, now taking a value of $\frac{1}{1+\epsilon}$ instead of 1. This alteration can only increase the choice probability of any product $i \in S$, since for all $m \in [n]$,

$$\begin{aligned}
\pi_m(i,S) &= \Pr\left[[\text{rank}(i,S) \leq m] \wedge [U_i > U_0]\right] \\
&= \Pr\left[\text{rank}(i,S) \leq m\right] \cdot \Pr\left[U_0 < U_i | \text{rank}(i,S) \leq m\right] ,
\end{aligned}$$

where $\Pr[\text{rank}(i,S) \leq m]$ is unrelated to the no-purchase weight, and $\Pr[U_0 < U_i | \text{rank}(i,S) \leq m]$ is decreasing in w_0. As a result, we have $\mathcal{R}^{\langle w^{\downarrow\downarrow} \rangle}(S) \geq \mathcal{R}^{\langle w^{\downarrow} \rangle}(S)$. Now let $w^{\uparrow} = (1+\epsilon) \cdot w^{\downarrow\downarrow}$ be the weight vector that results from scaling up each coordinate of $w^{\downarrow\downarrow}$ by a factor of $1+\epsilon$. It is easy to see that, similarly to the standard MNL model, the choice probabilities under the MP-MNL model are scale-invariant, meaning that $\pi^{\langle w^{\uparrow} \rangle}(i,S) = \pi^{\langle w^{\downarrow\downarrow} \rangle}(i,S)$ for any product $i \in S$. In turn, we have $\mathcal{R}^{\langle w^{\uparrow} \rangle}(S) = \mathcal{R}^{\langle w^{\downarrow\downarrow} \rangle}(S) \geq \mathcal{R}^{\langle w^{\downarrow} \rangle}(S)$, and to conclude the proof, it suffices to show that $\mathcal{R}^{\langle w \rangle}(S) \geq \frac{1}{1+2\epsilon} \cdot \mathcal{R}^{\langle w^{\uparrow} \rangle}(S)$. However, noting that $w_i \in [(1-\epsilon) \cdot w_i^{\uparrow}, w_i^{\uparrow}]$ for every $i \in S$, our previous argument for the revenue effect of rounding down weights can be reapplied, to obtain $\mathcal{R}^{\langle w \rangle}(S) \geq (1-\epsilon) \cdot \mathcal{R}^{\langle w^{\uparrow} \rangle}(S) \geq \frac{1}{1+2\epsilon} \cdot \mathcal{R}^{\langle w^{\uparrow} \rangle}(S)$ for $\epsilon \in (0, 1/2)$.

D.3. Proof of Lemma 8

Let S^* be an optimal assortment. Suppose that there is some weight class $\hat{q} \in [\mathcal{Q}]_0$ that violates the price-ordered structure given in the lemma's statement, meaning that $S^* \cap \mathcal{W}_{\hat{q}} \neq \mathcal{W}_{\hat{q}}[|S^* \cap \mathcal{W}_{\hat{q}}|]$. Let us define a new assortment, $\hat{S} = (S^* \setminus \mathcal{W}_{\hat{q}}) \cup \mathcal{W}_{\hat{q}}[|S^* \cap \mathcal{W}_{\hat{q}}|]$, which is identical to S^* over all classes except for $\hat{q}$; in the latter, $S^* \cap \mathcal{W}_{\hat{q}}$ is swapped for $\mathcal{W}_{\hat{q}}[|S^* \cap \mathcal{W}_{\hat{q}}|]$. We proceed by showing that $\hat{S}$ must be an optimal assortment as well. For this purpose, note that by symmetry, $\pi(i_1, S^*) = \pi(i_2, \hat{S})$ for any pair of products $i_1 \in S^* \cap \mathcal{W}_q$ and $i_2 \in \hat{S} \cap \mathcal{W}_q$, for every $q \in [\mathcal{Q}]_0$. Therefore, letting $\pi_{\hat{q}}^*$ stand for the latter choice probability with respect to $\mathcal{W}_{\hat{q}}$, we have

$$\begin{aligned}
\mathcal{R}(\hat{S}) &= \sum_{q \in [\mathcal{Q}]_0 : q \neq \hat{q}} \sum_{i \in \hat{S} \cap \mathcal{W}_q} r_i \cdot \pi(i, \hat{S}) + \sum_{i \in \hat{S} \cap \mathcal{W}_{\hat{q}}} r_i \cdot \pi(i, \hat{S}) \\
&= \sum_{q \in [\mathcal{Q}]_0 : q \neq \hat{q}} \sum_{i \in S^* \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) + \pi_{\hat{q}}^* \cdot \sum_{i \in \hat{S} \cap \mathcal{W}_{\hat{q}}} r_i \\
&\geq \sum_{q \in [\mathcal{Q}]_0 : q \neq \hat{q}} \sum_{i \in S^* \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) + \pi_{\hat{q}}^* \cdot \sum_{i \in S^* \cap \mathcal{W}_{\hat{q}}} r_i \\
&= \sum_{q \in [\mathcal{Q}]_0 : q \neq \hat{q}} \sum_{i \in S^* \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) + \sum_{i \in S^* \cap \mathcal{W}_{\hat{q}}} r_i \cdot \pi(i, S^*) \\
&= \mathcal{R}(S^*) \ .
\end{aligned}$$

The inequality above holds since $\sum_{i \in S^* \cap \mathcal{W}_q} r_i \leq \sum_{i \in \mathcal{W}_{\hat{q}}[|S^* \cap \mathcal{W}_{\hat{q}}|]} r_i$, as $\mathcal{W}_{\hat{q}}[|S^* \cap \mathcal{W}_{\hat{q}}|]$ is precisely the collection of $|S^* \cap \mathcal{W}_{\hat{q}}|$ highest priced items in weight class $\mathcal{W}_{\hat{q}}$. Repeating this swapping argument for all weight classes that violate the price-ordered property eventually produces an optimal assortment that satisfies the lemma's statement.

D.4. Proof of Lemma 9

We first argue that the assortment $\tilde{S} = \bigcup_{q \geq q_{\min}} \mathcal{W}_q[\tilde{k}_q]$ is necessarily a subset of the optimal assortment S^*. For this purpose, since S^* satisfies Lemma 8, we know in particular that $S^* \cap \mathcal{W}_q = \mathcal{W}_q[k_q^*]$ for every $q \geq q_{\min}$, and it remains to explain why $\tilde{k}_q \leq k_q^*$. To this end, by the way $\tilde{k}_q$ is chosen, it suffices to show that $\sum_{i \in \mathcal{W}_q[k_q^*]} r_i \geq \frac{\tilde{\mathcal{R}}_q}{\tilde{\pi}_q}$, which is indeed the case since

$$\begin{aligned}
\tilde{\pi}_q \cdot \sum_{i \in \mathcal{W}_q[k_q^*]} r_i &\geq \pi_q^* \cdot \sum_{i \in \mathcal{W}_q[k_q^*]} r_i \\
&= \sum_{i \in S^* \cap \mathcal{W}_q} r_i \cdot \pi(i, S^*) \\
&= \mathcal{R}_q(S^*) \\
&\geq \tilde{\mathcal{R}}_q ,
\end{aligned}$$

where the first and second inequalities hold since $\tilde{\pi}_q \geq \pi_q^*$ and $\tilde{\mathcal{R}}_q \leq \mathcal{R}_q(S^*)$, respectively.

Given this relation, to establish the desired revenue guarantee, note that

$$\begin{aligned}
\mathcal{R}(\tilde{S}) &= \sum_{q \geq q_{\min}} \sum_{i \in \mathcal{W}_q[\tilde{k}_q]} r_i \cdot \pi(i, \tilde{S}) \\
&\geq \sum_{q \geq q_{\min}} \sum_{i \in \mathcal{W}_q[\tilde{k}_q]} r_i \cdot \pi(i, S^*) \\
&= \sum_{q \geq q_{\min}} \left(\pi_q^* \cdot \sum_{i \in \mathcal{W}_q[\tilde{k}_q]} r_i \right) \\
&\geq \sum_{q \geq q_{\min}} \pi_q^* \cdot \frac{\tilde{\mathcal{R}}_q}{\tilde{\pi}_q} \\
&\geq (1-\epsilon) \cdot \sum_{q \geq q_{\min}} \tilde{\mathcal{R}}_q .
\end{aligned}$$

Here, the first inequality holds since $\tilde{S} \subseteq S^*$, as explained above. The second inequality follows by recalling that $\tilde{k}_q$ was chosen to satisfy $\sum_{i \in \mathcal{W}_q[\tilde{k}_q]} r_i \geq \frac{\tilde{\mathcal{R}}_q}{\tilde{\pi}_q}$. The third inequality is obtained by noting that $\tilde{\pi}_q \leq (1+\epsilon) \cdot \pi_q^*$. Consequently, we can now relate between the expected revenues of $\tilde{S}$ and S^* by observing that

$$\begin{aligned}
\mathcal{R}(\tilde{S}) &\geq (1-\epsilon) \cdot \sum_{q \geq q_{\min}} \left(\mathcal{R}_q(S^*) - \frac{\epsilon}{\mathcal{Q}+1} \cdot \widetilde{\mathrm{OPT}} \right) \\
&\geq (1-\epsilon) \cdot \sum_{q \in [\mathcal{Q}]_0} \mathcal{R}_q(S^*) - \sum_{q < q_{\min}} \mathcal{R}_q(S^*) - \epsilon \cdot \widetilde{\mathrm{OPT}} \\
&\geq (1-2\epsilon) \cdot \mathcal{R}(S^*) - \sum_{q < q_{\min}} \left(\pi_q^* \cdot \sum_{i \in S^* \cap \mathcal{W}_q} r_i \right) \\
&\geq (1-2\epsilon) \cdot \mathcal{R}(S^*) - \frac{\epsilon}{n} \cdot \sum_{q < q_{\min}} \sum_{i \in S^* \cap \mathcal{W}_q} \left(\Pr[M > 0] \cdot \frac{w_{\min}}{1+w_{\min}} \cdot r_i \right)
\end{aligned}$$

$$\begin{aligned} &\geq (1-2\epsilon)\cdot\mathcal{R}(S^*) - \epsilon\cdot\frac{|S^*|}{n}\cdot\mathcal{R}(S^*) \\ &\geq (1-3\epsilon)\cdot\mathcal{R}(S^*) \ . \end{aligned}$$

Here, the first inequality holds since $\tilde{\mathcal{R}}_q \geq \mathcal{R}_q(S^*) - \frac{\epsilon}{\mathcal{Q}+1}\cdot\widetilde{\mathrm{OPT}}$. The third inequality follows by noting that $\widetilde{\mathrm{OPT}} \leq \mathcal{R}(S^*)$. To obtain the fourth inequality, we recall that $\pi^*_q < \frac{\epsilon}{n}\cdot\Pr[M>0]\cdot\frac{w_{\min}}{1+w_{\min}}$ for every $q < q_{\min}$, by definition of $q_{\min}$. Finally, the fifth inequality holds since, due to the optimality of S^*,

$$\mathcal{R}(S^*) \geq \mathcal{R}(\{i\}) = \Pr[M>0]\cdot\frac{w_i}{1+w_i}\cdot r_i \geq \Pr[M>0]\cdot\frac{w_{\min}}{1+w_{\min}}\cdot r_i \ .$$

D.5. Proof of Lemma 10

Focusing on a single product $i\in S$, we first observe that for any given $m\in[n]$,

$$\begin{aligned} \pi_m(i,S) &= \Pr\left[[U_i > U_0]\wedge[\mathrm{rank}(i,S)\leq m]\right] \\ &= \Pr\left[U_i>U_0\right]\cdot\Pr\left[\mathrm{rank}(i,S)\leq m | U_i > U_0\right] \\ &= \frac{w_i}{1+w_i}\cdot\Pr\left[\mathrm{rank}(i,S)\leq m | U_i > U_0\right] \ . \end{aligned}$$

Therefore, by definition of $\tilde{\pi}_m(i,S)$, we have

$$\begin{aligned} &\Pr\left[|\tilde{\pi}_m(i,S)-\pi_m(i,S)| \geq \frac{\epsilon}{n}\cdot\frac{w_i}{1+w_i}\right] \\ &\quad= \Pr\left[\left|\frac{1}{N}\cdot\sum_{\nu\in[N]}\mathbb{1}[\rho_\nu\leq m] - \Pr\left[\mathrm{rank}(i,S)\leq m|U_i>U_0\right]\right| \geq \frac{\epsilon}{n}\right] \\ &\quad\leq 2\exp\left(-\frac{2N\epsilon^2}{n^2}\right) \\ &\quad\leq \frac{\delta}{n^2} \ . \end{aligned}$$

Here, the first inequality follows from Hoeffding's inequality (1963), noting that the indicators $\{\mathbb{1}[\rho_\nu \leq m]\}_{\nu\in[N]}$ are independent, each with an expected value of $\Pr[\mathrm{rank}(i,S)\leq m|U_i>U_0]$. The second inequality is obtained by plugging in $N = \lceil\frac{n^2}{2\epsilon^2}\cdot\ln(\frac{2n^2}{\delta})\rceil$.

Now, a simple application of the union bound implies that, with probability at least $1-\delta$, we have $|\tilde{\pi}_m(i,S)-\pi_m(i,S)| \leq \frac{\epsilon}{n}\cdot\frac{w_i}{1+w_i}$ simultaneously for all $i\in S$ and $m\in[n]$. When this event occurs,

$$\begin{aligned} \left|\tilde{\mathcal{R}}(S)-\mathcal{R}(S)\right| &= \left|\sum_{i\in S}\left(r_i\cdot\sum_{m\in[n]}\Pr\left[M=m\right]\cdot(\tilde{\pi}_m(i,S)-\pi_m(i,S))\right)\right| \\ &\leq \sum_{i\in S}\left(r_i\cdot\sum_{m\in[n]}\Pr\left[M=m\right]\cdot|\tilde{\pi}_m(i,S)-\pi_m(i,S)|\right) \\ &\leq \frac{\epsilon}{n}\cdot\sum_{i\in S}\Pr\left[M\geq 1\right]\cdot\frac{r_iw_i}{1+w_i} \\ &\leq \epsilon\cdot\mathcal{R}(S^*) \ , \end{aligned}$$

where the third inequality holds since $\mathcal{R}(S^*) \geq \Pr[M\geq 1]\cdot\frac{r_iw_i}{1+w_i}$ for every product $i\in[n]$.

Appendix E: PTAS for General Instances

This section is dedicated to devising a polynomial-time approximation scheme for arbitrarily-structured instances, whose specific guarantees were stated in Theorem 2. At least intuitively, our approach begins by decomposing any given instance into a sequence of so-called clusters, each with a polynomially-bounded $\frac{w_{\max}}{w_{\min}}$ ratio, such that there are large gaps between the product weights of successive clusters. These gaps allow us to establish a position-reservation lemma, capturing the notion that products offered in "heavier" clusters almost always outrank those in "lighter" clusters. With the latter result in place, we argue that the assortment optimization problems associated with various clusters become independent, once they are suitably linked through cardinality constraints on the number of offered products.

E.1. Clustering

For simplicity of analysis, we assume without loss of generality that $n \geq 1/\epsilon$, as in the opposite case, the approximation scheme of Section 4 clearly forms a PTAS. Now, by rather standard shifting arguments, at the cost of losing an $O(\epsilon)$-fraction of the optimal revenue, we can assume that the underlying set of products is partitioned into a sequence of weight clusters $\mathcal{C}_1, \mathcal{C}_2, \dots$ that satisfy the next two properties:

1. Intra-cluster ratios: Within each cluster, the weights of any two products differ by a multiplicative factor of at most $n^{5/\epsilon}$.

2. Inter-cluster gaps: For every $t_1 < t_2$, the weight of any product in cluster $\mathcal{C}_{t_1}$ is greater than the weight of any product in $\mathcal{C}_{t_2}$ by a multiplicative factor of at least n^5.

The specifics of this construction are described in Appendix F.1.

Now, letting T be the maximal index of a cluster in which each product has a preference weight of at least n^4, we define the cluster $\mathcal{C}_\infty$ as the union of $\mathcal{C}_{T+1}, \mathcal{C}_{T+2}, \dots$. By property 1, the weight of any product in $\mathcal{C}_\infty$ resides within $(0, n^{4+5/\epsilon}]$, although the weight ratio between any two products in this cluster may be arbitrarily large. For convenience, rather than representing an assortment as a single subset of products, we use a separate notation for the products picked out of each cluster. In other words, an assortment will now be written as $S = (S_1, \dots, S_T, S_\infty)$, where S_t stands for the collection of products picked from cluster $\mathcal{C}_t$, for every $t \in [T]_\infty$.

E.2. The position-reservation lemma

At least intuitively, our motivation for introducing $\mathcal{C}_1, \dots, \mathcal{C}_T, \mathcal{C}_\infty$ is that, due to having large inter-clusters weight gaps, it is highly improbable that any product offered out of a given cluster has a utility larger than those of products offered out of lower-indexed clusters. Therefore, products of the latter type can be perceived as blocking the former from attaining a certain range of ranks, leading to a by-cluster problem decomposition. To formalize this notion, we begin by proving a surprising structural property: With respect to any assortment S, let $\sigma_{t-1}(S) = \sum_{\tau \in [t-1]} |S_\tau|$ be the total number of products offered by S out of clusters $\mathcal{C}_1, \dots, \mathcal{C}_{t-1}$. We show that the choice probability of any product $i \in S_t$ with respect to S matches, up to negligible terms, its choice probability when we offer S_t by itself and eliminate in advance $\sigma_{t-1}(S)$ purchases from the overall demand. The proof appears in Appendix F.2.

LEMMA 14. *Let $S = (S_1, \dots, S_T, S_\infty)$. Then, for any $t \in [T]_\infty$, any product $i \in S_t$, and any $m \in [n]$, we have*

$$\pi_{[m-\sigma_{t-1}(S)]^+}(i, S_t) \in \pi_m(i, S) \pm \frac{2}{n^3} \cdot \frac{w_i}{1+w_i} \ .$$

As an immediate consequence, letting S^* be an optimal assortment, we obtain an upper bound on the optimal expected revenue by decomposing this quantity into proxies for the contributions of different clusters. This claim is formalized below, with the convention that the random number of maximum purchases in question will appear as a subscript of the expected revenue function, meaning in particular that $\mathcal{R}_M$ identifies with $\mathcal{R}$.

LEMMA 15. $\mathcal{R}_M(S^*) \le (1+4\epsilon) \cdot \sum_{t \in [T]_\infty} \mathcal{R}_{[M-\sigma_{t-1}(S^*)]^+}(S^*_t)$.

Proof. By specializing Lemma 14 to S^*, it follows in particular that for every cluster index $t \in [T]_\infty$, product $i \in S_t$, and $m \in [n]$, we have

$$\begin{aligned} \pi_m(i, S^*) &\le \pi_{[m-\sigma_{t-1}(S^*)]^+}\left(i, S^*_t\right) + \frac{2}{n^3} \cdot \frac{w_i}{1+w_i} \\ &\le \pi_{[m-\sigma_{t-1}(S^*)]^+}\left(i, S^*_t\right) + \frac{2\epsilon}{n} \cdot \frac{w_i}{1+w_i} \ , \end{aligned}$$

where the second inequality holds since $n \ge 1/\epsilon$. Therefore,

$$\begin{aligned} \mathcal{R}_M\left(S^*\right) &= \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \sum_{t \in [T]_\infty} \sum_{i \in S^*_t} r_i \cdot \pi_m\left(i, S^*\right) \\ &\le \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \sum_{t \in [T]_\infty} \sum_{i \in S^*_t} r_i \cdot \left(\pi_{[m-\sigma_{t-1}(S^*)]^+}\left(i, S^*_t\right) + \frac{2\epsilon}{n} \cdot \frac{w_i}{1+w_i} \right) \\ &= \sum_{t \in [T]_\infty} \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \mathcal{R}_{[m-\sigma_{t-1}(S^*)]^+}\left(i, S^*_t\right) + \frac{2\epsilon}{n} \cdot \sum_{t \in [T]_\infty} \sum_{i \in S^*_t} \underbrace{\Pr\left[M > 0\right] \cdot \frac{r_i w_i}{1+w_i}}_{\le \mathcal{R}_M(S^*)} \\ &\le \sum_{t \in [T]_\infty} \mathcal{R}_{[M-\sigma_{t-1}(S^*)]^+}\left(S^*_t\right) + 2\epsilon \cdot \mathcal{R}_M\left(S^*\right) \ , \end{aligned}$$

and we attain the desired inequality by rearranging the one above. $\square$

E.3. The dynamic program

By inspecting how Lemma 15 decomposes the optimal expected revenue $\mathcal{R}_M(S^*)$, the crucial property we exploit is that the individual contribution $\mathcal{R}_{[M-\sigma_{t-1}(S^*)]^+}(S^*_t)$ of each cluster $\mathcal{C}_t$ depends on our choices for all other clusters only through $\sigma_{t-1}(S^*) = \sum_{\tau \in [t-1]} |S^*_\tau|$. Noting that this dependency involves the total number of products picked out of lower-indexed clusters and nothing more, we are now ready to tackle the assortment optimization problem by means of dynamic programming. Specifically, each state (t, N_t) of our dynamic program consists of the following parameters:

- The current cluster index, $t \in [T]_{0,\infty}$.
- An upper bound $N_t \in [n]_0$ on the total number of products offered out of clusters $\mathcal{C}_1, \dots, \mathcal{C}_t$.

With respect to these states, the value function $F(t, N_t)$ is defined through the recursive equations:

$$F(t, N_t) = \max_{N_{t-1} \le N_t} \left\{ F(t-1, N_{t-1}) + \mathrm{OPT}\left[t, [M - N_{t-1}]^+, N_t - N_{t-1}\right] \right\} \ . \tag{14}$$

Here, $\mathrm{OPT}[t,[M-N_{t-1}]^+,N_t-N_{t-1}]$ represents the optimal expected revenue for an assortment optimization instance where the underlying products are those in cluster $\mathcal{C}_t$, with respect to a random demand of $[M-N_{t-1}]^+$, subject to the additional constraint that at most N_t-N_{t-1} products can be offered. Terminal states correspond to $t=0$, in which case $F(0,\cdot)=0$. The next claim, whose proof is given in Appendix F.3, identifies state (∞,n) as the one for which we expect to obtain a near-optimal assortment.

Lemma 16. $F(\infty,n)\geq(1-4\epsilon)\cdot\mathcal{R}_M(S^*)$.

E.4. The Approximate program

In spite of this characterization, the value function F is defined through the recursive equations (14), which unfortunately involve the single-cluster term $\mathrm{OPT}[t,[M-N_{t-1}]^+,N_t-N_{t-1}]$. Since we are not aware of an efficient way to compute its exact value, we proceed by explaining how to obtain a $(1-\epsilon)$-approximate counterpart $\mathrm{APX}[t,[M-N_{t-1}]^+,N_t-N_{t-1}]$. For every $t\in[T]_\infty$, the latter will satisfy

$$\mathrm{APX}\left[t,[M-N_{t-1}]^+,N_t-N_{t-1}\right]\geq(1-\epsilon)\cdot\mathrm{OPT}\left[t,[M-N_{t-1}]^+,N_t-N_{t-1}\right]\ , \tag{15}$$

with probability at least 1/2. To this end, we distinguish between two cases, depending on the cluster in question:

- Case 1: $t\leq T$. In this case, property 1 of our clusters (see Appendix E.1) ensures that the weights of any two products in cluster $\mathcal{C}_t$ differ by a multiplicative factor of at most $n^{5/\epsilon}$. Therefore, the randomized approximation scheme we propose in Section 5.1 constitutes an $O(n^{O(1/\epsilon^3)})$-time PTAS for the single-cluster instance under consideration (see Theorem 3). We remind the reader that this approach works even subject to a cardinality constraint on the offered assortment, ensuring that we indeed pick at most N_t-N_{t-1} products.
- Case 2: $t=\infty$. As mentioned earlier, while the weight of any product in $\mathcal{C}_\infty$ resides within $(0,n^{4+5/\epsilon}]$, the weight ratio between any two products in this cluster may be arbitrarily large. However, when evaluating the value function at its final state (∞,n), we are clearly operating without a cardinality constraint. In Appendix F.4, we explain how to derive a randomized $O(n^{O(1/\epsilon^3)})$-time PTAS for these particular circumstances, by extending the basic approach of Section 5.1.

Now, by plugging this single-cluster approximation into the recursive equations (14), we have just formulated an approximate dynamic program F_{APX} over precisely the same set of states, given by

$$F_{\mathrm{APX}}(t,N_t)=\max_{N_{t-1}\leq N_t}\left\{F_{\mathrm{APX}}(t-1,N_{t-1})+\mathrm{APX}\left[t,[M-N_{t-1}]^+,N_t-N_{t-1}\right]\right\}\ .$$

As a side note, since our single-cluster procedure APX is successful with probability at least 1/2, its failure probability can be shrunk to $\frac{1}{2n^3}$ via $O(\log n)$ independent repetitions. With this observation, a simple application of the union bound (over at most n^3 calls to APX) implies that, with probability at least 1/2, the approximation guarantee in (15) holds throughout the recursive evaluation of $F_{\mathrm{APX}}(\infty,n)$.

E.5. Analysis

Approximation guarantee. Letting $S^{\mathrm{DP}}=(S_1^{\mathrm{DP}},\dots,S_T^{\mathrm{DP}},S_\infty^{\mathrm{DP}})$ be the resulting assortment, we conclude our analysis by arguing that its expected revenue $\mathcal{R}_M(S^{\mathrm{DP}})$ nearly matches the optimal revenue $\mathcal{R}_M(S^*)$.

Lemma 17. $\mathcal{R}_M(S^{\mathrm{DP}})\geq(1-7\epsilon)\cdot\mathcal{R}_M(S^*)$.

Proof. By specializing Lemma 14 to S^{DP}, it follows in particular that for every cluster index $t \in [T]_\infty$, product $i \in S_t^{\mathrm{DP}}$, and $m \in [n]$, we have

$$\begin{aligned}
\pi_m\left(i, S^{\mathrm{DP}}\right) &\geq \pi_{[m-\sigma_{t-1}(S^{\mathrm{DP}})]^+}\left(i, S_t^{\mathrm{DP}}\right) - \frac{2}{n^3} \cdot \frac{w_i}{1+w_i} \\
&\geq \pi_{[m-\sigma_{t-1}(S^{\mathrm{DP}})]^+}\left(i, S_t^{\mathrm{DP}}\right) - \frac{2\epsilon}{n} \cdot \frac{w_i}{1+w_i} \ ,
\end{aligned}$$

where the second inequality holds since $n \geq 1/\epsilon$. Therefore, the expected revenue $\mathcal{R}_M(S^{\mathrm{DP}})$ can be lower bounded in terms of $\mathcal{R}^M(S^*)$ by noting that

$$\begin{aligned}
\mathcal{R}_M\left(S^{\mathrm{DP}}\right) &= \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \sum_{t \in [T]_\infty} \sum_{i \in S_t^{\mathrm{DP}}} r_i \cdot \pi_m\left(i, S^{\mathrm{DP}}\right) \\
&\geq \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \sum_{t \in [T]_\infty} \sum_{i \in S_t^{\mathrm{DP}}} r_i \cdot \left(\pi_{[m-\sigma_{t-1}(S^{\mathrm{DP}})]^+}\left(i, S_t^{\mathrm{DP}}\right) - \frac{2\epsilon}{n} \cdot \frac{w_i}{1+w_i}\right) \\
&= \sum_{t \in [T]_\infty} \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \mathcal{R}_{[m-\sigma_{t-1}(S^{\mathrm{DP}})]^+}\left(S_t^{\mathrm{DP}}\right) - \frac{2\epsilon}{n} \cdot \sum_{t \in [T]_\infty} \sum_{i \in S_t^{\mathrm{DP}}} \underbrace{\Pr\left[M > 0\right] \cdot \frac{r_i w_i}{1+w_i}}_{\leq \mathcal{R}_M(S^*)} \\
&\geq \sum_{t \in [T]_\infty} \mathcal{R}_{[M-\sigma_{t-1}(S^{\mathrm{DP}})]^+}\left(S_t^{\mathrm{DP}}\right) - 2\epsilon \cdot \mathcal{R}_M\left(S^*\right) \\
&= F_{\mathrm{APX}}(\infty, n) - 2\epsilon \cdot \mathcal{R}_M\left(S^*\right) \\
&\geq (1-\epsilon) \cdot F(\infty, n) - 2\epsilon \cdot \mathcal{R}_M\left(S^*\right) \\
&\geq (1-7\epsilon) \cdot \mathcal{R}_M(S^*) \ ,
\end{aligned}$$

where the last inequality follows from Lemma 16. $\square$

Running time. We first observe that, as explained in Appendices E.1 and F.1, the sequence of clusters $\mathcal{C}_1, \mathcal{C}_2, \ldots$ can easily be constructed in polynomial time. Given these clusters, our approximate dynamic program F_{APX} has only $O(n^2)$ states to consider, where each such state is evaluated via its recursive equations in $O(n^{O(1/\epsilon^3)})$ time. Consequently, the overall running time of our approximation scheme is $O(n^{O(1/\epsilon^3)})$, precisely as stated in Theorem 2.

Appendix F: Proofs from Appendix E

F.1. Constructing the clusters $\mathcal{C}_1, \mathcal{C}_2, \ldots$

Construction. Our clustering method begins by partitioning the entire collection products into "intervals" $I_1, I_2, \ldots$, based on geometrically-increasing preference weights, by powers of $\frac{n^5}{\epsilon}$. Specifically, recalling that $w_{\max} = \max_{i \in [n]} w_i$ and $w_{\min} = \min_{i \in [n]} w_i$, we define:

$$I_1 = \left\{ i \in [n] : w_i \in \left(\frac{\epsilon}{n^5} \cdot w_{\max}, w_{\max} \right] \right\}, \quad I_2 = \left\{ i \in [n] : w_i \in \left(\left(\frac{\epsilon}{n^5}\right)^2 \cdot w_{\max}, \frac{\epsilon}{n^5} \cdot w_{\max} \right] \right\} \ ,$$

so on and so forth, where in general

$$I_\ell = \left\{ i \in [n] : w_i \in \left(\left(\frac{\epsilon}{n^5}\right)^\ell \cdot w_{\max}, \left(\frac{\epsilon}{n^5}\right)^{\ell-1} \cdot w_{\max} \right] \right\} \ .$$

Now, assuming without loss of generality that $1/\epsilon$ is an integer, let $\chi \sim U\{0, 1, \ldots, \frac{1}{\epsilon} - 1\}$ be a discrete uniform random variable. With respect to χ, we say that an interval I_ℓ is "dropped" when $(\ell \mod \frac{1}{\epsilon}) = \chi$; otherwise, this interval is "undropped". We denote by $\mathcal{C}_1^\chi$ the set of products belonging to (undropped) intervals that are indexed lower than the first dropped interval (i.e., $I_1, \ldots, I_{\chi-1}$). Then, $\mathcal{C}_2^\chi$ is the set of products belonging to (undropped) intervals that are indexed between the first and second dropped intervals (i.e., $I_{\chi+1}, \ldots, I_{\chi+\frac{1}{\epsilon}-1}$), so on and so forth.

Analysis. It is easy to verify that, for any possible realization of χ, its corresponding sequence of clusters $\mathcal{C}_1^\chi, \mathcal{C}_2^\chi, \ldots$ satisfies the two structural properties mentioned in Appendix E.1, namely, intra-cluster ratios and inter-cluster gaps, both with room to spare. Moreover, we argue that a negligible amount of revenue is lost in expectation, when unclustered products are ignored. To formalize this argument, recalling that S^* stands for a fixed optimal assortment, let us define $S^{*\chi}$ as its subset of clustered products, meaning that $S^{*\chi} = S^* \cap (\bigcup_{t\geq 1} \mathcal{C}_t^\chi)$. The next lemma shows that, by offering the latter assortment, only an $O(\epsilon)$-fraction of the optimal revenue is lost in expectation, where the latter is taken over the randomness in choosing χ.

LEMMA 18. *$\mathbb{E}_\chi[\mathcal{R}(S^{*\chi})] \geq (1-\epsilon)\cdot \mathcal{R}(S^*)$.*

Proof. For every product $i \in S^*$, let $I_{(i)}$ be the unique interval to which i belongs. Then, by definition of $S^{*\chi}$, we have

$$\begin{aligned}
\mathbb{E}_\chi\left[\mathcal{R}(S^{*\chi})\right] &= \sum_{i\in S^*} r_i \cdot \mathbb{E}_\chi\left[\pi(i, S^{*\chi})\right] \\
&= \sum_{i\in S^*} r_i \cdot \Pr_\chi\left[I_{(i)} \text{ is undropped}\right] \cdot \mathbb{E}_\chi\left[\pi(i, S^{*\chi}) | \, I_{(i)} \text{ is undropped}\right] \\
&\geq (1-\epsilon)\cdot \sum_{i\in S^*} r_i \cdot \pi(i, S^*) \\
&= (1-\epsilon)\cdot \mathcal{R}(S^*) \ .
\end{aligned}$$

The inequality above holds since any interval is undropped with probability $1-\epsilon$. In addition, for every product $i \in S^*$, when its corresponding interval $I_{(i)}$ is undropped, we have $\pi(i, S^{*\chi}) \geq \pi(i, S^*)$, since $S^{*\chi} \subseteq S^*$ with probability 1. $\square$

F.2. Proof of Lemma 14

We divide our analysis into two cases, depending on the relation between m and $\sigma_{t-1}(S)$.

Case 1: $m \leq \sigma_{t-1}(S)$. In this case, we clearly have $\pi_{[m-\sigma_{t-1}(S)]^+}(i, S_t) = 0$. However,

$$\begin{aligned}
\pi_m(i, S) &= \pi_m(i, (S_1, \ldots, S_\infty)) \\
&\leq \pi_m(i, (S_1, \ldots, S_{t-1}, \{i\}, \emptyset, \ldots, \emptyset)) \\
&\leq \Pr\left[\biguplus_{\tau\in[t-1]} \biguplus_{j\in S_\tau} [U_i > U_j]\right] && (16)\\
&\leq \sum_{\tau\in[t-1]} \sum_{j\in S_\tau} \frac{w_i}{w_i + w_j} \\
&\leq \frac{2}{n^3}\cdot \frac{w_i}{1+w_i} \ . && (17)
\end{aligned}$$

Here, inequality (16) holds since at least one of the sets $S_1, \ldots, S_{t-1}$ is non-empty, as $\sum_{\tau\in[t-1]} |S_\tau| = \sigma_{t-1}(S) \geq m \geq 1$. To derive inequality (17), we consider two cases:

- When $w_i \geq 1$: Recalling that, by property 2, the weight gap between successive clusters is at least n^5, it follows that $w_j \geq n^5 \cdot w_i$ for every product $j \in \biguplus_{\tau\in[t-1]} S_\tau$. Therefore, $\frac{w_i}{w_i+w_j} \leq \frac{1}{n^5} \leq \frac{2}{n^5}\cdot\frac{w_i}{1+w_i}$, where the last inequality holds since $w_i \geq 1$.

- When $w_i < 1$: By definition of T, since $w_i < 1$, product i necessarily belongs to cluster $\mathcal{C}_\infty$ (i.e., $t = \infty$). As a result, $w_j \geq n^4$ for every product $j \in \biguplus_{\tau\in[t-1]} S_\tau$, again by definition of T. Thus, $\frac{w_i}{w_i+w_j} \leq \frac{w_i}{n^4} \leq \frac{2}{n^4}\cdot\frac{w_i}{1+w_i}$, where the last inequality holds since $w_i < 1$.

Case 2: $m > \sigma_{t-1}(S)$. Here, we derive one direction by noting that

$$
\begin{aligned}
\pi_m(i,S) &\le \pi_m(i,(S_1,\dots,S_t,\emptyset,\dots,\emptyset)) \\
&= \Pr\left[[U_i > U_0] \wedge \left[\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le m\right]\right] \\
&= \Pr\left[[U_i > U_0] \wedge \left[\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le m\right] \wedge \overline{\biguplus_{\tau\in[t-1]} \biguplus_{j\in S_\tau} [U_i > U_j]}\right] \\
&\quad + \Pr\left[[U_i > U_0] \wedge \left[\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le m\right] \wedge \biguplus_{\tau\in[t-1]} \biguplus_{j\in S_\tau} [U_i > U_j]\right] \\
&\le \underbrace{\Pr\left[[U_i > U_0] \wedge [\mathrm{rank}(i,S_t) \le m - \sigma_{t-1}(S)]\right]}_{\pi_{m-\sigma_{t-1}(S)}(i,S_t)} \\
&\quad + \underbrace{\Pr\left[\biguplus_{\tau\in[t-1]} \biguplus_{j\in S_\tau} [U_i > U_j]\right]}_{\le \frac{2}{n^3}\cdot\frac{w_i}{1+w_i} \text{ by (17)}} \\
&\le \pi_{m-\sigma_{t-1}(S)}(i,S_t) + \frac{2}{n^3}\cdot\frac{w_i}{1+w_i} \ .
\end{aligned}
$$

Now, in order to derive the second direction, we observe that

$$
\begin{aligned}
\pi_m(i,S) &\ge \Pr\left[[U_i > U_0] \wedge [\mathrm{rank}(i,S) \le m] \wedge \overline{\biguplus_{\tau\ge t+1} \biguplus_{j\in S_\tau} [U_j > U_i]}\right] \\
&= \Pr\left[[U_i > U_0] \wedge \left[\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le m\right] \wedge \overline{\biguplus_{\tau\ge t+1} \biguplus_{j\in S_\tau} [U_j > U_i]}\right] \\
&\ge \Pr\left[[U_i > U_0] \wedge \left[\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le m\right]\right] - \underbrace{\Pr\left[\biguplus_{\tau\ge t+1} \biguplus_{j\in S_\tau} [U_j > U_i]\right]}_{\le \frac{2}{n^4}\cdot\frac{w_i}{1+w_i} \text{ (see below)}} \qquad (18) \\
&\ge \Pr\left[[U_i > U_0] \wedge [\mathrm{rank}(i,S_t) \le m - \sigma_{t-1}(S)]\right] - \frac{2}{n^4}\cdot\frac{w_i}{1+w_i} \\
&= \pi_{m-\sigma_{t-1}(S)}(i,S_t) - \frac{2}{n^4}\cdot\frac{w_i}{1+w_i} \ .
\end{aligned}
$$

Here, we obtain an upper bound on the second term in (18) by noting that, when $t = \infty$, this probability is clearly 0. When $t \le T$, due to having a weight gap of at least n^5 between successive clusters, it follows that $w_i \ge n^5 \cdot w_j$ for every product $j \in \biguplus_{\tau\ge t+1} S_\tau$. Therefore,

$$
\Pr\left[\biguplus_{\tau\ge t+1} \biguplus_{j\in S_\tau} [U_j > U_i]\right] \le n \cdot \frac{1}{1+n^5} \le \frac{1}{n^4} \le \frac{2}{n^4}\cdot\frac{w_i}{1+w_i} \ ,
$$

where the last inequality holds since $i \in S_t \subseteq \mathcal{C}_t$ and $t \le T$, implying in particular that $w_i \ge n^4 \ge 1$. As for the first term in (18), note that in any realization,

$$
\mathrm{rank}\left(i, \biguplus_{\tau\in[t]} S_\tau\right) \le \mathrm{rank}(i,S_t) + \sum_{\tau\in[t-1]} |S_\tau| = \mathrm{rank}(i,S_t) + \sigma_{t-1}(S) \ .
$$

F.3. Proof of Lemma 16

The important observation is that, given the optimal assortment S^*, the sequence of states and actions

$$(0,0) \xrightarrow{|S_1^*|} (1,\sigma_1(S^*)) \xrightarrow{|S_2^*|} (2,\sigma_2(S^*)) \xrightarrow{|S_3^*|} \cdots$$

forms a feasible path to be traversed by our dynamic program. With respect to this sequence, we clearly have $\mathrm{OPT}[t,[M-\sigma_{t-1}(S^*)]^+,\sigma_t(S^*)-\sigma_{t-1}(S^*)] \geq \mathcal{R}_{[M-\sigma_{t-1}(S^*)]^+}(S_t^*)$ for every $t \in [T]_\infty$, since S_t^* is a feasible solution to this single-cluster instance. As a result,

$$\begin{aligned} F(\infty,n) &\geq \sum_{t\in[T]_\infty} \mathrm{OPT}\left[t,[M-\sigma_{t-1}(S^*)]^+,\sigma_t(S^*)-\sigma_{t-1}(S^*)\right] \\ &\geq \sum_{t\in[T]_\infty} \mathcal{R}_{[M-\sigma_{t-1}(S^*)]^+}(S_t^*) \\ &\geq (1-4\epsilon)\cdot\mathcal{R}^M(S^*) \ , \end{aligned}$$

where the last inequality follows from Lemma 15.

F.4. PTAS for cluster $\mathcal{C}_\infty$

In this section, we explain how our previously-developed methods can be extended to obtain a randomized PTAS for single-cluster instances defined over $\mathcal{C}_\infty$.

Creating a random hole. Recalling that the weight of any product in $\mathcal{C}_\infty$ is known to reside within $(0,n^{4+5/\epsilon}]$, our first step consists of creating a random "hole" within this interval. To this end, let us define a collection of $1/\epsilon$ intervals, produced by starting at $\frac{\epsilon}{n^5}$ and jumping down by powers of n^5:

$$\left[\frac{\epsilon}{n^{5+5/\epsilon}},\frac{\epsilon}{n^{5/\epsilon}}\right),\ldots,\left[\frac{\epsilon}{n^{15}},\frac{\epsilon}{n^{10}}\right),\left[\frac{\epsilon}{n^{10}},\frac{\epsilon}{n^5}\right] \ .$$

We pick one of these intervals uniformly at random, and eliminate all products whose weight falls within the chosen interval. It is easy to verify that we are losing in expectation only on $O(\epsilon)$-fraction of the optimal revenue, since each product in an optimal assortment is eliminated with probability at most ϵ. As a result, the cluster $\mathcal{C}_\infty$ breaks into $\mathcal{C}_\infty^-$ and $\mathcal{C}_\infty^+$, corresponding to the products whose weights are smaller than or greater than those in the deleted interval, respectively. From this point on, let S^* be an optimal assortment for the post-deletion instance.

Constructing a near-optimal assortment. Even though the weight ratio between any two products in $\mathcal{C}_\infty^+ \uplus \mathcal{C}_\infty^-$ may be arbitrarily large, let us temporarily overlook running time considerations, and assume that the randomized approximation scheme of Section 5.1 is applied with respect to $\mathcal{C}_\infty^+ \uplus \mathcal{C}_\infty^-$. Then, its resulting assortment S would have an expected revenue of $\mathcal{R}_M(S) \geq (1-\epsilon)\cdot\mathcal{R}_M(S^*)$. In addition, since our enumeration procedure is operating with respect to S^*, one can augment it with an extra cardinality constraint, stating that at most $|S^{*+}|$ products may be offered out of $\mathcal{C}_\infty^+$. Note that, due to being dependent on the optimal assortment S^*, the quantity $|S^{*+}|$ is generally unknown, and we therefore guess its value a-priori.

Now, to attain a polynomial-time approach, the crucial observation is that the weight of each product in $\mathcal{C}_\infty^+$ is necessarily bounded within $[\frac{\epsilon}{n^{5/\epsilon}},n^{4+5/\epsilon}]$, regardless of which interval was actually deleted, meaning that the weight ratio in $\mathcal{C}_\infty^+$ is $O(n^{O(1/\epsilon)})$. Therefore, our enumeration procedure has only $O(n^{O(1/\epsilon^3)})$ possible choices of product offerings from $\mathcal{C}_\infty^+$. One of these choices, say S^+, consists of at most $|S^{*+}|$ products and can be completed with an additional set of products $S^- \subseteq \mathcal{C}_\infty^+$ into an assortment with an expected revenue of $\mathcal{R}_M(S^+ \cup S^-) \geq (1-\epsilon) \cdot \mathcal{R}_M(S^*)$. Rather than enumerating over all possible choices for S^-, of which there could be exponentially-many, we make a trivial choice for the latter and include each and every product in $\mathcal{C}_\infty^+$. The following lemma shows that the assortment $S^+ \cup \mathcal{C}_\infty^+$ indeed nearly matches the optimal expected revenue.

Lemma 19. $\mathcal{R}_M(S^+ \cup \mathcal{C}_\infty^-) \geq (1-3\epsilon) \cdot \mathcal{R}_M(S^*)$.

Proof. Our analysis is based on establishing the next two claims, whose proofs are deferred to Appendices F.5 and F.6, respectively.

Claim 7. $\sum_{i \in S^+} r_i \cdot \pi_M(i, S^+ \cup \mathcal{C}_\infty^-) \geq \mathcal{R}_M(S^+) - \epsilon \cdot \mathcal{R}_M(S^*)$.

Claim 8. *For every product $i \in S^-$, we have*

$$\pi_M\left(i, S^+ \cup \mathcal{C}_\infty^-\right) \geq (1-\epsilon) \cdot \pi_M\left(i, S^+ \cup S^-\right) - \frac{\epsilon}{n} \cdot \Pr\left[M > 0\right] \cdot \frac{w_i}{1+w_i} \ .$$

As a result, the expected revenue of $S^+ \cup \mathcal{C}_\infty^-$ can be lower-bounded by noting that

$$\begin{aligned}
\mathcal{R}_M\left(S^+ \cup \mathcal{C}_\infty^-\right) &= \sum_{i \in S^+} r_i \cdot \pi_M(i, S^+ \cup \mathcal{C}_\infty^-) + \sum_{i \in \mathcal{C}_\infty^-} r_i \cdot \pi_M(i, S^+ \cup \mathcal{C}_\infty^-) \\
&\geq \sum_{i \in S^+} r_i \cdot \pi_M(i, S^+) - \epsilon \cdot \mathcal{R}_M\left(S^*\right) \\
&\quad + \sum_{i \in S^-} r_i \cdot \left((1-\epsilon) \cdot \pi_M\left(i, S^+ \cup S^-\right) - \frac{\epsilon}{n} \cdot \Pr\left[M > 0\right] \cdot \frac{w_i}{1+w_i} \right) \\
&\geq \sum_{i \in S^+} r_i \cdot \pi_M(i, S^+ \cup S^-) + (1-\epsilon) \cdot \sum_{i \in S^-} r_i \cdot \pi_M\left(i, S^+ \cup S^-\right) \\
&\quad - \epsilon \cdot \mathcal{R}_M\left(S^*\right) - \frac{\epsilon}{n} \cdot \underbrace{\sum_{i \in S^-} \Pr\left[M > 0\right] \cdot \frac{r_i w_i}{1+w_i}}_{\leq \mathcal{R}_M(S^*)} \\
&\geq (1-\epsilon) \cdot \mathcal{R}_M\left(S^+ \cup S^-\right) - 2\epsilon \cdot \mathcal{R}_M\left(S^*\right) \\
&\geq (1-3\epsilon) \cdot \mathcal{R}_M\left(S^*\right) \ .
\end{aligned}$$

Here, the first inequality follows from the conjunction of Claims 7 and 8. The second inequality holds since $\pi_M(i, S^+) \geq \pi_M(i, S^+ \cup S^-)$ for every $i \in S^+$. The last inequality is obtained by recalling that $\mathcal{R}_M(S^+ \cup S^-) \geq (1-\epsilon) \cdot \mathcal{R}_M(S^*)$. $\square$

F.5. Proof of Claim 7

We first note that for every product $i \in S^+$ and $m \in [n]$,

$$\begin{aligned}
\pi_m\left(i, S^+ \cup \mathcal{C}_\infty^-\right) &\geq \Pr\left[[U_i > U_0] \wedge \left[\text{rank}(i, S^+) \leq m\right] \wedge \overline{\biguplus_{j \in \mathcal{C}_\infty^-} [U_j > U_i]} \right] \\
&\geq \Pr\left[[U_i > U_0] \wedge \left[\text{rank}(i, S^+) \leq m\right] \right] - \Pr\left[[U_i > U_0] \wedge \biguplus_{j \in \mathcal{C}_\infty^-} [U_j > U_i] \right] \\
&= \pi_m\left(i, S^+\right) - \Pr\left[\biguplus_{j \in \mathcal{C}_\infty^-} [U_j > U_i > U_0] \right] \\
&\geq \pi_m\left(i, S^+\right) - \frac{\epsilon}{n^9} \cdot \frac{w_i}{1+w_i} \ .
\end{aligned}$$

Here, the last inequality holds since, for every product $j \in \mathcal{C}_\infty^-$,

$$\Pr\left[U_j > U_i > U_0\right] = \frac{w_j}{w_j + w_i + 1} \cdot \frac{w_i}{1+w_i} \leq w_j \cdot \frac{w_i}{1+w_i} \leq \frac{\epsilon}{n^{10}} \cdot \frac{w_i}{1+w_i} \ ,$$

where the last inequality follows by observing that the weight of any product in $\mathcal{C}_\infty^-$ is at most $\frac{\epsilon}{n^{10}}$, by definition of $\mathcal{C}_\infty^-$. Consequently,

$$\begin{aligned}
\sum_{i \in S^+} r_i \cdot \pi_M\left(i, S^+ \cup \mathcal{C}_\infty^-\right) &= \sum_{i \in S^+} r_i \cdot \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \pi_m\left(i, S^+ \cup \mathcal{C}_\infty^-\right) \\
&\geq \sum_{i \in S^+} r_i \cdot \sum_{m \in [n]} \Pr\left[M = m\right] \cdot \left(\pi_m\left(i, S^+\right) - \frac{\epsilon}{n^9} \cdot \frac{w_i}{1+w_i}\right) \\
&\geq \mathcal{R}_M\left(S^+\right) - \frac{\epsilon}{n} \cdot \underbrace{\sum_{i \in S^+} \Pr\left[M > 0\right] \cdot \frac{r_i w_i}{1+w_i}}_{\leq \mathcal{R}_M(S^*)} \\
&\geq \mathcal{R}_M\left(S^+\right) - \epsilon \cdot \mathcal{R}_M\left(S^*\right) \ .
\end{aligned}$$

F.6. Proof of Claim 8

Focusing on a single product $i \in S^-$, we show that for every $m \in [n]$, one has

$$\pi_m\left(i, S^+ \cup \mathcal{C}_\infty^-\right) \geq (1-\epsilon) \cdot \pi_m\left(i, S^+ \cup S^-\right) - \frac{\epsilon}{n} \cdot \frac{w_i}{1+w_i} \ .$$

The desired claim would then follow when both sides are multiplied by $\Pr[M = m]$ and summed over all $m \in [n]$. We divide our analysis to two cases, depending on how m and $|S^+|$ are related.

Case 1: $m \leq |S^+|$. In this case,

$$\begin{aligned}
\pi_m\left(i, S^+ \cup \mathcal{C}_\infty^-\right) &\geq \Pr\left[\left[U_i > U_0\right] \wedge \left[\operatorname{rank}\left(i, S^+ \cup \mathcal{C}_\infty^-\right) \leq m\right] \wedge \overline{\biguplus_{j_1 \in S^+} \biguplus_{j_2 \in \mathcal{C}_\infty^- \setminus \{i\}} \left[U_{j_2} > U_{j_1}\right]}\right] \\
&= \Pr\left[\left[U_i > U_0\right] \wedge \left[\operatorname{rank}\left(i, S^+ \cup \{i\}\right) \leq m\right] \wedge \overline{\biguplus_{j_1 \in S^+} \biguplus_{j_2 \in \mathcal{C}_\infty^- \setminus \{i\}} \left[U_{j_2} > U_{j_1}\right]}\right] \\
&\geq \Pr\left[\left[U_i > U_0\right] \wedge \left[\operatorname{rank}\left(i, S^+ \cup \{i\}\right) \leq m\right]\right] \\
&\quad - \Pr\left[\left[U_i > U_0\right] \wedge \biguplus_{j_1 \in S^+} \biguplus_{j_2 \in \mathcal{C}_\infty^- \setminus \{i\}} \left[U_{j_2} > U_{j_1}\right]\right] \\
&= \pi_m\left(i, S^+ \cup \{i\}\right) - \Pr\left[U_i > U_0\right] \cdot \Pr\left[\biguplus_{j_1 \in S^+} \biguplus_{j_2 \in \mathcal{C}_\infty^- \setminus \{i\}} \left[U_{j_2} > U_{j_1}\right]\right] \\
&\geq \pi_m\left(i, S^+ \cup \{i\}\right) - \frac{1}{n^3} \cdot \frac{w_i}{1+w_i} \\
&\geq \pi_m\left(i, S^+ \cup S^-\right) - \frac{\epsilon}{n} \cdot \frac{w_i}{1+w_i} \ .
\end{aligned}$$

Here, the first equality is obtained by noting that $m \leq |S^+|$. The second equality holds since the events $[U_i > U_0]$ and $\biguplus_{j_1 \in S^+} \biguplus_{j_2 \in \mathcal{C}_\infty^- \setminus \{i\}} [U_{j_2} > U_{j_1}]$ are independent. The next-to-last inequality holds since, due to having a multiplicative gap of at least n^5 between weights in $\mathcal{C}_\infty^+$ and $\mathcal{C}_\infty^-$, for every $j_1 \in S^+ \subseteq \mathcal{C}_\infty^+$ and $j_2 \in \mathcal{C}_\infty^-$ we have $\Pr[U_{j_2} > U_{j_1}] = \frac{w_{j_2}}{w_{j_1} + w_{j_2}} \leq \frac{1}{n^5}$.

Case 2: $m > |S^+|$. On the one hand,

$$\pi_m\left(i, S^+ \cup S^-\right) \leq \pi_m\left(i, S^+ \cup \{i\}\right) = \Pr\left[U_i > U_0\right] = \frac{w_i}{1+w_i} ,$$

where the first equality holds since $m > |S^+|$. On the other hand,

$$\begin{aligned} \pi_m\left(i, S^+ \cup \mathcal{C}_\infty^-\right) &\geq \Pr\left[\mathrm{rank}\left(i, \mathcal{C}_\infty^- \cup \{0\}\right) = 1\right] \\ &= \frac{w_i}{1 + w_i + w(\mathcal{C}_\infty^- \setminus \{i\})} \\ &\geq \frac{w_i}{1 + w_i + \epsilon/n^9} \\ &\geq (1-\epsilon) \cdot \frac{w_i}{1+w_i} , \end{aligned}$$

where the second inequality is obtained by noting that $w(\mathcal{C}_\infty^- \setminus \{i\}) \leq \frac{\epsilon}{n^9}$, since the weight of any product in $\mathcal{C}_\infty^-$ is at most $\frac{\epsilon}{n^{10}}$, by definition of $\mathcal{C}_\infty^-$. Now, combining both bounds, it follows that $\pi_m(i, S^+ \cup \mathcal{C}_\infty^-) \geq (1-\epsilon) \cdot \pi_m(i, S^+ \cup S^-)$.

Appendix G: Traditional Choice Models in a Multi-Purchase Setting

In this section, we present numerical experiments that reveal the pitfalls of employing traditional single-purchase choice models in settings where multi-purchase behavior is present. Specifically, we show that mixed-MNL models, which are arguably the most general single-purchase RUM-based choice model (McFadden and Train 2000), poorly perform in their attempt to recapture MP-MNL models from sales data. Our intent is to highlight the importance of explicitly modeling multi-purchase dynamics when such patterns are prevalent.

To accomplish the above-mentioned goal, we start by randomly generating a diverse collection of MP-MNL models, which capture varying degrees of multi-purchase tendencies. For each distinct model, we generate synthetic sales data under the assumption that customers make purchasing decisions according to the given MP-MNL model. We then fit mixed-MNL models with up to five customer segments to this sales data via maximum likelihood estimation; here, in order to adapt this traditional estimation procedure to a multi-purchase setting, we make the natural simplifying assumption that all purchase events are mutually independent. Next, we measure the mixed-MNL model's ability to adapt to this multi-purchase setting via the following two accuracy metrics. First, we compute the absolute relative error of their predicted choice probabilities for every product within each potential offer set. Second, we assess the extent to which the mixed-MNL fits lead to profitable assortment recommendations.

G.1. Experimental Set-up

Generating sales data. To start, we generate an extensive collection of MP-MNL models, each with $n = 20$ products, which will serve as the ground-truth models throughout this section. For each distinct model, we sample its preference weights from a log-normal distribution with location parameter 0 and scale parameter 1, and then set the revenue of each product to be $r_i = w_{\max} - w_i$, so that larger weight products are associated with smaller revenues. We note that the product revenues will not play a role in generating the sales data, but will come into play later on, in assessing the mixed-MNL model fits. For each MP-MNL model, we assume that $\Pr[M = m] = 1$ for some $m \in \{2, 3, 4\}$. More precisely, for each $m \in \{2, 3, 4\}$, we generate ten MP-MNL models, giving us a total of 30 distinct instantiations of our multi-purchase model of interest. Next, for each of these 30 models, we simulate the arrival of $T = 1000$ customers. (We have also conducted experiments with $T \in \{2500, 5000\}$, whose results remained unchanged, and hence for brevity, we decided to present only the results for $T = 1000$.) We assume that the retailer offers assortment S_t of size exactly C to customer t, which is uniformly sampled from the $\binom{20}{C}$ possible options; we vary $C \in \{6, 8\}$. Customers choose according to the given MP-MNL, where we use $Z_t \subseteq S_t \cup \{0\}$ to denote the set of products purchased by customer t. The end result consists of 60 distinct streams of choice data (one for each model and C-value), each of which can be concisely written as $\mathrm{CD} = \{(S_t, Z_t) : t \in [T]\}$, to which we fit mixed-MNL models using the procedure described next.

Fitting mixed-MNL models. Under the mixed-MNL model, the customer population is partitioned into G segments, each making purchasing decisions according to a distinct MNL model. As such, the probability that customer t purchases product $i \in S_t$ can be expressed as

$$\pi_{\mathrm{MMNL}}(i, S_t) = \sum_{g \in [G]} \theta_g \cdot \frac{w_{ig}}{1 + \sum_{j \in S_t} w_{jg}},$$

where θ_g and w_{ig} are the respective arrival probability and product-i weight associated with segment g. Moving forward, we use $\theta = \{\theta_g : g \in [G]\}$ to denote the set of all arrival probabilities, and $w_g = \{w_{ig} : i \in [n]\}$ to denote the set of weights associated with segment g.

We employ maximum likelihood estimation (MLE) to fit the mixed-MNL models. However, this procedure requires a careful handling of each multi-purchase event, corresponding to customers $t \in [T]$ for which $|Z_t| > 1$, since the mixed-MNL model cannot explicitly account for such purchasing dynamics. To side-step this issue, we make the natural assumption that each purchase is the result of an MNL-based choice from among all offered products, implying that the likelihood of observing (S_t, Z_t) can be written as

$$\mathcal{L}(w_1, \ldots, w_G, \theta \mid (S_t, z_t)) = \prod_{i \in Z_t} \pi_{\mathrm{MMNL}}(i, S_t).$$

Given this assumption, for choice data CD, we can write the log-likelihood of the mixed-MNL model as

$$\mathcal{LL}(w_1, \ldots, w_G, \theta \mid \mathrm{CD}) = \sum_{t \in [T]} \sum_{i \in Z_t} \log\left(\pi_{\mathrm{MMNL}}(i, S_t)\right).$$

Therefore, the MLE problem of interest is

$$\begin{array}{ll} \max\limits_{\theta \in [0,1]^G, w} & \mathcal{LL}(w_1, \ldots, w_G, \theta \mid \mathrm{CD}) \\ \text{s.t.} & \sum\limits_{i \in G} \theta_i = 1 \end{array} \tag{MLE-MMNL}$$

For each stream of choice data, we solve (MLE-MMNL) for every $G \in \{1, \ldots, 5\}$ by directly maximizing the log-likelihood function using MATLAB's constrained nonlinear solver *fmincon* (MATLAB Optimization Toolbox). Although this function is known to be concave only for $G = 1$, we found that *fmincon* generally converged to a local optima. Moreover, as can be gleaned from Table 2, which displays the average training log-likelihoods for each value of G, we do indeed see the in-sample fitting accuracy improve as we increase the number of customer segments within the fitted mixed-MNL models. This trend supports the notion that the local optima we uncover are reasonably good.

	Avg. Log-Like	
G	$C=6$	$C=8$
1	-1785	-1998
2	-1774	-1987
3	-1764	-1979
4	-1756	-1972
5	-1752	-1969

Table 2 Average training log-likelihoods of the mixed-MNL fits.

Accuracy metrics. In what follows, we formally define two accuracy metrics, which are utilized to assess the mixed-MNL model fits. For this purpose, we let $\pi(i,S)$ denote the choice probabilities attributed to the ground-truth MP-MNL model, and use $\mathcal{R}(S)=\sum_{i\in S} r_i\cdot\pi(i,S)$ to denote the expected revenue earned from offering the assortment S. Our first accuracy metric is the average percent absolute error in the choice probabilities, computed as

$$\frac{100}{\binom{20}{C}}\cdot\sum_{S:|S|=C}\sum_{i\in S}\frac{|\pi_{\mathrm{MMNL}}(i,S)-\pi(i,S)|}{\pi(i,S)}.$$

Next, we examine whether the mixed-MNL fits lead to profitable assortment recommendations. For this purpose, we first compute $S^*_{\mathrm{MMNL}}=\arg\max_{S:|S|=C}\sum_{i\in S} r_i\cdot\pi_{\mathrm{MMNL}}(i,S)$, which corresponds to the optimal assortment if customers were to choose according to the fitted mixed-MNL model. Additionally, we also compute $S^*=\arg\max_{S:|S|=C}\mathcal{R}(S)$, which is the true optimal assortment. We note that both S^*_{MMNL} and S^* can be efficiently computed via brute force enumeration, since we consider instances with only 20 products. We then measure the revenue performance of S^*_{MMNL} through its percent optimality gap, given by $100\cdot\frac{\mathcal{R}(S^*)-\mathcal{R}(S^*_{\mathrm{MMNL}})}{\mathcal{R}(S^*)}$.

G.2. Results

The results of our experiments are presented in Tables 3a and 3b, where we show the average optimality gap and average absolute error for the five fitted mixed-MNL models. Somewhat surprisingly, we observe that the standard MNL fits ($G=1$) performed best across both accuracy metrics. This trend is likely a consequence of how the mixed-MNL fits approximate each multiple-purchase event within the sales data as a sequence of independent single-purchase events. This extra layer of approximation in turn provides intuition for why a better fit to the modified sales data does not necessarily imply better fits to the ground truth model. That said, it is very clear that these MNL fits are far from being accurate, with a worst case average optimality gap of 10.35% ($C=6$ and $m=4$) and a worst case absolute error of 50.56% ($C=8$ and $m=4$). Altogether, the observation that MNL fits perform best, while having the worst in-sample accuracy (see Table 2), indicates that completely ignoring the dynamics of multi-purchase behavior, upon fitting the mixed-MNL model, can be quite detrimental. Additionally, we generally observe that the average performance of each mixed-MNL model deteriorates as m is increased. In other words, the performance of these traditional single-purchase models suffers as multi-purchase events become more prevalent in the sales data, which further emphasizes the importance of explicitly modeling multi-purchase behavior.

G	Avg. % Opt. Gap $m=2$	$m=3$	$m=4$	Avg. % Abs. Error $m=2$	$m=3$	$m=4$
1	2.02	6.19	10.35	23.02	37.04	42.63
2	5.13	10.83	15.88	35.25	48.38	52.28
3	11.24	10.16	23.10	39.18	53.02	54.35
4	12.97	15.82	22.48	39.97	55.06	57.51
5	8.96	13.14	19.23	41.02	54.58	57.64

(a) Results for $C=6$.

G	Avg. % Opt. Gap $m=2$	$m=3$	$m=4$	Avg. % Abs. Error $m=2$	$m=3$	$m=4$
1	1.06	2.62	6.73	29.31	44.99	50.56
2	3.92	5.94	9.30	44.10	55.56	61.04
3	7.45	15.56	12.50	48.51	62.52	64.76
4	16.87	6.88	18.39	54.91	60.78	67.82
5	16.48	7.86	12.85	54.60	63.86	65.70

(b) Results for $C=8$.

Table 3 Accuracy metrics for the fitted mixed-MNL models

Appendix H: Computational Experiments

In this section, we conduct extensive computational experiments aimed at both measuring the efficacy of the approximation scheme we propose in Section 4, as well as at assessing the revenue-potential of making assortment decisions that account for multi-purchase behavior in comparison to those that overlook this phenomenon. To accomplish this goal, we employ our PTAS, along with two heuristic approaches, on a wide variety of randomly generated problem instances with up to 60 products. In addition, we take a deep-dive into the make-up of the assortments recommended under each approach, so as to better understand higher level trends that account for differences in revenue performance.

H.1. Instance generator

We randomly generate cardinality-constrained instances of the assortment optimization problem under the MP-MNL model with $n \in \{15, 30, 60\}$ products. Specifically, motivated by typical e-commerce product recommendation pages, we enforce an additional side constraint, stating that the offered assortment should include exactly $C \in \{8, 12\}$ products. Since we focus on the limited purchases setting of Section 4, we truncate the maximum number of purchases at $m_{\max} \in \{2, 3\}$. The preference weight and price of each product and the distribution of M are generated as follows.

- *Weights:* For each product $i \in [n]$, we generate its preference weight w_i from the log-normal distribution with location-scale $(0, 1)$, so as to ensure that we obtain a heterogeneous collection of weights that are unlikely to drastically vary within each instance. Let $w_{\max}$ be the largest weight generated for a particular instance.
- *Prices:* The price of each product $i \in [n]$ is set to be $r_i = w_{\max} - w_i$, so that higher priced items have smaller associated weights. At least intuitively, the assortment problem becomes more difficult to handle when the prices and weights are inversely related, since in this case, choosing an assortment that balances market share and per-purchase revenues is a trickier task.
- *Distribution of M:* We consider two distinct types of distributions for the random number of potential purchases, M. The first is generated by independently sampling a uniform random variable $\lambda_m \sim U[0, 1]$ for each $m \in [m_{\max}]$. These values are then normalized so that we are left with a valid distribution, where $\Pr[M = m] = \frac{\lambda_m}{\sum_{\mu \in [m_{\max}]} \lambda_\mu}$; we denote the resulting distribution as M_{rand}. The second distribution, denoted as M_{dec}, is generated by sorting the above-mentioned random samples in decreasing order, such that $\lambda_{(1)} \geq \cdots \geq \lambda_{(m_{\max})}$. As before, these values are normalized to obtain $\Pr[M = m] = \frac{\lambda_{(m)}}{\sum_{\mu \in [m_{\max}]} \lambda_\mu}$. The latter scenario is reflective of the purchasing patterns observed in Table 1, which exhibits a decreasing trend of multi-purchase events for customers participating in the campaigns of our e-retailer collaborator.

Given this generative model, for each of the 16 possible combinations of $n \in \{15, 30\}$, $C \in \{8, 12\}$, $m_{\max} \in \{2, 3\}$, and $M \in \{M_{\text{rand}}, M_{\text{dec}}\}$, we create a distinct test case. For each test case, we randomly generate 50 problem instances, as explained above, each characterized by a unique set of weights, prices, and λ-values. Additionally, to keep the overall CPU time of our simulations within a manageable scale, we carry out a smaller set of experiments with $n = 60$ products. Here, we fix $C = 12$ and only vary $m_{\max} \in \{2, 3\}$ and $M \in \{M_{\text{rand}}, M_{\text{dec}}\}$, thus yielding four test cases. For each of these test cases, we generate 10 problem instances. The intent of this smaller collection of experiments is to demonstrate that our PTAS scales to instances with a relatively large number of products.

H.2. Tested algorithms

For each problem instance generated through the above-mentioned parameter configurations, we employ the following algorithmic approaches:

1. *Cardinality-constrained PTAS.* This approach is the approximation scheme outlined in Section 4, with a slight adaptation to account for an exact cardinality constraint. In this setting, it is easy to verify that Lemma 5 remains fully intact, and hence, the dynamic-programming-based FPTAS for problem (5) should be extended to account for the added restriction that exactly $C - |\mathcal{H}^*|$ light products should be chosen. This extension is straightforward, simply by adding an extra state parameter that tracks the total number of light products added thus far. Moreover, to ensure an efficient implementation, we execute our guessing procedure only for the $\frac{m_{\max}}{\epsilon}$ highest weight products (rather than $\frac{6m_{\max}^2}{\epsilon}$). While this feature degrades our worst-case theoretical guarantee to $(1-\epsilon)^{6m_{\max}}$, our experimental results show that the practical performance of our approximation scheme is near-optimal.

2. *Cardinality-constrained MNL.* As explained in Section 3, this approach assumes that our representative customer is willing to purchase exactly one product. Thus, the resulting "proxy" problem reduces to cardinality-constrained assortment optimization under the standard MNL model, which is known to be polynomial-time solvable (Rusmevichientong et al. 2010, Sumida et al. 2021). Due to efficiency considerations, we have implemented the LP-based algorithm proposed by Sumida et al. (2021). It is worth noting that Lemma 3 is no longer valid in the cardinality-constrained setting, where revenue-ordered assortments can be $\Omega(n)$-factor away from attaining optimal revenues in the worst case. This fact, however, does not diminish the importance of assessing the potential revenue loss that may result from ignoring multi-purchase behavior in practice (rather than with respect to worst-case constructions).

3. *Greedy heuristic.* In this approach, starting with an empty set of products, we employ a sequence of augmentations. In each step, the product added to the current assortment is one whose inclusion maximizes the expected revenue of the resulting assortment. This iterative procedure continues up until the selection of C products.

Technically speaking, our approximation scheme was implemented in Java 8, whereas the cardinality-constrained MNL approach and the greedy heuristic were both implemented in Python 3, with Gurobi 9.0.1 employed as an LP-solver. All experiments were executed on a standard desktop computer, equipped with an Intel Core i7 CPU and 64GB of RAM.

Performance metrics. For the remainder of this section, we make use of the shorthands PTAS, MNL, and GR for the cardinality-constrained PTAS, cardinality-constrained MNL approach, and greedy heuristic, respectively. Furthermore, for any algorithm $\mathcal{A} \in \{\text{PTAS}, \text{MNL}, \text{GR}\}$, let $S^{\mathcal{A}}$ be its returned assortment for a given instance. In order to assess the revenue guarantee of each approach, for test cases with $n = 15$ products, we employ exhaustive enumeration over all feasible assortments to recover the optimal assortment S^*. In this case, for each algorithm $\mathcal{A}$, we define its optimality gap with respect to a given instance as $\frac{\mathcal{R}(S^*) - \mathcal{R}(S^{\mathcal{A}})}{\mathcal{R}(S^*)}$. The performance of each approach on these small-scale instances is reported as basic statistics of the optimality gap, which ultimately allows us to show that our PTAS performs near-optimally even when seeded with a large accuracy level. For test cases with $n \in \{30, 60\}$ products, recovering the optimal expected revenue by means of exhaustive enumeration becomes computationally infeasible, and hence, we benchmark all approaches against PTAS. Specifically, for $\mathcal{A} \in \{\text{MNL}, \text{GR}\}$, we examine the percentage improvement attained by our approximation scheme, which is computed as $\frac{\mathcal{R}(S^{\text{PTAS}}) - \mathcal{R}(S^{\mathcal{A}})}{\mathcal{R}(S^{\mathcal{A}})}$.

H.3. Results

The results of our experiments for the test cases with $n = 15$, $n = 30$, and $n = 60$ products are respectively presented in Tables 4, 5, and 6. In each of these tables, we have a distinct row for each test case, with the first three columns listing the defining combination of the parameters $(C, m_{\max}, M)$. In Table 4, the next nine columns report the average, 95$^{\text{th}}$ percentile, and maximum optimality gaps for each of three implemented approaches across all 50 problem instances. In Table 5, the latter columns report these same metrics, albeit in relation to the percentage improvement of PTAS over GR and MNL, as explained earlier. Table 6 reports only the average and maximum percentage improvements of PTAS, since only 10 problem instances are generated for each test case in this larger-scale setting, which makes the 95$^{\text{th}}$ percentile identify with the maximal gap. Finally, we note that our approximation scheme was implemented with $\epsilon = 0.5$ for $n = 15$, $\epsilon = 0.6$ for $n = 30$, and $\epsilon = 0.75$ for $n = 60$. For these three sets of test cases, the respective per-instance computation times were approximately 15-30 seconds, 5-7 minutes, and 45-60 minutes.

			Avg. % Opt. Gap			95$^{\text{th}}$ Perc. % Opt. Gap			Max % Opt. Gap		
C	$m_{\max}$	M	PTAS	GR	MNL	PTAS	GR	MNL	PTAS	GR	MNL
8	2	M_{dec}	0.012	0.98	0.11	0.0	3.18	0.56	0.72	5.20	1.25
8	2	M_{rand}	0.0	0.92	0.14	0.0	2.81	0.90	0.0	5.15	1.39
12	2	M_{dec}	0.012	0.31	0.01	0.0	1.79	0.02	0.74	4.04	0.17
12	2	M_{rand}	0.0	0.19	0.0	0.0	1.71	0.01	0.0	3.10	0.02
8	3	M_{dec}	0.0	0.43	0.52	0.0	2.07	1.91	0.0	4.44	2.98
8	3	M_{rand}	0.0	0.43	0.99	0.0	1.80	2.77	0.0	3.06	5.18
12	3	M_{dec}	0.0	0.077	0.06	0.0	0.29	0.12	0.0	2.31	1.52
12	3	M_{rand}	0.0	0.16	0.05	0.0	1.30	0.44	0.0	3.19	1.34

Table 4 Average optimality gaps of PTAS, GR, and MNL for test cases with $n = 15$ products.

The results reported in all three tables provide strong evidence of the superiority of our approximation scheme in terms of revenue guarantees. First, Table 4 shows that this PTAS is near-optimal across all test cases, even when seeded with accuracy level of $\epsilon = 0.5$. In fact, its worst optimality gap across all test cases

			Avg. % Improve.		95$^{\text{th}}$ Perc. % Improve.		Max % Improve.	
C	$m_{\max}$	M	GR	MNL	GR	MNL	GR	MNL
8	2	M_{dec}	2.26	0.13	3.10	0.47	3.40	0.65
8	2	M_{rand}	2.22	0.38	3.21	0.92	3.66	2.04
12	2	M_{dec}	2.28	0.14	3.35	0.42	3.52	0.77
12	2	M_{rand}	2.30	0.26	3.07	0.81	3.77	1.23
8	3	M_{dec}	1.71	0.61	2.74	1.65	3.91	2.16
8	3	M_{rand}	1.46	1.10	2.44	2.06	2.63	3.09
12	3	M_{dec}	1.75	0.50	2.98	1.36	3.26	2.00
12	3	M_{rand}	1.68	0.88	2.76	2.18	3.12	2.78

Table 5 Percent improvement of PTAS over GR and MNL for test cases with $n = 30$ products.

with $n = 15$ is only 0.74%, while the maximum optimality gap for both GR and MNL exceeds 5% for these cases. Table 5 reveals that when we move to larger problem instances, consisting of 30 products, PTAS continues to outperform both GR and MNL by 0.5-2% on average, and by 3-4% in some cases. Finally, Table 6 illustrates that our PTAS can still be implemented even on instances with $n = 60$ products, where it continues to be superior even when we set $\epsilon = 0.75$.

			Avg. % Improve.		Max % Improve.	
C	$m_{\max}$	M	GR	MNL	GR	MNL
12	2	M_{dec}	3.14	0.12	3.70	0.21
12	2	M_{rand}	3.13	0.24	3.55	0.45
12	3	M_{dec}	2.92	0.40	3.47	0.93
12	3	M_{rand}	2.93	0.85	3.60	1.60

Table 6 Percent improvement of PTAS over GR and MNL for test cases with $n = 60$ products.

While our results suggest that PTAS is consistently the best performing approach, we also observe that MNL performs surprisingly well given that it completely ignores multi-purchase behavior. Across all test cases with $n = 30$ products, PTAS outperforms MNL by approximately 0.1-1% on average. However, as one can only expect, we observe that MNL's worst performance comes for test cases in which $m_{\max} = 3$, where the multi-purchase effects are more prominent. All-in-all, taking into account the results of our experiments as well as the theoretical guarantees developed in Lemmas 2 and 3, there is indeed some practical appeal in reducing the problem in question to a single-purchase setting in specific scenarios when the expected number of possible purchases is known to be very small, especially when an extremely efficient implementation is needed.

H.4. Analyzing the structure of the recommended assortments

In this section, we examine whether the revenue gains of PTAS over MNL can be attributed to fundamental differences in the make-up of the recommended assortments under each approach. For succinctness of presentation, we restrict our analysis to test cases with $n = 30$ products, noting that similar trends can be observed for the additional test cases.

To address our question of interest, we consider the subset of recommended products that is unique to each approach, and tease out salient differences in their composition. Specifically, for each problem instance, we first compute $S^{\mathrm{PTAS}} \setminus S^{\mathrm{MNL}}$ and $S^{\mathrm{MNL}} \setminus S^{\mathrm{PTAS}}$, whose cardinality is each 2.13 on average. Next, we examine the distributions of the preference weights of products contained in these two sets. As shown in Figure 1, it is very clear that S^{PTAS} differentiates itself from S^{MNL} through the inclusion of a few high-weight low-revenue products (recall that $r_i = w_{\max} - w_i$). Somewhat surprisingly, these high-weight products lead to only small gains in the overall number of expected purchases. Specifically, Figure 2 shows the distributions of the expected number of items purchased under the assortments S^{PTAS} and S^{MNL}, which only slightly differ. Additionally, to give a sense of the relative frequency with which customers select the no-purchase option, we also include the distribution of $\mathbb{E}[M]$ in these plots. Alternatively, it appears that these high-weight products allow our approximation scheme to capture significantly more revenue than MNL on purchases made beyond each customer's first choice product. In particular, Figure 3 shows exactly how the revenues of S^{PTAS} and S^{MNL} compare as customers expand their baskets; defining $\bar{R}_m(S) = \sum_{k\in[m]} \Pr[M=k] \cdot \mathcal{R}_m(S)$ to be the cumulative revenue earned from purchases up to the m-th most preferred product, we plot $\bar{R}_m(S^{\mathrm{PTAS}}) - \bar{R}_m(S^{\mathrm{MNL}})$ for $m \in [m_{\max}]$ across all problem instances. This figure clearly indicates that our PTAS exploits customers with $M > 1$ far more effectively than MNL, which again, is likely attributable to the inclusion of a few high-weight low-revenue products, as mentioned above.

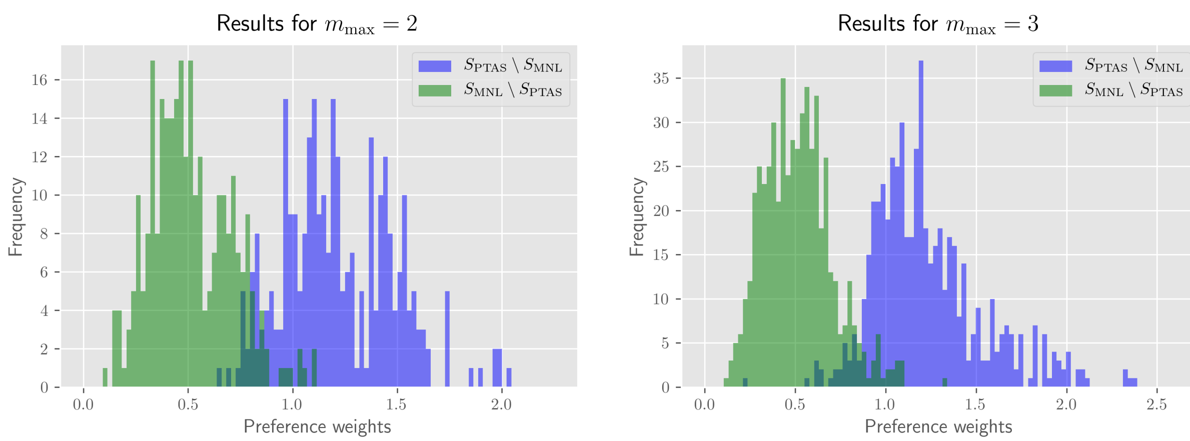

Figure 1 Preference weight distributions for recommended products unique to S^{PTAS} and S^{MNL}.

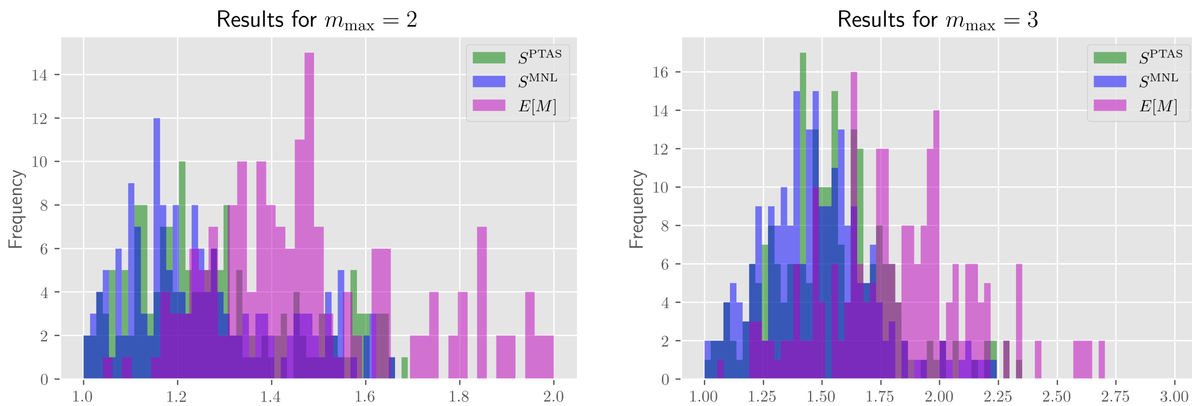

Figure 2 **Distributions of the expected number of items purchased under S^{PTAS} and S^{MNL}.**

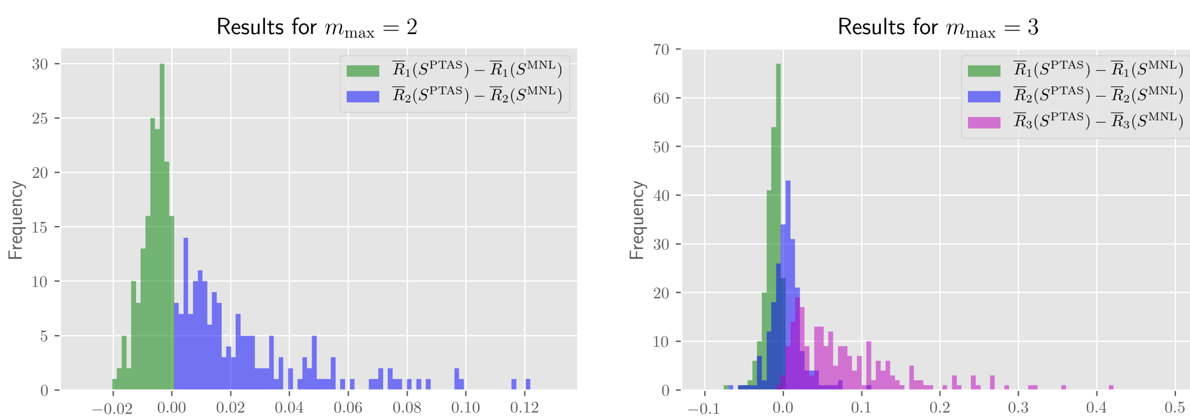

Figure 3 **Distributions of cumulative revenues under S^{PTAS} and S^{MNL}.**