Greedy-Like Algorithms for Dynamic Assortment Planning Under Multinomial Logit Preferences

Ali Aouad* Retsef Levi† Danny Segev‡

Abstract

We study the joint assortment planning and inventory management problem, where stock-out events elicit *dynamic substitution* effects, described by the Multinomial Logit (MNL) choice model. Special cases of this setting have extensively been studied in recent literature, notably the *static* assortment planning problem. Nevertheless, the general formulation is not known to admit efficient algorithms with analytical performance guarantees prior to this work, and most of its computational aspects are still wide open.

In this paper, we devise the first provably-good approximation algorithm for dynamic assortment planning under the MNL model, attaining a constant-factor guarantee for a broad class of demand distributions, that satisfy the *increasing failure rate* property. Our algorithm relies on a combination of greedy procedures, where stocking decisions are restricted to specific classes of products and the objective function takes modified forms. We demonstrate that our approach substantially outperforms state-of-the-art heuristic methods in terms of performance and speed, leading to an average revenue gain of 4% to 12% in computational experiments. In the course of establishing our main result, we develop new algorithmic ideas that may be of independent interest. These include weaker notions of submodularity and monotonicity, shown sufficient to obtain constant-factor worst-case guarantees, despite using noisy estimates of the objective function.

Keywords: Inventory management, multinomial logit, submodularity, stochastic modeling, approximation algorithm.

*London Business School, London NW1 4SA, United Kingdom. Email: aouad.ma@gmail.com.

†Sloan School of Management, Massachusetts Institute of Technology,77 Massachusetts Avenue, Cambridge, Massachusetts 02139, Email: retsef@mit.edu.

‡Department of Statistics, University of Haifa, Haifa 31905, Israel, Email: segevd@stat.haifa.ac.il.

1 Introduction

In the last two decades, growing product differentiation in brick-and-mortar retailing, airlines, and consumer-goods has motivated a paradigm shift in demand modeling from independent demand models to choice-based models. The latter capture the substitution effects between competing products in a given category by modeling customers as agents with random preferences. An important line of research in revenue management has focused on incorporating such choice behavior into more realistic decision models for inventory management and assortment planning. In particular, the *static assortment optimization* model takes the perspective of a retailer who wishes to determine the best-possible collection of products to maximize her expected revenue. In this setting, in order to model the customers' purchase preferences, the Multinomial Logit model (MNL) has gained widespread popularity among practitioners, since it can be estimated efficiently, even in data-scarce environments (Ford 1957, McFadden 1973), and it leads to tractable assortment planning formulations (Talluri and Van Ryzin 2004, Rusmevichientong et al. 2010, 2014).

However, this static modeling approach overlooks the inventory limitations faced by retailers and the resulting mismatch between supply and demand under uncertainty. Due to product proliferation, *out-of-stock* and *over-stock* situations are pervasive in retailing, leading to billions of dollars of opportunity loss and reputation cost (IHL 2015), and consequently, retailers need to manage their supply chain to ensure the availability of products to end customers. Thus, a more general class of optimization models, known as *dynamic assortment planning*, was proposed in the seminal work of Mahajan and van Ryzin (2001) to capture the customers' dynamic substitution behavior elicited by such stock-out events. Here, the inventory decisions are constrained by a joint capacity limitation across the assortment in the form of a cardinality constraint. Rather than focusing on a single representative customer, the consumption process is modeled through a random sequence of arriving customers, each having random preferences over the products on stock upon arrival. In this setting, the initial assortment and inventory decisions made by the retailer need to be robust (revenue-wise) to the stock-out events. From a computational standpoint, in addition to the underlying choice model, the problem formulation now hinges on describing the distribution of the number of customer arrivals, named the *demand* hereafter. A more formal description of the optimization model depicted above is given in Section 1.2.

Such dynamic problems are computationally challenging and notoriously difficult to analyze (Mahajan and van Ryzin 2001, Honhon et al. 2010, Aouad et al. 2015), due to the complex assortment dynamics induced by the stochastic nature of the customer arrivals and their respective random preferences. In fact, even the efficient evaluation of the expected revenue generated by given assortment and inventory decisions is an open question for most choice models of interest, including the MNL model. As a result, the vast majority of computational questions regarding dynamic assortment planning are still wide open, and were previously treated by heuristic methods or stylized models, as summarized in Section 1.3. Despite the centrality of the Multinomial Logit choice model in revenue management applications, obtaining efficient algorithms with analytical guarantees for MNL-based dynamic assortment planning models has been a long-standing open question.

1.1 Results and techniques

The main contribution of this paper is to devise the first provably-good approximation algorithm for dynamic assortment planning under the Multinomial Logit choice model. Specifically, our approach guarantees a constant-factor worst-case approximation for a broad class of demand distributions (standing for the total number of arriving customers) commonly used in operations management, that satisfy the *increasing failure rate* (IFR) property. Moreover, we show that this algorithm has a superior empirical performance in comparison to existing heuristics on synthetic instances. Against existing state-of-the-art methods, our algorithm leads to substantial gains in the expected revenue, ranging from 4% to 12%, with better computational efficiency and robustness. Our algorithmic approach relies on a combination of greedy procedures, where stocking decisions are restricted to specific assortments, and the objective function takes a modified form. Our theoretical analysis along with the experimental results provide evidence that such restrictive policies could in fact be more effective than general-purpose methods, that consider stocking decisions across *all* products. Along the way, we develop a number of novel technical ideas that could very well contribute to studying additional combinatorial optimization problems and to assortment planning methodologies in particular.

Restricted-submodular maximization. At the core of our analysis, we develop new concepts of submodularity and monotonicity, called the *restricted-submodular* and *restricted-non-decreasing* properties, that are weaker than their standard counterparts. Specifically, when optimizing certain set functions under a cardinality constraint, the objective function could generally violate the submodularity property, while still having a submodular-like behavior within the feasible collection of sets, i.e., those satisfying the cardinality constraint. Thus, in the restricted-submodular setting, the structural inequalities defining submodularity and monotonicity are not required uniformly over all sets, and instead, we restrict attention to the feasible region only. We show that the classic analysis of greedy algorithms extends to this broader setting, and obtain a $(0.318 - \epsilon)$-approximation. Moreover, this worst-case guarantee holds with high probability even when the greedy procedure is given access to noisy estimates of the objective function at each step.

Algorithmic approach and performance guarantees. For ease of presentation, we describe our approach in an incremental way, where a simplified setting is first examined, prior to addressing the most general case, thereby establishing the following worst-case guarantees.

- *Core algorithm with evaluation oracle.* As previously mentioned, in dynamic substitution models, it is generally unknown how to efficiently compute the expected revenue generated by given initial inventory levels. To bypass this difficulty, we first operate under the *efficient oracle assumption*, where we temporarily assume that the expected revenue function can be efficiently evaluated with high probability by some (unspecified) oracle procedure. Under this assumption, we devise in Section 3 a polynomial-time algorithm with a constant-factor worst-case guarantee, for any demand distribution with increasing failure rates. Specifically, for any error parameters $\epsilon \in (0, 1/4)$ and $\delta > 0$, our randomized algorithm attains a $(0.139 - \epsilon)$-approximation with probability at least $1 - \delta$. Moreover, our methods are amenable to tighter analysis under more restrictive settings, allowing us to obtain an approximation guarantee of $0.179 - \epsilon$ under the plausible assumption that the number of inventory units (or capacity) exceeds the number of products, and $0.632 - \epsilon$ when only products within an optimal static assortment can be stocked. The latter result holds for general (non-IFR) demand distributions.

Technically speaking, for large enough capacity values, the algorithm concurrently runs two greedy procedures: each restricts attention to a specific class of products, and the inventory levels are chosen greedily over the residual set of products, using a modified objective function. Our analysis relies in large part on the restricted-submodular and restricted-non-decreasing properties mentioned earlier. Indeed, after we interpret one residual problem in terms of optimizing a set function, we show that, while the latter generally violates the standard properties of submodularity and monotonicity, it still satisfies their weaker (restricted) version. The proof is based on novel probabilistic coupling ideas, allowing us to compare the dynamic substitution patterns driven by the MNL model. A particularly interesting byproduct of our analysis is showing that a commonplace heuristic, which stocks the optimal static assortment and scales inventory proportional to the expected sales, has a provable performance guarantee with respect to a restricted class of products.

- *General approximation algorithm.* In Appendix A, we bypass the efficient oracle assumption, and derive a general constant-factor approximation for dynamic assortment planning under the MNL choice model, with increasing failure rate demand distributions. For any $\epsilon \in (0, 1/4)$ and $\delta > 0$, we devise a randomized polynomial-time algorithm attaining a worst-case guarantee of $0.122 - \epsilon$ with probability at least $1 - \delta$, which is improved to $0.151 - \epsilon$ when the capacity exceeds the number of products.

Empirical evaluation. While our theoretical worst-case guarantees might look unsatisfactory for practical purposes, we present in Section 4 extensive computational experiments, showing that the resulting algorithm largely outperforms existing heuristics in terms of performance and speed. These experiments employ our algorithm on randomly-generated instances, concurrently to the following heuristics: (i) a local-search heuristic based on greedily exchanging units between pairs of products, similar to Goyal et al. (2016); (ii) a gradient-descent approach based on a continuous extension of the revenue function, similar in spirit to the work of Mahajan and van Ryzin (2001); (iii) exact dynamic programming for two variants of the problem formulated by Topaloglu (2013), based on a Poisson and a normal approximation of the demand process; (iv) the deterministic relaxation heuristic proposed by Honhon et al. (2010), implemented using a commercial integer programming solver; (v) a discrete-greedy algorithm, where in each step a single unit is added to the product with the largest marginal expected revenue. Against these benchmarks, our algorithm attains expected revenues that are better by a factor ranging between 4% and 12%, and simultaneously dominates all methods in 62% of the instances tested. We also report that the proportional scaling heuristic, used as a subroutine in Section 3.3, is outperformed by the overall algorithm on average by 5.5%. In addition, the running time of our algorithm is significantly shorter than the above-mentioned heuristics, at the exception of the normal-based dynamic program.

1.2 Problem formulation

We are given n products, where each product $i \in [n]$ is associated with a preference weight w_i and a per-unit selling price r_i. In addition, there is a capacity bound of C on the total number of units to be stocked. In the dynamic assortment planning problem, the retailer has to jointly decide on an assortment, i.e., a subset of products to be offered, as well as on the initial inventory levels of these products, which are not replenished later on. That is, a feasible solution specifies the initial inventory levels of all products, represented by an integer-valued vector $U = (u_1, \dots, u_n)$ that meets the capacity constraint, $\sum_{i=1}^n u_i \le C$.

Stochastic MNL-based consumption process. We proceed by providing the additional model ingredients that describe the process according to which customers arrive and purchase products over time. A random number of customers M arrive sequentially, where the distribution of M is known to the decision-maker. Upon the arrival of a customer, suppose that $S \subseteq [n]$ is the subset of products that are currently available, due to being initially stocked, and not depleted until now. Then, this customer either:

- Picks a random product out of S and purchases a single unit, where the probability for choosing product $i \in S$ is $w_i/(1 + w(S))$. Here, $w(S)$ stands for the total weight of the products in S, i.e., $w(S) = \sum_{j \in S} w_j$.
- Leaves without purchasing any product, which happens with probability $1/(1 + w(S))$.

Objective function. When the sequence of customer arrivals ends, we use $\mathcal{R}(U)$ to denote the revenue resulting from an initial inventory vector U. This revenue is clearly random, due to the stochasticity in the number of customers and in their choice of products to purchase. The objective is to compute a feasible inventory vector, so that the expected revenue is maximized,

$$\max_{(u_1,\dots,u_n) \in \mathbb{Z}_+^n} \left\{ \mathbb{E}\left[\mathcal{R}(u_1, \dots, u_n)\right] : \sum_{i=1}^n u_i \le C \right\} .$$

The IFR assumption. As mentioned in Section 1.1, the distribution of the number of customers M is assumed to have an increasing failure rate (IFR), meaning that the sequence $\Pr[M = k] / \Pr[M \ge k]$ is non-decreasing over the integer domain. For definitions of stochastic orders and stochastic monotonicity, we refer the reader to Shaked and Shanthikumar (1994). It is worth mentioning that the IFR property is satisfied by many distributions considered in operations management applications, including Normal, Exponential, Geometric, Poisson, and Beta (for certain parameters).

Remark. Without an additional capacity constraint (i.e., when $C = \infty$), there is an infinite supply of inventory for each product, in which case it is not difficult to verify that the problem reduces to its static assortment planning formulation. Essentially, since there are no stock-out events, we can offer the optimal static assortment to each arriving customer.

1.3 Related literature

The MNL choice model. The Multinomial Logit (MNL) model is arguably the most widespread approach for modeling choice among practitioners, as reflected by seminal studies in transportation (McFadden 1980, Ben-Akiva and Lerman 1985) and marketing (Guadagni and Little 1983, Grover and Vriens 2006, Chandukala et al. 2008). This model, proposed independently by Luce (1959) and Plackett (1975), is grounded in economic theory of utility maximization, and describes the probabilistic choice outcomes of a representative agent who maximizes his utility over different alternatives, through a noisy evaluation of the utility they procure. The popularity of this model was notably driven by its simple estimation procedures (McFadden 1973, Talluri and Van Ryzin 2004, Maystre and Grossglauser 2015), even with limited data (Ford 1957, Negahban et al. 2012), as well as by its computational tractability in decision-making problems. Indeed, the static assortment planning problem is well-understood in the context of MNL choice preferences. For the uncapacitated variant, where any number of products can be offered, Talluri and Van Ryzin (2004) showed that the optimal assortment consists of the k-highest price products, for some k, and can therefore be computed efficiently. Rusmevichientong et al. (2010) devised a polynomial-time algorithm for the capacitated variant, where an upper bound is imposed on the number of products offered. These results were further advanced to handle more general settings (Rusmevichientong and Topaloglu 2012, Rusmevichientong et al. 2014), including a linear programming approach proposed by Davis et al. (2013) and a local-ratio framework developed by Désir et al. (2015).

Challenges in dynamic assortment planning. Under multiple stochastic arrivals, the problem we study becomes considerably more challenging than its static counterpart, due to the additional 'dynamic' aspect. Indeed, the assortment is altered along the sequence of arrivals due to stock-out events, as customers purchase the most preferred product available according to a probabilistic choice model. Therefore, the substitution behavior of customers depends on each sample-path realization, and a large number of samples is generally needed to obtain accurate estimates of the expected revenue function. In addition, this function violates several well-behaved properties. For instance, under a general model of choice, for a continuous relaxation of the dynamic assortment problem, Mahajan and van Ryzin (2001) showed that the revenue function is not even quasiconcave. Aouad et al. (2015) demonstrated through various counterexamples that this function (in modified form) is not submodular, even for very simple choice modeling approaches.

Existing methods. As a result, most of the work we are aware of in the context of dynamic assortment planning develops heuristics based on continuous relaxations and probabilistic assumptions (Smith and Agrawal 2000, Mahajan and van Ryzin 2001, Gaur and Honhon 2006, Nagarajan and Rajagopalan 2008, Honhon et al. 2010, Honhon and Seshadri 2013). These approaches either give rise to exponential-time algorithms, apply to more stylized models, or converge to local optima, such as the gradient-descent method proposed by Mahajan and van Ryzin (2001).

Similar to the present setting, Topaloglu (2013) studied a joint assortment and inventory management model with sequential customer arrivals, MNL preferences, and exogenous per-unit costs rather than capacities. This model was shown to admit an efficient approximate dynamic programming formulation, based on strong separability properties. However, this setting has several restrictions: the dynamic substitution effects are overlooked and the demand follows a Poisson process, while on the other hand, the retailer is allowed to utilize a mixed assortment strategy. Goyal et al. (2016), Segev (2015), and Aouad et al. (2015) considered dynamic assortment planning models, with a fully stochastic consumption process, for which they devised polynomial-time algorithms with provable approximation guarantees. However, the choice models considered in these papers have simple combinatorial structures, that impose a very specific order by which products are consumed and depleted. This property is crucial to the design of low-dimensional dynamic programs for revenue evaluation and optimization. In contrast, the choice outcomes described by the MNL model do not impose any particular (deterministic) pattern on stock-out events. Consequently, dynamic optimization in this context appears to be significantly more challenging.

2 Preliminaries

In what follows, we establish a number of technical results that were briefly discussed in Section 1.1. These are instrumental for our algorithmic approach and its analysis.

2.1 Extensions of submodular maximization

The crux of our algorithm resides in exploiting new notions of submodularity and monotonicity, respectively termed as *restricted submodularity* and *restricted monotonicity*. Intuitively, these properties require that the structural inequalities defining submodularity and monotonicity are satisfied as long as the sets involved are within the feasible region, formed by a cardinality constraint. Although weaker than the standard notions, we show that these properties are sufficient for the design of constant worst-case approximations, even with noisy estimates of the objective function.

Restricted submodularity and monotonicity. We begin by defining the notion of restricted submodularity. A set function $f : 2^{[n]} \to \mathbb{R}$ is said to be restricted-s-submodular for some integer $s \in \mathbb{N}$ if, for any subset $S \subseteq [n]$ of cardinality at most $s-2$ and elements $i \neq j \in [n] \setminus S$, we have

$$f(S \cup \{i,j\}) - f(S \cup \{j\}) \leq f(S \cup \{i\}) - f(S) \ .$$

By a similar extension of conventional definitions, we say that a set function f is restricted-s-non-decreasing if $f(S) \leq f(T)$ for any pair of subsets $S \subseteq T$ of cardinality at most s. In what follows, the parameter s is always equal to the capacity C, and therefore, we simply say that a set function is *restricted-submodular* or *restricted-non-decreasing*.

The efficient oracle assumption. A particularly useful extension of the standard submodular maximization setting is to assume that the objective function f cannot be evaluated exactly in an efficient way, and instead, we are given access to a noisy estimation oracle. Formally, the efficient oracle assumption states that, for any error parameter $\epsilon > 0$ and for any confidence level $\delta > 0$, there exists an efficient procedure that, given any subset $S \subseteq [n]$, computes a random estimate $\tilde{f}(S)$ of $f(S)$ such that

$$\Pr\left[(1-\epsilon) \cdot f(S) \leq \tilde{f}(S) \leq (1+\epsilon) \cdot f(S)\right] \geq 1 - \delta \ .$$

The running time of this procedure is assumed to be polynomial in the input size, $1/\epsilon$, and $1/\delta$.

By leveraging classic techniques for approximately maximizing monotone submodular functions (see, e.g., Nemhauser et al. (1978)), we derive a constant-factor approximation for non-negative restricted-non-decreasing and restricted-submodular functions, as stated in the following claim. To avoid deviating from our general theme, we present the proof in Appendix C.1.

Lemma 2.1. *Under the efficient oracle assumption, for any $\epsilon \in (0, 1/4)$ and $\delta > 0$, the problem of maximizing a non-negative restricted-non-decreasing and restricted-submodular set function under a cardinality constraint can be approximated within factor $0.318 - \epsilon$, with probability at least $1 - \delta$. The running time of our algorithm is polynomial in the input size, $n^{1/\epsilon}$, and $1/\delta$.*

2.2 Subadditivity of the expected revenue function

The next lemma, whose proof appears in Appendix C.2, asserts that the expected revenue function in the MNL-based dynamic assortment model is subadditive.

Lemma 2.2 (Subbaditivity). *For any inventory vectors U_1 and U_2, we have $\mathbb{E}[\mathcal{R}(U_1 + U_2)] \leq \mathbb{E}[\mathcal{R}(U_1)] + \mathbb{E}[\mathcal{R}(U_2)]$.*

To better understand the implications of this claim, let U^* be an optimal inventory vector. For any subset of products $S \subseteq [n]$, we use U^*_S to designate the projection of U^* on S, i.e., U^*_S is the vector obtained from U^* by setting the inventory levels of all products in $[n] \setminus S$ to zero. Now suppose that the collection of products $[n]$ is partitioned into the subsets $\mathcal{S}_1, \ldots, \mathcal{S}_K$. Consequently, since $U^* = \sum_{k \in [K]} U^*_{\mathcal{S}_k}$, and the expected revenue function is subadditive, it follows that

$$\sum_{k \in [K]} \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{S}_k})\right] \geq \mathbb{E}[\mathcal{R}(U^*)] \ . \qquad (1)$$

From an algorithmic perspective, this bound can be utilized by treating each subset $\mathcal{S}_k$ as a separate subproblem for which a tailor-made algorithm is developed. Now, suppose we obtain a γ_k-approximation for each subproblem, i.e., an inventory vector $\tilde{U}_{\mathcal{S}_k}$ satisfying $\mathbb{E}[\mathcal{R}(\tilde{U}_{\mathcal{S}_k})] \geq \gamma_k \cdot \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{S}_k})\right]$. By picking the best solution (revenue-wise) out of the K resulting inventory vectors, for any $\alpha_1, \ldots, \alpha_K \geq 1$ with $\sum_{k \in [K]} \alpha_k = 1$, we obtain an expected revenue of

$$\max_{k \in [K]} \mathbb{E}\left[\mathcal{R}(\tilde{U}_{\mathcal{S}_k})\right] \geq \sum_{k \in [K]} \alpha_k \cdot \mathbb{E}\left[\mathcal{R}(\tilde{U}_{\mathcal{S}_k})\right] \geq \left(\min_{k \in [K]} \alpha_k \gamma_k\right) \cdot \sum_{k \in [K]} \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{S}_k})\right] \geq \left(\min_{k \in [K]} \alpha_k \gamma_k\right) \cdot \mathbb{E}\left[\mathcal{R}(U^*)\right] \ ,$$

where the last inequality holds by (1). As a result, we have just obtained an approximation ratio of $\min_{k \in [K]} \alpha_k \gamma_k$ for the original problem, which can be optimized by picking the best convex combination $\alpha_1, \ldots, \alpha_K$. This decomposition idea is exploited in Section 3 and in Appendix A.

3 Core Algorithm with Evaluation Oracle

In this section, we devise an efficient algorithm with a constant-factor worst-case guarantee, under IFR demand distributions. Since revenue evaluation is challenging by itself in the dynamic setting, we temporarily operate under the efficient oracle assumption described in Section 2.1. Specifically, we assume in the remainder of this section that, for any error parameter $\epsilon > 0$ and confidence level $\delta > 0$, there is an efficient procedure to estimate the expected revenue $\mathbb{E}[\mathcal{R}(U)]$ of any inventory vector U up to a multiplicative factor of $1 \pm \epsilon$, with probability at least $1 - \delta$.

In Appendix A, we explain how this assumption can be bypassed, losing a small constant factor in optimality, while utilizing the conditional approach developed here as a subroutine. The latter bears practical significance by itself, since simulation-based methods or surrogate models are commonly used to go around the computational difficulties of evaluating certain objective functions.

Theorem 3.1. *Under the efficient oracle assumption, for any $\epsilon \in (0, 1/4)$ and $\delta > 0$, the dynamic assortment planning problem under the Multinomial Logit choice model and IFR demand distribution can be approximated within a factor of $0.139 - \epsilon$ with probability at least $1 - \delta$, in time polynomial in the input size, $n^{1/\epsilon}$, and $1/\delta$. When $C \geq n$, this factor can be improved to $0.179 - \epsilon$.*

3.1 Overview of the algorithm

Preliminary step: price threshold. We begin by computing $\text{OPT}_{\text{static}}$, the optimal capacitated static revenue. There are several well-known polynomial-time algorithms (Megiddo 1979, Rusmevichientong et al. 2010, Davis et al. 2013) to solve the capacitated static variant of the problem, where there is only one customer arriving, still with an upper bound of C on the number of products to be stocked (rather than units). Hereafter, the corresponding optimal static assortment (which generates an expected revenue of $\text{OPT}_{\text{static}}$) is denoted by $\mathcal{A}^*$. Subsequently, we use $\text{OPT}_{\text{static}}$ as a price threshold to distinguish between *expensive* products, with price greater or equal to $\text{OPT}_{\text{static}}$, and *cheap* products, whose price is smaller than $\text{OPT}_{\text{static}}$. We let $\mathcal{E}$ and $\mathcal{C}$ designate the subsets of expensive and cheap products, respectively, thus forming a partition of the products $[n]$. In the sequel, our algorithm constructs inventory vectors that are exclusively composed of expensive products, whereas the optimal inventory vector could stock both cheap and expensive products.

Decomposition approach. Next, we utilize the decomposition idea described in Section 2.2. Specifically, we pick the most profitable among two candidate inventory vectors, denoted by $U_{\mathcal{E}}$ and $U_{\mathcal{C}}$. While both stocking only expensive products, these vectors are constructed to fulfill different purposes: $U_{\mathcal{E}}$ competes against the contributions of expensive products in the optimal expected revenue, while $U_{\mathcal{C}}$ competes against the revenue contributions of cheap products. By 'compete', we mean that $U_{\mathcal{E}}$ is guaranteed to generate a constant fraction of the expected revenue due to selling expensive products in the optimal solution, and an analogous property holds for $U_{\mathcal{C}}$ with respect to the cheap products. From this point on, we let U^* be a fixed optimal inventory vector, and recall that $U^*_{\mathcal{E}}$ designates the projection of U^* on the set of expensive products $\mathcal{E}$, i.e., the vector obtained from U^* by setting the inventory levels of all cheap products to zero. The vector $U^*_{\mathcal{C}}$ is defined in an analogous way. Our analysis relies on comparing the expected revenue of $U_{\mathcal{E}}$ and $U_{\mathcal{C}}$ with that of $U^*_{\mathcal{E}}$ and $U^*_{\mathcal{C}}$, respectively.

Competing against expensive products (Section 3.2). Since our analysis of the greedy algorithm for restricted-non-decreasing and restricted-submodular functions (see Section 2.1) results in an extra additive error that depends on $1/C$, we distinguish between two cases in order to construct $U_{\mathcal{E}}$. Specifically, when $C \geq 1/\epsilon$, the inventory levels of expensive products are determined by a greedy approach, where at each step a single unit of the product that generates the largest marginal increase in the expected revenue is picked until stocking C units. In this case, the above-mentioned additive error affects the multiplicative factor we obtain by a factor of only $O(\epsilon)$. In the opposite case, when $C < 1/\epsilon$, we resort to enumeration over all $O(n^{1/\epsilon})$ feasible inventory vectors. To analyze this approach, we prove that the restricted-non-decreasing and restricted-submodular properties are satisfied by the revenue function (in modified form), for the problem restricted to the collection of expensive products $\mathcal{E}$. By executing the evaluation oracle with the appropriate error and confidence parameters described in Lemma 2.1, it follows that the inventory vector $U_{\mathcal{E}}$ competes against the optimal expected revenue obtained from expensive products. Specifically, with probability at least $1-\delta$,

$$\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{E}}\right)\right] \geq (0.318-\epsilon) \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{E}}\right)\right] \ . \tag{2}$$

When $C \geq n$, we obtain an improved approximation ratio of $(1-\epsilon)\cdot(1-1/e)$.

Competing against cheap products (Section 3.3). We compete against $U^*_{\mathcal{C}}$ by stocking expensive products, rather than cheap products. Specifically, $U_{\mathcal{C}}$ is computed through a greedy procedure, where stocking decisions are restricted to the optimal static assortment $\mathcal{A}^*$, and the expected revenue function is replaced by a simplified objective function. This alternative objective is formed by neglecting the revenue generated by stock-out substitution, namely, assuming that customers do not substitute to less preferred options once their most preferred product is depleted. The lower bound thus obtained can be interpreted as the objective function of a multi-item newsvendor problem, that can be optimized greedily. By exploiting the IFR property, we show that $U_{\mathcal{C}}$ guarantees at least $1/4$ of the optimal expected revenue due to cheap products, i.e.,

$$\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{C}}\right)\right] \geq \frac{1}{4} \cdot \mathbb{E}[\mathcal{R}\left(U^*_{\mathcal{C}}\right)] \ . \tag{3}$$

Concluding the proof of Theorem 3.1. Before providing additional details on the above-mentioned algorithms and their respective performance, we argue that inequalities (2) and (3) are sufficient to prove the worst-case guarantee stated in Theorem 3.1, using the decomposition ideas of Section 2.2. Recall that, since the expected revenue function is subadditive, we have

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{E}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{C}}\right)\right] \geq \mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \ . \tag{4}$$

Now, for any $\alpha \in [0,1]$, picking the better vector out of $U_{\mathcal{E}}$ and $U_{\mathcal{C}}$ guarantees, with probability at least $1-\delta$, an expected revenue of

$$\begin{aligned}\max\left\{\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{E}}\right)\right], \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{C}}\right)\right]\right\} &\geq \alpha \cdot \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{E}}\right)\right] + (1-\alpha) \cdot \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{C}}\right)\right] \\ &\geq \alpha \cdot (0.318-\epsilon) \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{E}}\right)\right] + \frac{1-\alpha}{4} \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{C}}\right)\right] \\ &\geq (1-4\epsilon) \cdot \left(0.318 \cdot \alpha \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{E}}\right)\right] + \frac{1-\alpha}{4} \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{C}}\right)\right]\right) ,\end{aligned}$$

where the second inequality is an immediate consequence of (2) and (3). Thus, by choosing $\alpha = 0.25/0.568$,

$$\begin{aligned}\max\left\{\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{E}}\right)\right], \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{C}}\right)\right]\right\} &\geq (1-4\epsilon) \cdot 0.318 \cdot \alpha \cdot \left(\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{E}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{C}}\right)\right]\right) \\ &\geq (1-4\epsilon) \cdot 0.139 \cdot \mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] ,\end{aligned}$$

where the last inequality follows from the upper bound (4). As previously mentioned, when $C \geq n$, the inventory vector $U_{\mathcal{E}}$ actually satisfies $\mathbb{E}[\mathcal{R}(U_{\mathcal{E}})] \geq (1-\epsilon)(1-1/e) \cdot \mathbb{E}[\mathcal{R}(U^*_{\mathcal{E}})]$. By plugging-in this inequality instead of (2), and by picking $\alpha = e/(5e-4)$, we derive an improved constant-factor guarantee of $0.179-\epsilon$.

3.2 Competing against expensive products

In this section, we consider the *expensive-products problem*, that is, a modified instance only comprised of products in $\mathcal{E}$. We show that the restricted-submodular and restricted-non-decreasing properties are satisfied by the expected revenue function (when reformulated appropriately), although their standard counterparts do not hold in this context. The proof mainly relies on probabilistic coupling ideas, that allow us to compare the consumption process under different initial inventory level decisions. As a result, Lemma 2.1 entails the following theorem.

Theorem 3.2. *Under the efficient oracle assumption, for any* $\epsilon \in (0, 1/4)$ *and* $\delta > 0$*, the expensive-products problem can be approximated within factor* $0.318-\epsilon$ *with probability at least* $1-\delta$*. The running time of our algorithm is polynomial in the input size,* $n^{1/\epsilon}$*, and* $1/\delta$*.*

Set decision formulation. In order to establish the desired submodularity-like properties, the problem needs to be interpreted as the maximization of a set function under a cardinality constraint. To this end, each product is duplicated into C copies, each representing a distinct unit of that product. In the expensive-products problem, this transformation results in an extended set of $N = C \cdot |\mathcal{E}|$ distinct units. With this notation, the objective is to decide on a subset of the extended collection of units $S \subseteq [N]$, as a substitute to the inventory vector U.

Once an initial offer set $S \subseteq [N]$ is chosen, each arriving customer purchases one unit of her most preferred product available, according to the MNL choice model. Since units of the same product are identical, the realizations of the revenue random variable are invariant to the precise unit being purchased, which can thus be chosen arbitrarily (in the sequel, we often impose that a specific unit is purchased, for purposes of analysis). Finally, we define the objective function $f_Y(S)$ to be the expected revenue when initially stocking the subset of units S, where Y stands for the random number of arriving customers. Consequently, the original expensive-products problem translates to maximizing $f_M(S)$ over all subsets $S \subseteq [N]$ of cardinality at most C.

Simplified notation. In what follows, we allow mixed notation between products and their respective units. Specifically, w_i and r_i designate the preference weight and the selling price of the product corresponding to unit i. Unless specified otherwise, when the subset of units $S \subseteq [N]$ is fixed, the corresponding assortment of products is designated by $\mathcal{A} \subseteq \mathcal{E}$. We use the shorthand notation $\mathcal{A}^{+i}$ to denote the resulting assortment when a unit $i \in [N]$ is added to S, and $\mathcal{A}^{+ij}$ when two units $i, j \in [N]$ are added.

3.2.1 Probabilistic coupling

To establish the restricted-submodular and restricted-non-decreasing properties, we would have to compare the expected revenue of different subsets. For example, we wish to prove that, for any subset $S \subseteq [N]$ of cardinality at most $C - 2$, and any units $i \neq j \in [N] \setminus S$, we have

$$f_M\left(S \cup \{i, j\}\right) - f_M\left(S \cup \{j\}\right) \leq f_M\left(S \cup \{i\}\right) - f_M\left(S\right) \ .$$

To derive such inequalities, we implicitly need to compare the consumption process for the initial subsets S, $S \cup \{i\}$, $S \cup \{j\}$, and $S \cup \{i, j\}$. To this end, our coupling construction will introduce useful relationships between the probabilistic outcomes generated by these subsets. By design, in the construction below, units i and j correspond to two distinct products, both not stocked in S. The construction remains identical even in other settings, where units i and j are arbitrary.

Purchase random variables. We focus on the first arriving customer, and introduce several random variables to describe her purchase decision, when facing each of the above-mentioned subsets. Specifically, denoting the no-purchase option by product 0, with preference weight $w_0 = 1$ and selling price $r_0 = 0$, we define:

- P as the product purchased when the offered set is S, i.e., within the initial assortment $\mathcal{A} \cup \{0\}$.
- P_i as the product purchased when the offered set is $S \cup \{i\}$, i.e., within the initial assortment $\mathcal{A}^{+i} \cup \{0\}$.
- P_j as the product purchased when the offered set is $S \cup \{j\}$, i.e., within the initial assortment $\mathcal{A}^{+j} \cup \{0\}$.
- $P_{i,j}$ as the product purchased when the offered set is $S \cup \{i, j\}$, i.e., within the initial assortment $\mathcal{A}^{+ij} \cup \{0\}$.

Coupling construction. Rather than defining these random variables separately, in independent probabilistic spaces, we artificially correlate their random outcomes for purposes of analysis, while still preserving their MNL-based marginal probabilities. In other words, denoting by $Y \sim Z$ the equivalence in distribution between two random variables Y and Z, we

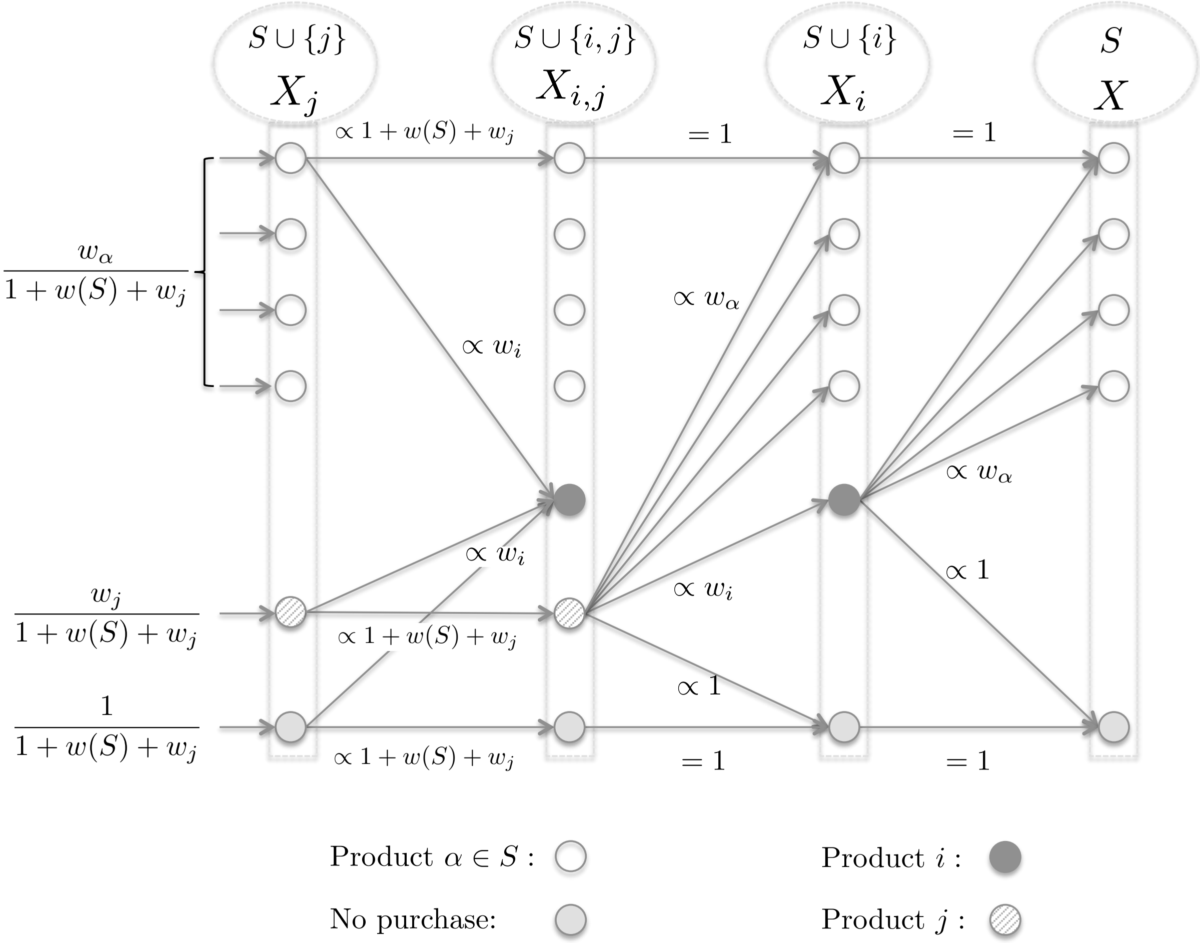

Figure 1: Markov chain representation of the coupling between the random variables X_j,$X_{i,j}$,X_i, and X. Here, random purchase events (or states) are represented by nodes, and each arc corresponds to a transition with positive probability. These transition probabilities are specified either exactly or in proportions (e.g., if a node has two outgoing arcs, one with $\propto 3$ and the other with $\propto 5$, the transition probabilities are 3/8 and 5/8, respectively).

construct a multivariate distribution for $(X_j, X_{i,j}, X_i, X)$, where $X_j \sim P_j$, $X_{i,j} \sim P_{i,j}$, $X_i \sim P_i$, and $X \sim P$. Our coupling approach relies on stipulating that the sequence $X_j, X_{i,j}, X_i, X$ forms a Markov chain, i.e., $X_i|(X_{i,j}, X_j) = X_i|X_{i,j}$ and $X|(X_i, X_{i,j}, X_j) = X|X_i$, whose transition probabilities are specified below, through the conditional random variables X_j, $X_{i,j}|X_j$, $X_i|X_{i,j}$, and $X|X_i$.

To illustrate the upcoming definitions, we provide in Figure 1 a schematic representation of the underlying transition graph, that can be used to derive useful probabilistic claims regarding the purchase random variables. For example, there is a single incoming edge to each white node of $X_{i,j}$, representing the purchase of a product $\alpha \in \mathcal{A}$. Given that this edge is horizontal, it describes the same product option for the variable X_j, and it follows that $\Pr[X_j = \alpha | X_{i,j} = \alpha] = 1$.

- *Defining X_j.* Here, we simply use the marginal probabilities prescribed by the MNL choice model for the purchases made by the first arriving customer under the initial assortment

$\mathcal{A}^{+j}$. That is, for any product $\alpha \in \mathcal{A}^{+j} \cup \{0\}$,

$$\Pr\left[X_j = \alpha\right] = \frac{w_\alpha}{1 + w(\mathcal{A}) + w_j} \ .$$

- *Defining* $X_{i,j}|X_j$. The initial set $S \cup \{i,j\}$ contains one more purchase option than $S \cup \{j\}$, namely product i. Intuitively, the event $\{X_{i,j} = i\}$ is defined by 'rescaling' uniformly the purchase probabilities of all other products in $\mathcal{A}^{+j} \cup \{0\}$, which are captured by the variable X_j. Formally, for any product $\alpha \in \mathcal{A}^{+ij} \cup \{0\}$, we define:

$$\Pr\left[X_{i,j} = i | X_j = \beta\right] = \frac{w_i}{1 + w(\mathcal{A}) + w_j + w_i} \text{ for } \beta \in \mathcal{A}^{+j} \cup \{0\} \ , \tag{5}$$

$$\Pr\left[X_{i,j} = \alpha | X_j = \alpha\right] = \frac{1 + w(\mathcal{A}) + w_j}{1 + w(\mathcal{A}) + w_j + w_i} \text{ if } \alpha \neq i \ , \tag{6}$$

$$\Pr\left[X_{i,j} = \alpha | X_j = \beta\right] = 0 \text{ if } \alpha \neq i \text{ and } \beta \neq \alpha \ . \tag{7}$$

- *Defining* $X_i|X_{i,j}$. We relate the purchases made in the assortment $\mathcal{A}^{+i} \cup \{0\}$ with the purchases made in $\mathcal{A}^{+ij} \cup \{0\}$. In contrast to the previous case, we now need to 'eliminate' the purchase option relative to product j. This is done by 'reallocating' the probability of the event $\{X_{i,j} = j\}$ to the purchases of other products, proportionally to their MNL weights. That is, for any product $\alpha \in \mathcal{A}^{+i} \cup \{0\}$, we define:

$$\Pr\left[X_i = \alpha | X_{i,j} = \alpha\right] = 1 \ , \tag{8}$$

$$\Pr\left[X_i = \alpha | X_{i,j} = j\right] = \frac{w_\alpha}{1 + w(\mathcal{A}) + w_i} \tag{9}$$

$$\Pr\left[X_i = \alpha | X_{i,j} = \beta\right] = 0 \text{ for } \beta \neq \alpha, j \ . \tag{10}$$

- *Defining* $X|X_i$. Our construction is similar to the previous case, and for any product $\alpha \in \mathcal{A} \cup \{0\}$, we define:

$$\Pr\left[X = \alpha | X_i = \alpha\right] = 1 \ , \tag{11}$$

$$\Pr\left[X = \alpha | X_i = i\right] = \frac{w_\alpha}{1 + w(\mathcal{A})} \ . \tag{12}$$

$$\Pr\left[X = \alpha | X_i = \beta\right] = 0 \text{ for } \beta \neq \alpha, i \ . \tag{13}$$

The next lemma, whose proof is given in Appendix C.3, states that this coupling method indeed preserves the (marginal) MNL purchase probabilities for each initial offer set.

Claim 3.3. *$X_j \sim P_j$, $X_{i,j} \sim P_{i,j}$, $X_i \sim P_i$, and $X \sim P$.*

In addition, we establish several equivalence properties that will prove useful for the analysis, stating that the purchase random variables X, X_i, and X_j are invariant in distribution when conditioned on appropriate events of $X_{i,j}$. The proof of the next lemma appears in Appendix C.4.

Claim 3.4. *$(X_j|X_{i,j} = i) \sim X_j$, $(X_i|X_{i,j} = j) \sim X_i$, $(X|X_{i,j} = i) \sim X$, and $(X|X_{i,j} = j) \sim X$.*

3.2.2 Proving restricted monotonicity

We prove that the revenue function f_M is restricted-non-decreasing by an inductive argument. To better understand which sufficient properties come into play, the proof is broken down into two lemmas: we first examine the case of a single arriving customer, before extending our arguments to any random variable M. It is worth mentioning that this property holds regardless of how M is distributed, whether IFR or not.

Lemma 3.5. *In the expensive-products setting, the static expected revenue function f_1 is restricted-non-decreasing.*

Proof. It is easy to verify that the restricted-non-decreasing property is equivalent to having $f_1(S \cup \{i\}) \geq f_1(S)$ for any subset S of size at most $C-1$ and any unit i. Observe that, if the product corresponding to i is stocked by S, we clearly have $f_1(S \cup \{i\}) = f_1(S)$. When this product is not stocked in S,

$$\begin{aligned} f_1(S \cup \{i\}) - f_1(S) &= \frac{r_i w_i}{1 + w(\mathcal{A}) + w_i} + \sum_{k \in S} r_k w_k \cdot \left(\frac{1}{1 + w(\mathcal{A}) + w_i} - \frac{1}{1 + w(\mathcal{A})} \right) \\ &= \frac{w_i}{1 + w(\mathcal{A}) + w_i} \cdot \left(r_i - \sum_{k \in S} \frac{r_k w_k}{1 + w(\mathcal{A})} \right) \\ &= \frac{w_i}{1 + w(\mathcal{A}) + w_i} \cdot (r_i - f_1(S)) \ . \end{aligned} \tag{14}$$

This proves the desired inequality, since $r_i \geq \text{OPT}_{\text{static}} \geq f_1(S)$, where the former inequality follows from i being an expensive product, and the latter holds since the assortment $\mathcal{A}$ stocked by S has fewer than C products (all expensive), implying that its static expected revenue is at most $\text{OPT}_{\text{static}}$, which stands for the maximum possible static revenue when we are allowed to stock at most C products (expensive and cheap). ■

Lemma 3.6. *For any instance of the dynamic assortment planning problem where the static revenue function f_1 is restricted-non-decreasing, the revenue function f_M is restricted-non-decreasing as well, for any demand random variable M.*

Proof. The preliminary observation is that, by the formula of conditional expectation, it is sufficient to prove the desired property for a deterministic value of M. Also, one can easily verify that it is sufficient to prove $f_M(S_1) \leq f_M(T_1)$ for any two initial offer sets $S_1 \subseteq T_1 \subseteq [N]$ with cardinality at most C, that differ by at most one unit, i.e., $|T_1 \setminus S_1| \leq 1$.

To this end, we leverage our coupling method for the purchase random variables, constructed in Section 3.2.1, to derive a coupling of the consumption processes under the initial subsets S_1 and T_1. Let $S_1, \ldots, S_M$ and $T_1, \ldots, T_M$ be the (random) residual subsets of inventory units facing customers $1, \ldots, M$, when respectively stocking the initial subsets S_1 and T_1. We denote by $\mathcal{A}_1, \ldots, \mathcal{A}_M$ and $\mathcal{B}_1, \ldots, \mathcal{B}_M$ the corresponding sequences of assortments.

We wish to define a coupling of these random variables such that $S_k \subseteq T_k$ and $|T_k \setminus S_k| \leq 1$, at each arrival $k \in [M]$. This coupling is constructed inductively over the sequence of arrivals, by refining at each step our probabilistic space with respect to the next arriving customer. Since the desired properties are clearly satisfied for the base case $k = 1$ by definition of S_1 and T_1, suppose the inductive hypothesis holds until the k-th arrival. If $S_k = T_k$, it is easy to see that the inductive property propagates to the next arrivals, since the purchases made in the two consumption processes are identical. Otherwise, let i be the (single) unit contained in $T_k \setminus S_k$. For the next arriving customer, we distinguish between two cases:

- *Product i is contained in $\mathcal{A}_k$.* As a result, the k-th arriving customer is facing the same assortment under both offer sets S_k and T_k, i.e., $\mathcal{A}_k = \mathcal{B}_k$. Here, the purchase random variable X (with $S = S_k$) defines a trivial coupling of the purchases made by the first arriving customer in both cases, in the sense that the purchases are identical and described by the outcomes of X with respect to $\mathcal{A}_k$. Consequently, the random residual sets facing the next arriving customer satisfy $S_{k+1} \subseteq T_{k+1}$, since $S_k \subseteq T_k$ and the same unit of product X can be purchased in both cases by the k-th arriving customer. In addition, we have $|T_{k+1} \setminus S_{k+1}| = |T_k \setminus S_k| \leq 1$.

- *Product i is not contained in $\mathcal{A}_k$.* In this case, taking $S = S_k$ and $\mathcal{A} = \mathcal{A}_k$, we use the coupling of X and X_i as a joint distribution for the purchases made by the first arriving customer under the sets S_k and T_k, respectively. By definition of $X|X_i$, observe that if the customer faced with T_k purchases a product in $\mathcal{A} \cup \{0\}$, i.e., $X_i \in \mathcal{A} \cup \{0\}$, then the customer faced with S_k purchases the same product, i.e., $X = X_i$ (in Figure 1, there is a single horizontal edge going into each white node of X, describing the same purchase option in X and X_i). Indeed, our coupling entails that $\Pr[X = X_i | X_i \in \mathcal{A} \cup \{0\}] = 1$ due to equation (11). As a result, since the k-th customer purchases the same unit in both cases, the inductive hypothesis implies that $S_{k+1} \subseteq T_{k+1}$. On the other hand, the event $\{X_i = i\}$ means that the customer faced with T_k purchases the last unit of product i, and thus, conditional to this event, the remaining set of units is necessarily $T_{k+1} = T_k \setminus \{i\} = S_k$, which clearly leads to $S_{k+1} \subseteq T_{k+1}$. In both cases, we have preserved the invariant $|T_{k+1} \setminus S_{k+1}| \leq 1$.

We have just obtained a coupling of the consumption processes such that $S_k \subseteq T_k$ and $|T_k \setminus S_k| \leq 1$ for every $k \in [M]$. By exploiting this inclusion property between subsets of units, we now prove that $f_M(S_1) \leq f_M(T_1)$. To this end, a natural transformation of the expected revenue function is

$$f_M(S_1) = \sum_{k=1}^{M} \mathbb{E}\left[f_1(S_k)\right] \ , \tag{15}$$

where the overall expected revenue breaks down into the sum of expected revenues generated by customers $1, \ldots, M$, faced by the random residual sets of units $S_1, \ldots, S_M$, respectively. Using a similar transformation for T_1, we have

$$f_M\left(T_1\right) - f_M\left(S_1\right) = \sum_{k=1}^{M} \mathbb{E}\left[f_1(T_k) - f_1(S_k)\right] \ .$$

Therefore, since f_1 is assumed to be restricted-non-decreasing, and $S_k \subseteq T_k$, the latter expression is non-negative, meaning that the restricted-non-decreasing property extends to f_M. ■

3.2.3 Proving restricted submodularity

We now show that the transformed revenue function f_M is also restricted-submodular, regardless of how M is distributed, by exploiting the coupling method described in Section 3.2.1. We first examine the static case, before extending the desired property to any demand random variable.

Lemma 3.7. *For any instance of the dynamic assortment planning problem where the static expected revenue function f_1 is restricted-non-decreasing, this function is restricted-submodular.*

Proof. Let $S \subseteq [N]$ be a set with $|S| \leq C - 2$, $\mathcal{A}$ is the assortment stocked by S and let $i \neq j$ be two units in $[N] \setminus S$. In order to prove that $f_1(S \cup \{i, j\}) - f_1(S \cup \{j\}) \leq f_1(S \cup \{i\}) - f_1(S)$, we distinguish between four cases:

1. *Product i is contained in $\mathcal{A}$.* Adding unit i to any subset of units containing S leaves us with the same assortment $\mathcal{A}$, meaning that $f_1(S \cup \{i, j\}) - f_1(S \cup \{j\}) = f_1(S \cup \{i\}) - f_1(S) = 0$.

2. *Product j is contained in $\mathcal{A}$, product i is not.* In this case, adding unit j to any subset containing S leaves us with the same assortment, thus $f_1(S \cup \{i, j\}) - f_1(S \cup \{i\}) = f_1(S \cup \{j\}) - f_1(S)$.

3. *Units i and j are of the same product, not contained in $\mathcal{A}$.* Here, we observe that $f_1(S \cup \{i, j\}) - f_1(S \cup \{j\}) = 0$ while $f_1(S \cup \{i\}) - f_1(S) \geq 0$ since f_1 is restricted-non-decreasing.

4. *Products i and j are different, and both not contained in $\mathcal{A}$.* By calculations similar to those leading to equation (14), we get

$$\begin{aligned} f_1(S \cup \{i, j\}) - f_1(S \cup \{j\}) &= \frac{w_i}{1 + w(\mathcal{A}) + w_i + w_j} \cdot (r_i - f_1(S \cup \{j\})) \\ &\leq \frac{w_i}{1 + w(\mathcal{A}) + w_i} \cdot (r_i - f_1(S)) \\ &= f_1(S \cup \{i\}) - f_1(S) \ , \end{aligned}$$

   where the inequality above holds since f_1 is restricted-non-decreasing, thus $f_1(S \cup \{j\}) \geq f_1(S)$.

■

Lemma 3.8. *For any instance of the dynamic assortment planning problem where the static expected revenue function f_1 is restricted-non-decreasing, the revenue function f_M is restricted-submodular as well, for any demand random variable M.*

Proof. By the formula of conditional expectations, we restrict attention to deterministic values of M without loss of generality, and prove the claim by induction on M. Since the case $M = 1$ corresponds to Lemma 3.7, suppose that restricted submodularity has been established for $M - 1$ arrivals. We show that, for any subset $S \subseteq [N]$ of cardinality at most $C - 2$, and units $i \neq j \in [N] \setminus S$,

$$f_M(S \cup \{i, j\}) - f_M(S \cup \{j\}) \leq f_M(S \cup \{i\}) - f_M(S) \ . \tag{16}$$

The proof consists of the same case analysis made for the proof of Lemma 3.7. In what follows, we only discuss the most difficult case, where products i and j are different and not contained in the assortment $\mathcal{A}$ stocked by S. The additional cases have nearly-identical proofs and are not presented here to avoid redundancy.

A particularly instructive observation is that, for any subset of units $T \subseteq [N]$, the expected revenue function decomposes into the contribution due to the purchase made by the first arriving customer, and that associated with the residual subset of units and the remaining customers arrivals. Formally, letting $R_M(T)$ denote the random revenue obtained after M arriving customer facing an initial subset T, and using Y to designate the product purchased by the first arriving customer, we have

$$f_M(T) = \mathbb{E}\left[R_M(T)\right] = \mathbb{E}\left[r_Y\right] + \mathbb{E}\left[R_{M-1}\left(T \setminus \{Y\}\right)\right] \ . \tag{17}$$

In what follows, for ease of notation, we denote $S^{+i} = S \cup \{i\}$, $S^{+j} = S \cup \{j\}$, and $S^{+ij} = S \cup \{i,j\}$. In order to derive inequality (16), we make use of the revenue decomposition (17) under different initial subsets S, S^{+i}, S^{+j}, and S^{+ij}, along with their corresponding purchase random variables X, X_i, X_j, and $X_{i,j}$.

By conditioning the revenue random variable with respect to $X_{i,j}$, we leverage our coupling method of Section 3.2.1 to explicitly compare the random purchases made by the first arriving customer under different initial subsets, and establish the next claim. Finally, the desired inequality (16) comes at an immediate consequence of this claim, combined with the formula of conditional expectations.

Claim 3.9. $\mathbb{E}[R_M(S^{+ij}) - R_M(S^{+j})|X_{i,j} = \alpha] \leq \mathbb{E}[R_M(S^{+i}) - R_M(S)|X_{i,j} = \alpha]$

Proof of Claim 3.9. We present here the case where $\alpha \in \mathcal{A} \cup \{0\}$, the other cases being treated through similar arguments in Appendix C.5. We begin by observing that, conditional on the event $\{X_{i,j} = \alpha\}$ where $\alpha \in \mathcal{A} \cup \{0\}$, we necessarily have $X_j = X_{i,j} = X_i = X = \alpha$. Indeed, in the transition graph of the Markov chain (Figure 1), observe that there is a single horizontal path going through each white node of $X_{i,j}$, that describes the same purchase option across all variables. Formally, by equation (8), observe that $\Pr\left[X_i = \alpha | X_{i,j} = \alpha\right] = 1$, while equation (11) implies that $\Pr\left[X = \alpha | X_i = \alpha\right] = 1$. Finally, $\Pr\left[X_j = \alpha | X_{i,j} = \alpha\right] = 1$ follows from Bayes rule, using equation (6) along with the marginal distributions of X_j and $X_{i,j}$, which are described by the MNL model (see Claim 3.3). Thus, we obtain:

$$\begin{aligned}
&\mathbb{E}\left[R_M\left(S^{+ij}\right) - R_M\left(S^{+j}\right)\middle| X_{i,j} = \alpha\right] \\
&\qquad = \mathbb{E}\left[r_{(X_{i,j}|X_{i,j}=\alpha)} + R_{M-1}\left(S^{+ij} \setminus \{\alpha\}\right)\right] - \mathbb{E}\left[r_{(X_j|X_{i,j}=\alpha)} + R_{M-1}\left(S^{+j} \setminus \{\alpha\}\right)\right] \\
&\qquad = \mathbb{E}\left[R_{M-1}\left(S^{+ij} \setminus \{\alpha\}\right) - R_{M-1}\left(S^{+j} \setminus \{\alpha\}\right)\right] \\
&\qquad \leq \mathbb{E}\left[R_{M-1}\left(S^{+i} \setminus \{\alpha\}\right) - R_{M-1}\left(S \setminus \{\alpha\}\right)\right] \\
&\qquad = \mathbb{E}\left[R_M\left(S^{+i}\right) - R_M\left(S\right)\middle| X_{i,j} = \alpha\right] \ ,
\end{aligned} \tag{18}$$

where the first equality proceeds from equation (17), the second equality holds since the terms $r_{(X_{i,j}|X_{i,j}=\alpha)} = r_{(X_j|X_{i,j}=\alpha)} = r_\alpha$ cancel out, the next inequality is due to the inductive hypothesis (16), and the last equality is analogous to the first two equalities (in reverse order). ■

■

3.2.4 Improved performance guarantees for special settings

A close examination of the statements of Lemmas 3.6, 3.7, and 3.8 reveals that, for any integer $s \in [N]$, the static expected revenue function f_1 being restricted-s-non-decreasing is a sufficient condition for f_M to be both restricted-s-non-decreasing and restricted-s-submodular. Hence, when f_1 is non-decreasing in the standard sense, it follows that the function f_M is non-decreasing and submodular. In such cases, the standard analysis of the greedy algorithm for monotone submodular maximization (Nemhauser et al. 1978), with appropriately-chosen error and confidence parameters for the evaluation oracle, provides an improved performance guarantee of $(1 - \epsilon) \cdot (1 - 1/e)$.

It is worth mentioning that the expensive-products problem naturally satisfies this condition when $C \geq n$. Indeed, by Lemma 3.5, we know that f_1 is restricted-n-non-decreasing. Now, consider a transformation that maps each set $S \subseteq [N]$ to the subset $\tilde{S} \subseteq S$ obtained by keeping the lowest-index unit of each product stocked by S. Clearly, this transformation preserves the static expected revenue, i.e., $f_1(S) = f_1(\tilde{S})$, while ensuring that $\tilde{S}$ has at most n units. Consequently, for any subsets $S \subseteq T \subseteq [N]$, we infer that $f_1(S) = f_1(\tilde{S}) \leq f_1(\tilde{T}) = f_1(T)$ since f_1 is restricted-n-non-decreasing and $\tilde{S} \subseteq \tilde{T}$.

This condition is also satisfied when the assortment of products has been chosen in advance by solving the static assortment planning problem under the MNL model, and it remains to set their inventory levels. It is easy to verify that f_1 is non-decreasing when the entire collection of products forms an optimal static assortment.

3.3 Competing against cheap products

In this section, we construct an inventory vector $U_{\mathcal{C}}$ that guarantees a constant fraction of the expected revenue of $U^*_{\mathcal{C}}$, which stands for the projection of the optimal inventory U^* on the cheap products. We begin by presenting the algorithm before stating the performance guarantee obtained.

3.3.1 Algorithm

Step 1: Computing an optimal static assortment. As explained in Section 3.1, we begin by optimally solving the static assortment planning problem, subject to a capacity constraint of C on the number of products offered (Megiddo 1979, Rusmevichientong et al. 2010, Davis et al. 2013). Recall that the corresponding optimal static assortment, that generates an expected revenue of $\mathrm{OPT}_{\mathrm{static}}$, is denoted by $\mathcal{A}^*$. We now highlight a basic property of optimal static assortments, claiming that only expensive products are being stocked.

Claim 3.10. $\mathcal{A}^* \subseteq \mathcal{E}$.

This property follows from existing work on the static capacitated assortment planning problem under the MNL model (see Proposition 1 of Talluri and Van Ryzin (2004) and endnote 2 of Rusmevichientong et al. (2010)). For completeness, we provide a short proof in Appendix C.6.

Step 2: Deriving a newsvendor-like lower bound. To derive a lower bound, we will neglect the revenue generated by stockout-based substitution, and consider a setting where customers purchase their most preferred product until it stocks out. Specifically, suppose that U is an initial inventory vector stocking only units in the assortment $\mathcal{A}^*$. Then, for any product $i \in \mathcal{A}^*$, the probability that an arriving customer purchases that product is at least $\psi_i = w_i/(1 + w(\mathcal{A}^*))$, regardless of the inventory levels of the other products, as long as product i has not stocked out. Indeed, as the inventory vector is depleted due to previously-arriving customers, this can only increase the probability of each remaining product to be consumed by the next customer. To better understand the latter claim, consider two assortments $\tilde{\mathcal{A}} \subseteq \mathcal{A}$. The probability that an arriving customer purchases a unit of product $i \in \mathcal{A}$, when faced with assortment $\mathcal{A}$, is $w_i/(1+w(\mathcal{A}))$. By inclusion, this quantity is smaller or equal to $w_i/(1+w(\tilde{\mathcal{A}}))$, namely the probability of picking i among the assortment $\tilde{\mathcal{A}}$.

Consequently, the number of units purchased from i if this product had an unlimited number of units is stochastically larger than the binomial random variable $Y_i \sim B(M, \psi_i)$. However, since product i has only u_i units, we will actually be considering the truncated random variable $\bar{Y}_i(u_i) = \min\{Y_i, u_i\}$. Therefore, we obtain the following lower bound:

$$\mathbb{E}\left[\mathcal{R}\left(U\right)\right] \geq \sum_{i \in \mathcal{A}^*} r_i \cdot \mathbb{E}\left[\bar{Y}_i(u_i)\right] \ . \tag{19}$$

This lower bound can be interpreted as the objective function of a multi-item newsvendor problem, where the demand is separable across the products of $\mathcal{A}^*$. In what follows, this function is denoted by $\mathcal{L}(U) = \sum_{i \in \mathcal{A}^*} r_i \cdot \mathbb{E}\left[\bar{Y}_i(u_i)\right]$.

Step 3: Computing $U_{\mathcal{C}}$ by greedily optimizing the lower bound. Finally, the inventory vector $U_{\mathcal{C}}$, used to compete against cheap products, is constructed by solving the multi-item newsvendor instance defined above. That is, we compute U that maximizes $\mathcal{L}(U)$, subject to $\sum_{i \in \mathcal{A}^*} u_i \leq C$. This optimization problem can be solved exactly by a standard greedy procedure (see Muckstadt and Sapra (2010, Chap. 5)). Namely, starting from an empty inventory vector, units are added iteratively until reaching the capacity C, by picking at each step the unit with largest marginal contribution to the objective function. Therefore, in contrast to the original expected revenue, for any inventory vector U the lower bound $\mathcal{L}(U)$ can be computed in polynomial time.

3.3.2 Performance guarantee

The remainder of this section is devoted to proving the next theorem, showing that $U_{\mathcal{C}}$ competes against the expected revenue of cheap products $U^*_{\mathcal{C}}$.

Theorem 3.11. $\mathbb{E}[\mathcal{R}(U_{\mathcal{C}})] \geq (1/4) \cdot \mathbb{E}[\mathcal{R}(U^*_{\mathcal{C}})]$.

Our analysis proceeds by comparing an upper bound on $\mathbb{E}[\mathcal{R}(U^*_{\mathcal{C}})]$ with a lower bound on $\mathbb{E}[\mathcal{R}(U_{\mathcal{C}})]$, using the IFR property. It bears some resemblance to the analysis of Aouad et al. (2015, Thm. 1), combined with additional structural properties of the MNL choice model.

Upper bound on the expected revenue of $U^*_{\mathcal{C}}$. The important observation is that, when initially stocking at most C units of cheap products, each arriving customer will generate an expected revenue of at most $\text{OPT}_{\text{static}}$. As a result, the expected revenue of $U^*_{\mathcal{C}}$ is upper bounded by $\mathbb{E}[M]\cdot\text{OPT}_{\text{static}}$. In addition, since all cheap products have by definition selling prices smaller than $\text{OPT}_{\text{static}}$, another upper bound on the expected revenue of $U^*_{\mathcal{C}}$ is $C\cdot\text{OPT}_{\text{static}}$. Therefore,

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{C}}\right)\right] \leq \text{OPT}_{\text{static}} \cdot \min\left\{C, \mathbb{E}\left[M\right]\right\} \ . \tag{20}$$

Lower bound on the expected revenue of $U_{\mathcal{C}}$. To define our lower bound, we begin by introducing an inventory vector $U^{\propto}$, where the inventory levels are scaled proportionally to their revenue contribution toward $\text{OPT}_{\text{static}}$. Ideally, for each product $i \in \mathcal{A}^*$, the inventory level of i represents a fraction of $r_i\psi_i/\text{OPT}_{\text{static}}$ of the total capacity C (recalling that $\psi_i = w_i/(1 + w(\mathcal{A}^*))$). Hence, we would have liked to define the vector $\tilde{U}$, where $\tilde{u}_i = (r_i\psi_i/\text{OPT}_{\text{static}})\cdot C$. However, this quantity may not be integral, and is therefore rounded up to the nearest integer, creating the vector $U^{\propto}$. That is, $u^{\propto}_i = \lceil \tilde{u}_i \rceil$ for every product $i \in \mathcal{A}^*$, and $u^{\propto}_i = 0$ otherwise. Due to our rounding procedure, the overall number of units stocked by $U^{\propto}$ exceeds the capacity C by a factor of at most 2 since

$$\sum_{i\in\mathcal{A}^*} u^{\propto}_i = \sum_{i\in\mathcal{A}^*} \lceil \tilde{u}_i \rceil \leq \sum_{i\in\mathcal{A}^*} \tilde{u}_i + |\mathcal{A}^*| \leq \frac{C}{\text{OPT}_{\text{static}}} \cdot \sum_{i\in\mathcal{A}^*} r_i\psi_i + |\mathcal{A}^*| \leq 2C \ . \tag{21}$$

Since the newsvendor-like objective function $\mathcal{L}$ has diminishing marginals (Muckstadt and Sapra 2010, Chap. 5), we infer that $\mathcal{L}(U^{\propto}) \leq 2\cdot\mathcal{L}(U_{\mathcal{C}})$ by observing that $U_{\mathcal{C}}$ is an optimal inventory vector for $\mathcal{L}$ with C units, while $U^{\propto}$ has at most $2C$ units. Therefore, by the lower bound (19),

$$\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{C}}\right)\right] \geq \mathcal{L}\left(U_{\mathcal{C}}\right) \geq \frac{\mathcal{L}\left(U^{\propto}\right)}{2} \ . \tag{22}$$

Comparing the upper bound (20) **with the lower bound** (22)**.** In the next claim, we leverage the structure of $U^{\propto}$ as well as the IFR property to obtain a lower bound on the marginal contribution of each product in $\mathcal{A}^*$ toward $\mathcal{L}(U^{\propto})$. The proof is given in Appendix C.7

Claim 3.12. *For every product $i \in \mathcal{A}^*$, $\mathbb{E}[\bar{Y}_i(u^{\propto}_i)] \geq (1/2)\cdot\min\{u^{\propto}_i, \mathbb{E}[Y_i]\}$.*

By plugging Claim 3.12 into the lower bound stated in (22), we conclude that

$$\begin{aligned}
\mathbb{E}\left[R\left(U_{\mathcal{C}}\right)\right] &\geq \frac{1}{4}\cdot\sum_{i\in\mathcal{A}^*} r_i\cdot\min\left\{u^{\propto}_i, \mathbb{E}\left[Y_i\right]\right\} \\
&\geq \frac{1}{4}\cdot\sum_{i\in\mathcal{A}^*} r_i\cdot\min\left\{\tilde{u}_i, \mathbb{E}\left[M\right]\cdot\psi_i\right\} \\
&= \frac{1}{4}\cdot\sum_{i\in\mathcal{A}^*} r_i\tilde{u}_i\cdot\min\left\{1, \frac{\mathbb{E}\left[M\right]\cdot\text{OPT}_{\text{static}}}{C\cdot r_i}\right\}
\end{aligned}$$

$$\begin{aligned} &\geq \frac{\mathrm{OPT}_{\mathrm{static}} \cdot \min\left\{C, \mathbb{E}\left[M\right]\right\}}{4C} \cdot \sum_{i \in \mathcal{A}^*} \tilde{u}_i \\ &\geq \frac{\mathbb{E}\left[R\left(U^*_{\mathcal{C}}\right)\right]}{4C} \cdot \sum_{i \in \mathcal{A}^*} \tilde{u}_i \\ &= \frac{\mathbb{E}\left[R\left(U^*_{\mathcal{C}}\right)\right]}{4} \ . \end{aligned}$$

Here, the second inequality holds since $u_i^{\propto} = \lceil \tilde{u}_i \rceil$, the next equality follows from the definition of $\tilde{u}_i$, the third inequality holds since all products in $\mathcal{A}^*$ are expensive by Claim 3.10, meaning that $r_i \geq \mathrm{OPT}_{\mathrm{static}}$, the fourth inequality is derived from the upper bound in (20), and the last equality holds since $\sum_{i \in \mathcal{A}^*} \tilde{u}_i = C$.

4 Computational Experiments

In this section, we show that our algorithmic approach has a superior empirical performance in comparison to existing heuristics on randomly-generated instances. In particular, substantial gains in the expected revenue are demonstrated against these heuristics, with better computational efficiency and robustness.

4.1 Generative models

Products and MNL parameters. Our simulations make use of $n = 20$ products and a capacity bound of C, taking values in $\{10, 25, 50, 100\}$. Instances of the MNL model are constructed by considering two alternative settings, with different levels of heterogeneity in revenues and preferences.

- *Setting A:* The preference weights w_i are i.i.d. samples of a uniform distribution over the interval $[0, 1]$. The per-unit selling prices r_i are i.i.d. random samples of a standard log-normal random variable (with $\mu = 0$ and $\sigma = 1$).
- *Setting B:* Here, we create instances having a greater dispersion of weights and prices. Specifically, the weights are generated through i.i.d. samples of a standard log-normal distribution, rescaled by a factor of 1/2 to remain on average equivalent to setting A. The per-unit selling prices r_i are now sampled from a log-normal distribution with $\mu = 0$ and $\sigma = 2$.

The demand. The random number of arriving customers M is generated through two families of distributions with finite support $0, \ldots, 100 = \bar{M}$: a truncated Poisson distribution and randomly generated nonparametric distributions. The former uses a random variable $\mathcal{P} \sim \mathrm{Poisson}(0.35 \cdot \bar{M})$, such that $M = \min\{\mathcal{P}, \bar{M}\}$. The latter nonparametric distributions are constructed as follows. To enforce the IFR property, we generate a decreasing sequence of $\bar{M}$ failure rates, each of at most 5%. To this end, we first draw $\bar{M}$ i.i.d. samples of the uniform distribution over the interval $[0, 0.04]$, which are next sorted by increasing values, to obtain a sequence $z_1 \leq \cdots \leq z_{\bar{M}}$. This sequence is used to specify the failure rate $\Pr[M = k | M \geq k] = z_{k+1}$ for every $k \in [0, \bar{M} - 1]$, and in addition, $\Pr[M = \bar{M} | M \geq \bar{M}] = 1$.

4.2 Tested heuristics

The performance of our algorithm is compared against five different heuristics, whose specifics are discussed in Appendix D: (i) a local search heuristic similar to that of Goyal et al. (2016); (ii) a gradient-descent approach based on a continuous extension of the revenue function, similar in spirit to the work of Mahajan and van Ryzin (2001); (iii) the dynamic programming formulations devised by Topaloglu (2013) for variants of our problem; (iv) the deterministic relaxation heuristic proposed by Honhon et al. (2010); and (v) a discrete-greedy heuristic. The latter forms a natural benchmark since our approach relies primarily on greedy decisions, with the main difference of stocking products within a restricted set, possibly with modified objective functions. In addition, we report the revenue performance of the subroutine used in Section 3.3 to compete against cheap products, that scales the inventory levels proportional to the expected sales within an optimal static assortment.

4.3 Additional technical details

We implemented our algorithm, as well as the above-mentioned heuristics, using the Python programming language. The experiments described in this section were conducted on a standard laptop with 2.5GHz Intel Core i5 processor and 8GB of RAM.

The number of tested instances for each combination of parameters is 20. We impose a time limit of 2000 seconds (per instance) for every algorithm. When this limit is reached before termination, we use an identical rounding procedure on the current solution. Specifically, letting U be the best inventory vector found after 2000 seconds, U is linearly scaled and rounded down to the nearest integral vector: $U'_i = \lfloor \frac{U_i}{\|U\|_1} \cdot C \rfloor$. Finally, units are greedily added to U' until an inventory vector of exactly C units is obtained.

To approximately evaluate the expected revenue function, each call to the random oracle results in 500 samples. Although the number of samples needed to derive our theoretical guarantee in Lemma A.3 could be significantly larger, we observed in preliminary experiments that a greater number of samples has negligible impact on the performance of the algorithms considered. Moreover, the gradient-descent approach and discrete-greedy algorithm become rapidly impractical for the instances tested when the number of samples is increased.

Relative performance. For each instance tested, obtaining an estimate of the optimal expected revenue through brute-force enumeration is computationally prohibitive. Furthermore, we are not aware of any good empirical upper bound on the optimal expected revenue. For example, using a sample average approximation method, the resulting problem can be formulated as an integer program. However, using a state-of-the-art commercial solver (Gurobi Optimization 2015), this IP incurred running times greater than 1 hour, even for the simplest instances with $n = \bar{M} = 20$, $C = 10$, and 500 samples. The latter approach can be made more tractable using a relaxation, where a custom solution is computed for each sample-path realization through a separate IP. However, this approach produces low quality approximations.

For these reasons, we do not estimate the exact optimality gap attained by each algorithm. Instead, the algorithms are compared on a relative basis where, for each instance, the benchmark is set as the expected revenue of the most profitable inventory vector obtained through all algorithms considered. Then, the *relative performance* of each algorithm is defined as the ratio between its expected revenue and that of the benchmark. For example, if our algorithm attains for a particular instance an expected revenue of 1, while all tested heuristics generate an expected revenue of 0.9, the relative performance is 100% for our algorithm, and 90% for the others.

4.4 Results

Practical performance. As shown in Table 1, our algorithm exhibits moderate to significant performance gains in comparison to the heuristics under consideration, except for two configurations (out of 16) with capacity $C = 10$, where the local search shows slightly better revenue performance. Specifically, the average performance gains of our algorithm range from -0.7% to 39.6%. Overall, the expected revenues are increased by an average factor of 5.5% in comparison to the proportional scaling heuristic (a subroutine of our algorithm to compete against cheap products), 4.8% in comparison to the Poisson-based dynamic program, 4.4% in comparison to the normal approximation-based dynamic program, 10.1% in comparison to the deterministic relaxation, 12.7% against the local search algorithm, 9.3% in comparison to gradient descent, and 6.1% in comparison to discrete-greedy. In addition, our algorithm is robust, as it outperforms all heuristics for 62% of the instances tested.

Table 1: Average revenue performance of the different algorithms tested.

Parameters				**Relative revenue performance (%)**							
M	Setting	n	C	ALG	PROP	P-DP	N-DP	DET	LS	GD	DG
Poisson	A	20	10	98.5	88.6	96.9	94.4	92.5	**99.3**	80.8	98.5
		20	25	**99.7**	98.3	98.2	96.6	80.2	91.4	93.4	98.1
		20	50	**99.8**	98.6	95.3	98.6	90.7	89.1	97.6	96.5
		20	100	**99.8**	97.7	96.9	97.1	85.8	85	96.9	92.1
	B	20	10	98	91.7	96.6	96.1	94.1	**98.5**	88.1	97.9
		20	25	**99.3**	97.9	94.9	96.3	91.2	91.4	95.7	93.1
		20	50	**99.7**	98.3	95.7	97.5	93.2	90.5	97.8	92.3
		20	100	**99.7**	97.8	96.6	97.9	86.4	89.7	96.6	88
Nonparametric	A	20	10	**98.8**	75.2	94.8	90.8	89.4	98.3	78.2	98
		20	25	**98.7**	90.4	88.1	92.5	91.9	86	83.7	98.4
		20	50	**99.8**	98.3	83.4	91.9	93.1	76.6	92	93.1
		20	100	**99.9**	97	95.1	94	88.6	60.3	88.5	83.6
	B	20	10	**98.2**	82	91.9	90.4	93.7	97.8	80.8	97.6
		20	25	**99**	94.9	97.7	92.3	91.7	84.7	88.5	94
		20	50	**99.5**	95	92.3	94.2	78.8	74.1	90.1	84.3
		20	100	**98.9**	96.5	94.9	95	83.4	70.3	88.7	83.1

Here, ALG designates our algorithm, PROP is the subroutine of our algorithm to compete against cheap products, P-DP is the dynamic program under a Poisson process, N-DP is the dynamic program under a normal approximation, DET is the deterministic relaxation, LS corresponds to the local search heuristic, GD designates the gradient-descent approach, and DG is the discrete-greedy algorithm.

Although discrete-greedy is closest in spirit to our algorithm, particularly for expensive products, the relative performance gap between the two approaches is significant. The improvements observed on the computational front can be explained in that the discrete-greedy algorithm is given access at each iteration to a larger space of incremental actions (augmenting the inventory level of any product) in comparison to how our greedy procedures operate (augmenting inventory levels within restrictive assortments). On the revenue performance front, our results are somewhat surprising, as one could expect that a more constrained decision space would limit the flexibility in constructing the stocking policy. However, the numerical results reported here provide empirical evidence that the structural restrictions we impose on the stocking policy are in fact beneficial, not only for purposes of analysis, but also in practical settings.

It is worth noting that, as the price and weight variabilities increase from setting A to setting B, the relative performance of the discrete-greedy and the Poisson-based dynamic program is negatively affected. On the other hand, our algorithm along with the deterministic relaxation, the normal approximation, and the local search algorithm have better robustness in the face of heterogeneity. Intuitively, in such settings, it is expected that near-optimal inventory vectors are concentrated over fewer products. Thus, it is not surprising that the deterministic relaxation and normal approximation turn out to be more accurate, as further corroborated by Table 2 below.

Random stock-outs vs. other models. Our experiments demonstrate the benefits of using a realistic modeling approach, that captures the stochastic nature of stock-out events, even though the resulting model is not solved optimally. The proportional scaling heuristic, based on the optimal static assortment (and used as a subroutine to compete against cheap products), has a rather satisfactory performance, which is not entirely surprising in light of the guarantees established in Section 3.3. However, its average revenue loss can be as large as 23.6%, thus supporting the value of jointly studying the assortment and inventory dynamics. As shown in Table 2, the model proposed by Topaloglu (2013) tends to be less accurate for larger capacity values as well as for larger prices and weight variabilities. These trends are more pronounced for the Poisson-based dynamic program, while the normal-based algorithm tends to be more robust. This observation suggests that the flexibility to vary the assortment over time provides greater value to the retailer in such regimes. Interestingly, the accuracy of the deterministic relaxation tends to vary in the opposite direction, as a function of the different parameters. Indeed, the quality of the approximation and the optimality gap often improve as we scale-up the different parameters. One possible intuitive explanation is that fluid approximation models become more relevant asymptotically, as well as more tractable.

Running time. As shown in Table 3, the proposed algorithm emerges as the fastest algorithm tested in our implementation. As mentioned earlier, the relative efficiency of our algorithm is mainly due to the restrictive greedy rules, which limit the actions examined prior to each incremental decision, in comparison to the local search and discrete-greedy heuristics. Observe however that the running times of the gradient-descent algorithm and the dynamic programs largely depend on the parameters chosen in our implementation (the step size and discretization parameter, respectively).

The gradient-descent algorithm is particularly inefficient from a computational perspective, presumably due to the likely existence of local minima, where the algorithm progresses at a

Table 2: Average absolute approximation errors and optimality gap of the mixed integer program (Poisson model).

Setting	n	C	P-DP	N-DP	DET	GAP
	20	10	3%	14%	10%	63%
A	20	25	7%	28%	4%	34%
	20	50	28%	32%	1%	9%
	20	100	231%	145%	5%	24%
	20	10	6%	20%	8%	20%
B	20	25	12%	42%	2%	9%
	20	50	79%	52%	3%	12%
	20	100	427%	219%	14%	14%

Here, GAP refers to the MIP optimality gap for the deterministic relaxation algorithm after termination (either reaching the time limit of 2000 seconds, or an optimality gap of at most 0.5%). The additional entries correspond to the average absolute errors, when comparing the optimal objective value of the models DP-P, N-DP, and DET to the actual expected revenue generated under our stochastic dynamic substitution model. For example, if the Poisson-based dynamic program, whose objective function is denoted by $\text{obj}_P(\cdot)$, returns the inventory vector U_P, its error is measured as $|\frac{\mathbb{E}[\mathcal{R}(U_P)]}{\text{obj}_P(U_P)} - 1|$.

Table 3: Average running time of the algorithms tested.

Parameters				**Average running time (sec.)***						
Model	Setting	n	C	ALG	P-DP	N-DP	DET	LS	GD	DG
		20	10	**26**	1174	465	2000	232	175	70
	A	20	25	**95**	1702	859	2000	719	775	218
Poisson		20	50	**155**	2000	1408	1625	509	1632	465
		20	100	**381**	2000	2000	1984	562	2000	971
		20	10	**13**	977	434	1601	110	451	67
	B	20	25	**49**	2000	881	1121	209	1170	219
		20	50	**101**	2000	1531	609	288	1603	457
		20	100	**259**	2000	2000	1331	422	2000	958
		20	10	**44**	985	445	2000	353	134	96
	A	20	25	**131**	2000	736	2000	481	1181	315
Nonparametric		20	50	**292**	2000	1319	1421	858	2000	781
		20	100	691	2000	1919	1365	**493**	2000	1647
		20	10	**27**	1865	309	2000	151	573	105
	B	20	25	**44**	985	445	2000	353	134	96
		20	50	**171**	2000	1112	1677	317	2000	732
		20	100	478	2000	2000	1504	**465**	2000	2000

*Recall that every algorithm is being run with a time limit of 2000 seconds.

slower rate towards the final solution. Furthermore, due to its parameter dependency (step size and stopping criterion), the gradient-descent algorithm poses several implementation challenges. Even though we used here the best parameter values found by trial and error, it is still possible that fine-tuned parameters for each configuration could improve the performance in terms of optimality gaps and running times. Interestingly, as mentioned above, the optimality gap of the deterministic MIP shrinks when we increase the capacity or the price and weights variability, possibly since the combinatorial aspects are likely to be mitigated in an asymptotic regime, where the LP relaxation becomes tighter.

5 Concluding Remarks

Applications of restricted properties. To derive our main result, the analysis in Section 3.2 unravels hidden submodularity-like properties satisfied by the expected revenue function, and utilizes new notions of monotonicity and submodularity. One interesting direction for future research is to investigate whether these weaker properties could be used for closely related models in dynamic assortment planning, such as the Markov chain choice model (Blanchet et al. 2016, Désir et al. 2015, Feldman and Topaloglu 2017), which generalizes MNL, or a mixture of Multinomial Logits (Bront et al. 2009, Méndez-Díaz et al. 2014, Rusmevichientong et al. 2014, Feldman and Topaloglu 2015) with fixed number of customer types.

Approximation guarantee without IFR. It would be interesting to determine whether a constant-factor approximation for the MNL-based dynamic assortment planning problem can be obtained under general (non-IFR) demand distributions. Here, we mention in passing that the methods developed in this paper allow us to obtain an $O(\log n)$-approximation in this general setting, using an appropriate decomposition of the underlying set of products, combined with greedy procedures. The specifics of the resulting algorithm and its analysis are given in Appendix B.

Open questions. A natural direction for future research is that of obtaining improved approximation guarantees, which seems particularly challenging through the techniques developed in this paper, specifically due to the optimality loss incurred by subadditivity-based bounds. Another important theoretical question is to establish hardness of approximation results for dynamic substitution models. In fact, it might be NP-hard to evaluate (even approximately) the expected revenue function at a given inventory vector. That being said, due to the stochastic nature of this problem, any complexity results along these lines would be very interesting to obtain. Along these lines, yet another fundamental direction to consider is that of obtaining reasonable upper bounds on the optimum value. Such bounds will be valuable in measuring the practical performance of future algorithms in this context.

Finally, given their greedy nature and scalability, the algorithms we present are applicable in a broad range of settings. These include, for instance, product-specific per-unit costs, general knapsack constraints for storage or display, matroid/extendibility constraints on the assortment offered, etc. For such settings (and combinations thereof), the only requirement is being able to efficiently solve the corresponding static formulation. However, in its current form, our worst-case analysis holds under a cardinality constraint, similar to previous analytical work on approximation algorithms for dynamic assortment planning (Goyal et al. 2016, Segev 2015, Aouad et al. 2015). An interesting open question is that of devising provably-good algorithms for more general constraint structures, which seem to require further technical developments.

A General Constant-Factor Approximation

The approximation algorithm proposed in Section 3 relies on the efficient oracle assumption in order to compute the expected revenue generated by any given inventory vector. However, for the Multinomial Logit choice model, whether or not the expected revenue function can be evaluated in polynomial time (even approximately) is still an open question. We work around this difficulty by decomposing the set of products beforehand, and arguing that the terms requiring more effort from an optimization standpoint admit a sampling-based evaluation oracle, compatible with the algorithm developed in Section 3. Consequently, we establish the following theorem.

Theorem A.1. *For any* $\epsilon \in (0, 1/4)$ *and* $\delta > 0$*, the dynamic assortment planning problem under the Multinomial Logit choice model with IFR demand distribution can be approximated within a factor of* $0.122 - \epsilon$*, with probability at least* $1 - \delta$*. The running time of our algorithm is polynomial in the input size,* $n^{1/\epsilon}$*, and* $1/\delta$*. When* $C \geq n$*, this factor can be improved to* $0.151 - \epsilon$*.*

High-level overview of the algorithm. To work around the estimation obstacle, we make use of the decomposition idea explained in Section 2.2. Somewhat informally, sampling procedures fail to estimate the expected revenue accurately when there are very low probability purchase events, that require an exponential number of samples to be observed. Such rare events are possible when there is large variability between the preference weights of different products. Thus motivated, as a preliminary step, we partition the set of products into two classes, *light* and *heavy*, based on their respective MNL preference weights.

As a result, our decomposition generates two subproblems: one instance exclusively formed by heavy products, and another instance comprised of light products. Using an appropriate estimator, we show that the expected revenue in the heavy products instance can be efficiently approximated through sampling. Consequently, the methods developed in Section 3 provide a polynomial-time randomized algorithm, with a constant-factor worst-case guarantee. On the other hand, we show that a relatively simple approximation scheme can be derived for the light products instance.

Partition of products. To formalize this approach, the collection of products $[n]$ is decomposed into two sets:

- The set $\mathcal{L}$ of light products, consisting of those with $w_i \in (0, \epsilon/n]$.
- The set $\mathcal{H}$ of heavy products, with $w_i \in (\epsilon/n, \infty)$.

Product elimination. We further restrict attention to a smaller subset of heavy products, by eliminating in advance certain products whose revenue contribution toward $\mathbb{E}[\mathcal{R}(U^*)]$ is negligible. Specifically, let $i_{\max}$ be the heavy product that maximizes the quantity $r_i w_i/(1+w_i)$ over all products $i \in \mathcal{H}$ stocked by $U^*_{\mathcal{H}}$. From an algorithmic perspective, $i_{\max}$ can be guessed by considering $|\mathcal{H}|$ options, and we can now define the residual collection of heavy products $\tilde{\mathcal{H}} = \{i \in \mathcal{H} : \frac{\epsilon^2 r_{i_{\max}}}{2n^2 C} \leq r_i \leq \frac{2n^2 C \cdot r_{i_{\max}}}{\epsilon^3}\}$.

Upper bound on the optimal expected revenue. We now argue that the classes of products $\mathcal{L}$ and $\tilde{\mathcal{H}}$ are sufficient to compete against U^*. Recall that $U^*_{\mathcal{L}}$ denotes the projection of the optimal inventory vector U^* on light products, i.e., the vector obtained from U^* by setting the inventory levels of $[n] \setminus \mathcal{L}$ to zero. The vector $U^*_{\tilde{\mathcal{H}}}$ is defined in an analogous way. By exploiting the subadditive nature of the expected revenue function (see Lemma 2.2), we derive an upper bound on $\mathbb{E}[\mathcal{R}(U^*)]$ in the next lemma, whose proof is given in Appendix C.8.

Lemma A.2. $\mathbb{E}[\mathcal{R}(U^*_{\mathcal{L}})] + \mathbb{E}[\mathcal{R}(U^*_{\tilde{\mathcal{H}}})] \geq (1-2\epsilon) \cdot \mathbb{E}[\mathcal{R}(U^*)]$.

A.1 Efficient oracle for heavy products

Here, we show that the subproblem restricted to the heavy products $\tilde{\mathcal{H}}$ admits an efficient oracle. That is, for any error parameter $\epsilon > 0$ and confidence level $\delta > 0$, we devise a procedure to evaluate the expected revenue within a multiplicative factor of $1 \pm \epsilon$, running in time polynomial in the input size, $1/\epsilon$, and $1/\delta$.

For this purpose, suppose we are given an inventory vector U that stocks at most C units of products in $\tilde{\mathcal{H}}$, and wish to estimate $\mathbb{E}[\mathcal{R}(U)]$. Our evaluation procedure samples $L = \lceil 64C^6 n^{10}/(\epsilon^{12}\delta) \rceil$ independent realizations $R_1, \ldots, R_L$ of the random variable $\mathcal{R}(U)$ conditional on $M \geq 1$. These conditional realizations are obtained by sampling from a modified instance, where the number of arrivals M is replaced by $M|M \geq 1$. Next, the expected revenue $\mathbb{E}[\mathcal{R}(U)]$ is estimated by the unbiased estimator

$$\tilde{R} = \Pr[M \geq 1] \cdot \frac{1}{L} \cdot \sum_{\ell=1}^{L} R_\ell \ . \tag{23}$$

Lemma A.3. *The estimator $\tilde{R}$ provides an efficient oracle for the expected revenue function, i.e.,*

$$\Pr\left[\left|\frac{\tilde{R}}{\mathbb{E}[\mathcal{R}(U)]} - 1\right| \geq \epsilon\right] \leq \delta \ .$$

Proof. The proof relies on bounding the variance of the conditional revenue relative to its expected value, before applying Chebyshev's inequality. Since U stocks at most C units, the random variable $\mathcal{R}(U)|M \geq 1$ is upper bounded by $C \cdot r_{i_1}$ for any realization, where $i_1 \in \tilde{\mathcal{H}}$ is the most expensive product stocked by U. Also, letting i_2 be the maximal preference weight product stocked by U, an immediate lower bound on the expectation of this random variable is given by

$$\mathbb{E}[\mathcal{R}(U)| M \geq 1] \geq \frac{r_{i_2} w_{i_2}}{1 + |\tilde{\mathcal{H}}| \cdot w_{i_2}} \geq \frac{\epsilon r_{i_2}}{2n} \ , \tag{24}$$

where the first inequality accounts for the expected revenue due to the first arriving customer, who purchases a unit of product i_2 with probability at least $w_{i_2}/(1+|\tilde{\mathcal{H}}|\cdot w_{i_2})$, given that i_2 the has maximal preference weight among all products stocked by U, and the second inequality holds since $w_{i_2} \geq \epsilon/n$. Note that, since $\mathbb{E}[\mathcal{R}(U)|M \geq 1] = \mathbb{E}[\mathcal{R}(U)]/\Pr[M \geq 1]$, we have:

$$\Pr\left[\left|\tilde{R} - \mathbb{E}\left[\mathcal{R}\left(U\right)\right]\right| \geq \epsilon \cdot \mathbb{E}\left[\mathcal{R}\left(U\right)\right]\right] = \Pr\left[\left|\frac{1}{L} \cdot \sum_{\ell=1}^{L} R_\ell - \mathbb{E}\left[\mathcal{R}\left(U\right)\middle| M \geq 1\right]\right| \geq \epsilon \cdot \mathbb{E}\left[\mathcal{R}\left(U\right)\middle| M \geq 1\right]\right] .$$

Hence, by Chebyshev's inequality,

$$\Pr\left[\left|\tilde{R} - \mathbb{E}\left[\mathcal{R}\left(U\right)\right]\right| \geq \epsilon \cdot \mathbb{E}\left[\mathcal{R}\left(U\right)\right]\right] \leq \frac{\mathrm{var}((1/L) \cdot \sum_{\ell=1}^{L} R_\ell)}{\epsilon^2 \cdot (\mathbb{E}\left[\mathcal{R}\left(U\right)\middle| M \geq 1\right])^2} \leq \frac{4C^2 n^2}{\epsilon^2 L} \cdot \frac{r_{i_1}^2}{r_{i_2}^2} \leq \delta ,$$

where the second inequality follows from (24), along with the upper bound of $C^2 \cdot r_{i_1}^2$ on the second moment of each sample R_ℓ, and the last inequality holds since $L = \lceil 64C^6 n^{10}/(\epsilon^{12}\delta) \rceil$ while $r_{i_1}/r_{i_2} \leq 4C^2 n^4/\epsilon^5$, by definition of $\tilde{\mathcal{H}}$. ■

A.2 Approximation scheme for light products

The approach for handling light products $\mathcal{L}$ relies on identifying a newsvendor-like lower bound, in the spirit of Section 3.3. The important observation is that, when we are restricted to stocking only light products, each arriving customer faces a random assortment $S \subseteq \mathcal{L}$ with total weight $w(S) \leq |S|\cdot\epsilon/n \leq \epsilon$. Thus, as long as product $i \in \mathcal{L}$ is available, it is purchased by an arriving customer with probability at least $w_i/(1+\epsilon) \geq (1-\epsilon)\cdot w_i$, regardless of what the other available products are. Hence, at least intuitively, at the cost of losing a negligible factor in optimality, one could view the contribution of each product to the expected revenue as if it depends only on the initial number of units stocked.

Algorithm. To turn this intuition into a concrete argument, suppose that U is an inventory vector that stocks only light products. Then, the number of units purchased from each product $i \in \mathcal{L}$ is stochastically larger than the random variable $\bar{Y}_i(u_i) = \min\{Y_i, u_i\}$, where $Y_i \sim B(M, (1-\epsilon)\cdot w_i)$. Therefore,

$$\mathbb{E}\left[\mathcal{R}(U)\right] \geq \sum_{i \in \mathcal{L}} r_i \cdot \mathbb{E}\left[\bar{Y}_i(u_i)\right] . \tag{25}$$

Our algorithm optimizes the latter newsvendor-like lower bound, by computing an optimal solution to the following problem:

$$\max_U \left\{ \sum_{i \in \mathcal{L}} r_i \cdot \mathbb{E}\left[\bar{Y}_i(u_i)\right] : \sum_{i \in \mathcal{L}} u_i \leq C \right\} . \tag{26}$$

As explained in Section 3.3, an optimal solution to this problem can be computed efficiently by means of a greedy procedure. Note that the expectation $\mathbb{E}[\bar{Y}_i(u_i)]$ can be computed in polynomial time with respect to C and the maximum value of M, using a simple dynamic program.

The next lemma shows that the inventory vector $U_{\mathcal{L}}$, obtained by solving problem (26), guarantees a $(1-\epsilon)$-approximation with respect to the inventory vector $U^*_{\mathcal{L}}$.

Lemma A.4. $\mathbb{E}\left[\mathcal{R}(U_{\mathcal{L}})\right] \geq (1-\epsilon) \cdot \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{L}})\right]$.

Proof. First, observe that the lower bound (25) can be complemented by an upper bound on the expected revenue of $U^*_{\mathcal{L}}$. Specifically, letting $Y^*_i \sim B(M, w_i)$ and $\bar{Y}^*_i = \min\{Y^*_i, u^*_i\}$, we have

$$\mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{L}})\right] \leq \sum_{i \in \mathcal{L}} r_i \cdot \mathbb{E}\left[\bar{Y}^*_i\right] \ . \tag{27}$$

Based on inequalities (25) and (27), since $U_{\mathcal{L}}$ is an optimal solution to problem (26), it remains to show that the objective value of $U^*_{\mathcal{L}}$ with respect to the latter problem is at least $(1-\epsilon) \cdot \sum_{i \in \mathcal{L}} r_i \cdot \mathbb{E}[\bar{Y}^*_i]$. This follows by observing that $\mathbb{E}[\bar{Y}_i(u^*_i)] \geq (1-\epsilon) \cdot \mathbb{E}[\bar{Y}^*_i]$ for any product $i \in \mathcal{L}$, where the latter inequality is an immediate consequence of the next claim (proven in Appendix C.9), specialized for $\theta = 1-\epsilon$.

Claim A.5. *Let M be a non-negative integer-valued random variable, and suppose that $X \sim B(M, \alpha)$ and $Y \sim B(M, \theta\alpha)$, where $\alpha \in [0,1]$ and $\theta \in [0,1]$. For some integer C, let $\bar{X} = \min\{X, C\}$ and $\bar{Y} = \min\{Y, C\}$. Then, $\mathbb{E}[\bar{Y}] \geq \theta \cdot \mathbb{E}[\bar{X}]$.*

■

A.3 Conclusion

To summarize, our algorithm computes two approximate inventory vectors, corresponding to the weight classes $\mathcal{L}$ and $\tilde{\mathcal{H}}$, and eventually picks the one with maximal expected revenue.

- *Heavy products.* We employ the algorithm described in Section 3, for the subproblem restricted to the heavy products $\tilde{\mathcal{H}}$. This algorithm relies on the efficient oracle assumption, and therefore, we utilize the efficient sampling-based procedure described in Appendix A.1, running in time polynomial in the input size, $n^{1/\epsilon}$, and $1/\delta$. By Theorem 3.1, the random vector $U_{\tilde{\mathcal{H}}}$ returned by this algorithm satisfies $\mathbb{E}[\mathcal{R}(U_{\tilde{\mathcal{H}}})] \geq (0.139-\epsilon) \cdot \mathbb{E}[\mathcal{R}(U^*_{\tilde{\mathcal{H}}})]$, with probability at least $1-\delta$.

- *Light products.* The vector $U_{\mathcal{L}}$, returned by the algorithm described in Section A.2, is a $(1-\epsilon)$-approximation with respect to the expected revenue of $U^*_{\mathcal{L}}$, i.e., $\mathbb{E}[\mathcal{R}(U_{\mathcal{L}}] \geq (1-\epsilon) \cdot \mathbb{E}[\mathcal{R}(U^*_{\mathcal{L}}]$. In addition, this guarantee applies to a lower bound, that can be efficiently computed through dynamic programming.

Establishing Theorem A.1. Since we pick the best vector out of $U_{\mathcal{L}}$ and $U_{\tilde{\mathcal{H}}}$, with probability at least $1-\delta$, for any $\alpha \in [0,1]$ we obtain an expected revenue of at least

$$\begin{aligned} &\max\left\{\mathbb{E}\left[\mathcal{R}(U_{\mathcal{L}})\right], \mathbb{E}\left[\mathcal{R}(U_{\tilde{\mathcal{H}}})\right]\right\} \\ &\quad \geq \alpha \cdot \mathbb{E}\left[\mathcal{R}(U_{\mathcal{L}})\right] + (1-\alpha) \cdot \mathbb{E}\left[\mathcal{R}(U_{\tilde{\mathcal{H}}})\right] \\ &\quad \geq (1-8\epsilon) \cdot \left(\alpha \cdot \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{L}})\right] + 0.139 \cdot (1-\alpha)\mathbb{E}\left[\mathcal{R}(U^*_{\tilde{\mathcal{H}}})\right]\right) \ . \end{aligned}$$

By choosing $\alpha = 0.139/1.139 \approx 0.122$, we have

$$\begin{aligned}
&\max\left\{\mathbb{E}\left[\mathcal{R}(U_{\mathcal{L}})\right], \mathbb{E}\left[\mathcal{R}(U_{\tilde{\mathcal{H}}})\right]\right\} \\
&\quad \geq 0.122 \cdot (1-8\epsilon) \cdot \left(\mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{L}})\right] + \mathbb{E}\left[\mathcal{R}(U^*_{\tilde{\mathcal{H}}})\right]\right) \\
&\quad \geq (0.122 - 2\epsilon) \cdot \mathbb{E}\left[\mathcal{R}(U^*)\right] \ ,
\end{aligned}$$

where the last inequality is due to Lemma A.2. In the special case where $C \geq n$, an improved guarantee of $0.151 - \epsilon$ is derived by plugging the refined approximation ratio of $0.179 - \epsilon$ given by Theorem 3.1 for the heavy products vector $U_{\tilde{\mathcal{H}}}$.

B $O(\log n)$-approximation for non-IFR demand distributions

In what follows, recall that U^* is an optimal inventory vector, and for any subset of products $S \subseteq [n]$ we use U^*_S to denote the projection of U^* on S.

Step 1: Decomposition. Similar to the algorithm described in Appendix A, we begin by partitioning the set of products into the weight classes $\mathcal{L}$ and $\mathcal{H}$, by specifically choosing $\epsilon = 1/4$. Since our approximation algorithm for light products (Section A.2) does not rely on the IFR property, the resulting inventory vector $U_{\mathcal{L}}$ still attains a performance guarantee of 3/4 with respect to $U^*_{\mathcal{L}}$.

Now, let $i_{\max} \in \mathcal{H}$ be the most expensive heavy product. From an algorithmic perspective, this product can be guessed by considering $|\mathcal{H}|$ options. With this definition at hand, we construct the subset of products $\mathcal{H}^+ \subseteq \mathcal{H}$ whose selling price is at least $r_{i_{\max}}/(8n)$, and designate by $\mathcal{H}^-$ the remaining heavy products.

Step 2: Competing against cheap heavy products. Let $U_{\mathcal{H}^-}$ be the inventory vector that stocks C units of product $i_{\max}$. In the next claim, whose proof is deferred to the end of this section, we argue that $U_{\mathcal{H}^-}$ is at least as good revenue-wise as $U^*_{\mathcal{H}^-}$.

Claim B.1. $\mathbb{E}\left[\mathcal{R}(U_{\mathcal{H}^-})\right] \geq \mathbb{E}\left[\mathcal{R}(U^*_{\mathcal{H}^-})\right]$.

Step 3: Competing against expensive heavy products. We further decompose $\mathcal{H}^+$ into $K = \lceil \log(8n) \rceil$ nearly-uniform price classes $\mathcal{H}^+_1, \ldots, \mathcal{H}^+_K$, where $\mathcal{H}^+_k = \{i \in \mathcal{H}^+ : \frac{r_{i_{\max}}}{2^k} < r_i \leq \frac{r_{i_{\max}}}{2^{k-1}}\}$. Next, for every $k \in [K]$, our algorithm computes an inventory vector $U_{\mathcal{H}^+_k}$ that compete against $U^*_{\mathcal{H}^+_k}$. To this end, consider the subproblem where only products in $\mathcal{H}^+_k$ can be stocked, and let f^k be the corresponding expected revenue set function, that specifies the expected revenue associated with subsets of the extended collection of units $\mathcal{H}^+_k \times [C]$ (see Section 3.2). By rounding up the selling prices of products in $\mathcal{H}^+_k$ to $r_{i_{\max}}/2^{k-1}$, the resulting expected revenue set function $\tilde{f}^k$ clearly satisfies, for any subset of units S,

$$\frac{1}{2} \cdot \tilde{f}^k(S) \leq f^k(S) \leq \tilde{f}^k(S) \ . \tag{28}$$

On the other hand, it is easy to verify that, when all selling prices are equal, the static expected revenue function associated with an instance of the MNL model is non-decreasing, implying in particular that $\tilde{f}_1^k$ is non-decreasing. As a result, the problem of maximizing $\tilde{f}^k(S)$ over subsets S of at most C units falls within the special setting discussed in Section 3.2.4. Therefore, the standard greedy algorithm, combined with the sampling-based oracle of Appendix A (with appropriate error and confidence parameters), computes an inventory vector $U_{\mathcal{H}_k^+}$ such that, with probability at least $1-\delta/K$,

$$\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{H}_k^+}\right)\right] \geq \left(\frac{1}{2}\cdot\left(1-\frac{1}{e}\right)-\epsilon\right)\cdot\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}_k^+}\right)\right] \ , \tag{29}$$

where the latter performance guarantee follows from the approximation ratio of Section 3.2.4 combined with (28).

Step 4: Picking the most profitable inventory vector. Finally, the algorithm selects the most profitable inventory vector out of $U_{\mathcal{L}}, U_{\mathcal{H}^-}, U_{\mathcal{H}_1^+}, \ldots, U_{\mathcal{H}_K^+}$. Since the corresponding expected revenues are unknown, these vectors are compared using the randomized oracle (for $U_{\mathcal{H}^-}, U_{\mathcal{H}_1^+}, \ldots, U_{\mathcal{H}_K^+}$), and the previously-mentioned lower bound for $U_{\mathcal{L}}$. Given the subadditivity of the expected revenue function (see Lemma 2.2), since $\mathcal{L}, \mathcal{H}^-, \mathcal{H}_1^+, \ldots, \mathcal{H}_K^+$ form a partition of $[n]$, it follows that

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{L}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] + \sum_{k\in[K]}\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}_k^+}\right)\right] \geq \mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \ .$$

Therefore, by the union bound, with probability at least $1-\delta$ we obtain that

$$\begin{aligned}
\max\ &\left\{\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{L}}\right)\right], \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{H}^-}\right)\right], \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{H}_1^+}\right)\right], \ldots, \mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{H}_K^+}\right)\right]\right\} \\
&\geq \frac{(1/2)\cdot(1-1/e)-\epsilon}{K+2}\cdot\left(\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{L}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] + \sum_{k\in[K]}\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}_k^+}\right)\right]\right) \\
&\geq \frac{1}{4(K+2)}\cdot\mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \\
&= \Omega\left(\frac{1}{\log n}\right)\cdot\mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \ .
\end{aligned}$$

where the first inequality holds due to the performance guarantees stated in Lemma A.4, Claim B.1, and inequality (29).

Proof of Claim B.1.. Since the selling price of every product in $\mathcal{H}^-$ is at most $r_{i_{\max}}/(8n)$, an upper bound on the expected revenue of $U^*_{\mathcal{H}^-}$ is given by

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] \leq \frac{r_{i_{\max}}}{8n}\cdot\mathbb{E}\left[\min\{M,C\}\right] \ .$$

On the other hand, when initially stocking the inventory vector $U_{\mathcal{H}^-}$, until product $i_{\max}$ stocks out, each arriving customer purchases a unit of $i_{\max}$ with probability $w_{i_{\max}}/(1+w_{i_{\max}}) \geq \epsilon/(2n) = 1/(8n)$, where the latter inequality holds since $w_{i_{\max}} \geq \epsilon/n$, given that $i_{\max}$ is a heavy product. Consequently, letting $Y \sim B(M, 1/(8n))$, we have

$$\mathbb{E}\left[\mathcal{R}\left(U_{\mathcal{H}^-}\right)\right] \geq r_{i_{\max}}\cdot\mathbb{E}\left[\min\{Y,C\}\right] \geq \frac{r_{i_{\max}}}{8n}\cdot\mathbb{E}\left[\min\{M,C\}\right] \geq \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] \ .$$

where the second inequality follows from Claim A.5 specialized with $\theta = 1/(8n)$. ■

C Additional Proofs

C.1 Proof of Lemma 2.1

Algorithm. For a small cardinality value, i.e., $C < 1/\epsilon$, one could simply enumerate over all $O(n^{1/\epsilon})$ possible subsets, and pick the one with largest estimated objective value, according to an $(\epsilon/2, \delta/n^{1/\epsilon})$-oracle. It is not difficult to verify that, by the union bound, this enumerative procedure returns a $1-\epsilon$-approximate solution with probability at least $1-\delta$. For large cardinality values ($C \geq 1/\epsilon$), the algorithm is a standard greedy procedure. Starting with the empty set $S_0 = \emptyset$, we add in each step the element that generates the largest increase in the objective function, among all unpicked elements. At each step, we call the random $(\epsilon/(2C), \delta/(nC))$-oracle to evaluate the objective value associated with each unpicked element. Let $S_0, S_1, \ldots, S_C$ be the sequence of subsets corresponding to the different steps in the algorithm, and let S^* be a fixed optimal subset. We assume without loss of generality that $|S^*| = C$, as f is restricted-non-decreasing.

Analysis. Since the greedy algorithm makes at most nC calls to the randomized oracle, by the union bound the relative error associated with all estimates returned by the $(\epsilon/(2C), \delta/(nC))$-oracle is upper bounded by $\epsilon/(2C)$ with probability at least $1-\delta$. From this point on, we establish the desired approximation guarantee under this condition. Below, for any subsets S and T, we use $f_S(T)$ to denote the marginal variation in f when S is augmented by T, i.e., $f_S(T) = f(S \cup \{T\}) - f(S)$.

Claim C.1. *Let S and T be disjoint subsets, with $|S| + |T| \leq C$ and $T \neq \emptyset$. Then, for every $0 \leq k \leq |T|$, there exists $T_k \subseteq T$ with $|T_k| = k$ and $f_S(T_k) \geq (k/|T|) \cdot f_S(T)$.*

Proof. The proof follows by an inductive argument over k. The base case $k = 0$ is clearly satisfied by $T_0 = \emptyset$. For the general case, by the induction hypothesis, there exists $T_k \subseteq T$ with $|T_k| = k$ and $f_S(T_k) \geq (k/|T|) \cdot f_S(T)$. Letting $T \setminus T_k = \{e_1, \ldots, e_{|T|-k}\}$, we have

$$f_{S\cup T_k}(T \setminus T_k) = \sum_{j=0}^{|T|-k-1} f_{S\cup T_k \cup\{e_1,\ldots,e_j\}}(e_{j+1}) \leq \sum_{j=0}^{|T|-k-1} f_{S\cup T_k}(e_{j+1}) \ ,$$

where the latter inequality holds since f is restricted-submodular, by observing that $|S \cup T_k \cup \{e_1, \ldots, e_{|T|-k-1}\}| = |S| + |T| - 1 \leq C - 1$. Consequently, there exists $1 \leq j \leq |T| - k$ such that $f_{S\cup T_k}(e_j) \geq f_{S\cup T_k}(T \setminus T_k)/(|T| - k)$. As a result, by defining $T_{k+1} = T_k \cup \{e_j\}$, we obtain

$$\begin{aligned} f_S(T_{k+1}) &= f_S(T_k) + f_{S\cup T_k}(e_j) \\ &\geq f_S(T_k) + \frac{1}{|T|-k} \cdot f_{S\cup T_k}(T \setminus T_k) \\ &= \left(1 - \frac{1}{|T|-k}\right) \cdot f_S(T_k) + \frac{1}{|T|-k} \cdot f_S(T) \\ &\geq \left(\frac{k}{|T|} \cdot \left(1 - \frac{1}{|T|-k}\right) + \frac{1}{|T|-k}\right) \cdot f_S(T) \end{aligned}$$

$$= \quad \frac{k+1}{|T|} \cdot f_S(T) \ ,$$

where the second equality holds since $f_{S \cup T_k}(T \setminus T_k) = f_S(T) - f_S(T_k)$, and the second inequality proceeds from the inductive hypothesis. ■

It immediately follows from Claim C.1 that, for every $0 \leq k \leq C$, there exists a subset $S^*_{C-k} \subseteq S^*$ such that $|S^*_{C-k}| = C - k$, and

$$f_\emptyset \left(S^*_{C-k}\right) \geq \frac{C-k}{C} \cdot f_\emptyset \left(S^*\right) \ . \tag{30}$$

We can now analyze the sequence of subsets $S_0, \ldots, S_C$ produced by our (random) greedy procedure, where we use e_{k+1} to denote the unique element of $S_{k+1} \setminus S_k$, for every $0 \leq k \leq C-1$. To establish lower bounds on $f_{S_k}(e_{k+1})$ for every $0 \leq k \leq C-1$, we make the following case disjunction:

- *Case A: $S^*_{C-k} \setminus S_k \neq \emptyset$.* Observe that S_k and $S^*_{C-k} \setminus S_k$ are disjoint, and $|S_k| + |S^*_{C-k} \setminus S_k| \leq C$. By Claim C.1, since $S^*_{C-k} \setminus S_k \neq \emptyset$ by the case hypothesis, it follows that there exists an element $e \in S^*_{C-k} \setminus S_k$ such that

$$\begin{aligned} f_{S_k}(e) &\geq \frac{1}{|S^*_{C-k} \setminus S_k|} \cdot f_{S_k}\left(S^*_{C-k} \setminus S_k\right) \\ &\geq \frac{1}{C-k} \cdot \left(f_\emptyset\left(S_k \cup S^*_{C-k}\right) - f_\emptyset\left(S_k\right)\right) \\ &\geq \frac{1}{C-k} \cdot \left(f_\emptyset\left(S^*_{C-k}\right) - f_\emptyset\left(S_k\right)\right) \\ &\geq \frac{1}{C} \cdot f_\emptyset\left(S^*\right) - \frac{1}{C-k} \cdot f_\emptyset\left(S_k\right) \ , \end{aligned} \tag{31}$$

  where the third inequality holds since f is restricted-non-decreasing and $|S_k \cup S^*_{C-k}| \leq C$, while the last inequality proceeds from (30).

- *Case B: $S^*_{C-k} \setminus S_k = \emptyset$.* In this case, $S^*_{C-k} \subseteq S_k$, and therefore

$$f_\emptyset\left(S_k\right) \geq f_\emptyset\left(S^*_{C-k}\right) \geq \frac{C-k}{C} \cdot f_\emptyset\left(S^*\right) \ , \tag{32}$$

  where the first inequality holds since f is restricted-non-decreasing, and the last inequality follows from (30).

Concluding the analysis. Let $\mu \in [0,1]$ be a parameter that will be optimized later on, and let $L = \lfloor (1-\mu) \cdot C \rfloor$. When $f\left(S_C\right) \geq \mu \cdot f\left(S^*\right)$, our algorithm attains a μ-approximation. When $f\left(S_C\right) < \mu \cdot f\left(S^*\right)$, for every $k \leq (1-\mu) \cdot C$ we necessarily have $S^*_{C-k} \setminus S_k \neq \emptyset$, or otherwise

$$f_\emptyset\left(S_C\right) \geq f_\emptyset\left(S_k\right) \geq \frac{C-k}{C} \cdot f_\emptyset\left(S^*\right) \geq \mu \cdot f_\emptyset\left(S^*\right) \ ,$$

where the first inequality is due to f being restricted-non-decreasing, and the second inequality follows from (32); since $f(\emptyset) \geq 0$, the latter observation would imply that $f(S_C) \geq \mu \cdot f(S^*)$.

As a result, in this setting we have

$$\begin{aligned} f_{\emptyset}(S_C) &\geq f_{\emptyset}(S_L) \\ &= \sum_{k=0}^{L-1} f_{S_k}(e_{k+1}) \\ &\geq \sum_{k=0}^{L-1} \left(\frac{1}{C} \cdot f_{\emptyset}(S^*) - \frac{1}{C-k} \cdot f_{\emptyset}(S_k) - \frac{\epsilon}{C} \cdot f(S^*) \right) \\ &\geq \frac{L}{C} \cdot f_{\emptyset}(S^*) - \mu \cdot f_{\emptyset}(S^*) \cdot \sum_{k=0}^{L-1} \frac{1}{C-k} - \epsilon \cdot f(S^*) \\ &= \left(\frac{L}{C} - \mu \cdot (H_C - H_{C-L}) \right) \cdot f_{\emptyset}(S^*) - \epsilon \cdot f(S^*) \ , \end{aligned}$$

where $H_m = \sum_{k=1}^{m} \frac{1}{k}$ is the m-th harmonic number. Here, the first inequality holds since f is restricted-non-decreasing. The second inequality follows from (31), given that all estimates of the evaluation oracle are accurate up to a relative error of $\epsilon/(2C)$. The next inequality holds since $f(S_k) < \mu \cdot f(S^*)$ by hypothesis and since $f(\emptyset) \geq 0$.

Claim C.2. $L/C - \mu \cdot (H_C - H_{C-L}) \geq 1 - \mu - \mu \ln \mu - O(\epsilon)$.

Proof. Since $|H_n - \ln n| = \gamma + O(1/n)$, where γ is the EulerMascheroni constant, we have

$$\begin{aligned} \frac{L}{C} - \mu \cdot (H_C - H_{C-L}) &\geq \frac{L}{C} - \mu \cdot \ln \frac{C}{C-L} - \mu \cdot O\left(\frac{1}{C-L} \right) \\ &= \frac{\lfloor (1-\mu) \cdot C \rfloor}{C} + \mu \cdot \ln \left(1 - \frac{\lfloor (1-\mu) \cdot C \rfloor}{C} \right) - \mu \cdot O\left(\frac{1}{C - \lfloor (1-\mu) \cdot C \rfloor} \right) \\ &\geq 1 - \mu - \frac{1}{C} + \mu \ln \mu - O\left(\frac{1}{C} \right) \\ &= 1 - \mu + \mu \ln \mu - O(\epsilon) \ , \end{aligned}$$

where the first equality is obtained by substituting $L = \lfloor (1-\mu) \cdot C \rfloor$, and the last equality holds since $C \geq 1/\epsilon$. ■

Using the above claim, it follows that $f_{\emptyset}(S_C) \geq (1 - \mu - \mu \ln \mu - O(\epsilon)) \cdot f_{\emptyset}(S^*) - \epsilon \cdot f(S^*)$, and since $f(\emptyset) \geq 0$, we have $f(S_C) \geq (1 - \mu - \mu \ln \mu - O(\epsilon)) \cdot f(S^*)$. Therefore, our algorithm attains an overall approximation ratio of $\min\{\mu, 1 - \mu - \mu \ln \mu\} - O(\epsilon)$. The latter constant is optimized by picking $\mu^* \approx 0.318$, in which case we obtain a performance guarantee of $0.318 - O(\epsilon)$.

C.2 Proof of Lemma 2.2

To show that the expected revenue function is subadditive, it is sufficient to prove that, in a given inventory vector, the deletion of any unit can only increase the probability of every other unit to be purchased. Indeed, if $U = U_1 + U_2$, starting from U, we can iteratively delete units to obtain U_1 or U_2 while increasing the consumption probabilities of all remaining units at each step. This immediately implies that the expected revenue generated by remaining units may only increase as well. Finally, by combining the units of U_1 and U_2, the total expected revenue should be at least as large as that of U.

To formalize the above statement, for any unit v stocked by some inventory vector, we use $\mathcal{C}_v$ to denote the event "unit v is consumed". With this definition, it remains to establish the following claim.

Claim C.3. *Let U be some inventory vector, and let U^- be a vector obtained by deleting a single unit from U. Then, for any remaining unit v stocked by U^-, we have* $\Pr[\mathcal{C}_v|U^-] \geq \Pr[\mathcal{C}_v|U]$.

Let x be the product of which one unit was deleted in order to obtain U^- from U. There are three cases:

1. The unit v belongs to product x.
2. The unit v does not belong to product x, and no additional units of x are stocked (i.e., $U_x = 1$ and $U_x^- = 0$).
3. The unit v does not belong to product x, and at least one additional unit of x is stocked (i.e., $U_x \geq 2$ and $U_x^- = U_x - 1$).

In what follows, we prove the claim for case 2, noting that the remaining cases can be proven in a nearly-identical way. Moreover, by the formula of conditional expectations, it is sufficient to establish the claim for a deterministic demand variable M. For any event E, we use $\Pr_M[E|U]$ to denote the probability of E with M arriving customers and the initial inventory vector U. Finally, to simplify the notation, we make use of product 0 to designate the no-purchase option, with preference weight $w_0 = 0$.

The proof is by induction on $\sum_{i=1}^n u_i + M$. The base case, corresponding to $\sum_{i=1}^n u_i + M = 1$, implies that $M = 0$. Hence, $\Pr_M[\mathcal{C}_v|U] = \Pr_M[\mathcal{C}_v|U^-] = 0$.

In the general case, consider the random product X picked by the first arriving customer, including the no-purchase option 0. Then,

$$\begin{aligned}
\Pr_M[\mathcal{C}_v|U] &= \Pr_M[X=0|U] \cdot \underbrace{\Pr_M[\mathcal{C}_v|X=0,U]}_{\text{(I)}} \\
&\quad + \Pr_M[X=x|U] \cdot \Pr_M[\mathcal{C}_v|X=x,U] \\
&\quad + \sum_{i \in S(U)\setminus\{x\}} \Pr_M[X=i|U] \cdot \underbrace{\Pr_M[\mathcal{C}_v|X=i,U]}_{\text{(II)}} \\
&\leq \Pr_M[X=0|U] \cdot \Pr_M\left[\mathcal{C}_v|X=0,U^-\right] \\
&\quad + \Pr_M[X=x|U] \cdot \Pr_M[\mathcal{C}_v|X=x,U] \\
&\quad + \sum_{i \in S(U)\setminus\{x\}} \Pr_M[X=i|U] \cdot \Pr_M\left[\mathcal{C}_v|X=i,U^-\right] \ . \qquad (33)
\end{aligned}$$

Here, we use $S(U)$ to denote the set of products stocked by the vector U, i.e., $S(U) = \{i \in [n] : u_i > 0\}$. The inequality above hold since by the induction hypothesis,

$$\text{(I)} = \Pr_M[\mathcal{C}_v|X=0,U] = \Pr_{M-1}[\mathcal{C}_v|U] \leq \Pr_{M-1}\left[\mathcal{C}_v|U^-\right] = \Pr_M\left[\mathcal{C}_v|X=0,U^-\right] \ .$$

In addition, if v is the first available unit of product i to be purchased,

$$\text{(II)} = \Pr_M[\mathcal{C}_v|X=i,U] = 1 = \Pr_M\left[\mathcal{C}_v|X=i,U^-\right] \ ,$$

and otherwise,

$$(\text{II}) = \Pr_M \left[\mathcal{C}_v | X = i, U\right] = \Pr_{M-1} \left[\mathcal{C}_v | U_{-i}\right] \leq \Pr_{M-1} \left[\mathcal{C}_v | U^-_{-i}\right] = \Pr_M \left[\mathcal{C}_v | X = i, U^-\right] ,$$

where U_{-i} and U^-_{-i} stand for the residual inventory vectors after a unit of product i is consumed in U and U^-, respectively. On the other hand,

$$\begin{aligned} \Pr_M \left[\mathcal{C}_v | U^-\right] &= \Pr_M \left[X = 0 | U^-\right] \cdot \Pr_M \left[\mathcal{C}_v | X = 0, U^-\right] \\ &\quad + \sum_{i \in S(U^-)} \Pr_M \left[X = i | U^-\right] \cdot \Pr_M \left[\mathcal{C}_v | X = i, U^-\right] . \end{aligned} \tag{34}$$

To conclude the proof, note that since $S(U) = S(U^-) \uplus \{x\}$, by equation (33) and (34), we have

$$\begin{aligned} &\Pr_M \left[\mathcal{C}_v | U^-\right] - \Pr_M \left[\mathcal{C}_v | U\right] \\ &\quad \geq \left(\Pr_M \left[X = 0 | U^-\right] - \Pr_M \left[X = 0 | U\right]\right) \cdot \Pr_M \left[\mathcal{C}_v | X = 0, U^-\right] \\ &\quad + \sum_{i \in S(U^-)} \left(\Pr_M \left[X = i | U^-\right] - \Pr_M \left[X = i | U\right]\right) \cdot \Pr_M \left[\mathcal{C}_v | X = i, U^-\right] \\ &\quad - \Pr_M \left[X = x | U\right] \cdot \Pr_M \left[\mathcal{C}_v | X = x, U\right] \\ &\quad \geq \Pr_{M-1} \left[\mathcal{C}_v | U^-\right] \cdot \left(\sum_{i \in S(U^-) \cup \{0\}} \Pr_M \left[X = i | U^-\right] - \sum_{i \in S(U) \cup \{0\}} \Pr_M \left[X = i | U\right] \right) \\ &\quad = 0 , \end{aligned}$$

where the first inequality holds since $\Pr_M \left[X = 0 | U^-\right] \geq \Pr_M \left[X = 0 | U\right]$ and $\Pr_M \left[X = i | U^-\right] \geq \Pr_M \left[X = i | U\right]$ by the choice probabilities of the MNL model, combined with the fact that $\Pr_M \left[\mathcal{C}_v | X = 0, U^-\right] = \Pr_{M-1} \left[\mathcal{C}_v | U^-\right]$ and $\Pr_M \left[\mathcal{C}_v | X = x, U\right] = \Pr_{M-1} \left[\mathcal{C}_v | U^-\right]$, while $\Pr_M \left[\mathcal{C}_v | X = i, U^-\right] = \Pr_{M-1} \left[\mathcal{C}_v | U^-_{-i}\right] \geq \Pr_{M-1} \left[\mathcal{C}_v | U^-\right]$ due to the inductive hypothesis. The last equality proceeds from observing that the two sums of probabilities are both equal to 1.

C.3 Proof of Claim 3.3

By construction, the marginal purchase probabilities of the random variable X_j coincide with the MNL probabilities given by P_j. It remains to show that this property propagates to the random variables $X_{i,j}$, X_i, and X through the chain of conditional distributions $X_{i,j} | X_j$, $X_i | X_{i,j}$, and $X | X_i$. To avoid redundancy, we only present the proof for the variable $X_{i,j}$; those of X_i and X are based on similar ideas.

Recall that $P_{i,j}$ is the product purchased by the first arriving customer in the assortment stocked by $S \cup \{i, j\}$. Thus, we need to show that $\Pr[X_{i,j} = \alpha] = w_\alpha / (1 + w(\mathcal{A}) + w_i + w_j)$ for any product $\alpha \in \mathcal{A}^{+ij} \cup \{0\}$. For any product $\alpha \in \mathcal{A}^{+j} \cup \{0\}$, we have

$$\begin{aligned} \Pr\left[X_{i,j} = \alpha\right] &= \sum_{\beta \in \mathcal{A}^{+j} \cup \{0\}} \Pr\left[X_j = \beta\right] \cdot \Pr\left[X_{i,j} = \alpha | X_j = \beta\right] \\ &= \Pr\left[X_j = \alpha\right] \cdot \Pr\left[X_{i,j} = \alpha | X_j = \alpha\right] \\ &= \frac{w_\alpha}{1 + w(\mathcal{A}) + w_j} \cdot \frac{1 + w(\mathcal{A}) + w_j}{1 + w(\mathcal{A}) + w_j + w_i} \end{aligned}$$

$$= \frac{w_\alpha}{1 + w(\mathcal{A}) + w_j + w_i}$$

where the second equality proceeds from equation (7), that guarantees $\Pr[X_{i,j} = \alpha | X_j = \beta] = 0$ for $\alpha \neq i$ and $\beta \neq \alpha$, and the next equality holds since the distribution of X_j is given by the MNL model with respect to products $\mathcal{A}^{+j} \cup \{0\}$, combined with equation (6). In addition,

$$\begin{aligned}
\Pr[X_{i,j} = i] &= \sum_{\beta \in \mathcal{A}^{+j} \cup \{0\}} \Pr[X_j = \beta] \cdot \Pr[X_{i,j} = i | X_j = \beta] \\
&= \frac{w_i}{1 + w(\mathcal{A}) + w_j + w_i} \cdot \sum_{\beta \in \mathcal{A}^{+j} \cup \{0\}} \Pr[X_j = \beta] \\
&= \frac{w_i}{1 + w(\mathcal{A}) + w_j + w_i}
\end{aligned}$$

where the second equality is due to our definition of $X_{i,j} | X_j$ (equation (5)).

C.4 Proof of Claim 3.4

To see why $(X_j | X_{i,j} = i) \sim X_j$, observe that the event $\{X_{i,j} = i\}$ is independent of the outcomes of X_j as stated by equation (5). Similarly, given equation (9) along with the equivalence $X_i \sim P_i$ shown in Claim 3.3, we infer that $(X_i | X_{i,j} = j) \sim X_i$.

To establish the next equivalence, $(X | X_{i,j} = i) \sim X$, observe that

$$X \sim (X | X_i = i) \sim (X | X_i = X_{i,j} = i) \sim (X | X_{i,j} = i) \ ,$$

where the first equivalence holds since the distributions of X and $X | X_i = i$ are both prescribed by the MNL model with respect to $\mathcal{A}$ (see equation (12) and Claim 3.3), and the second equivalence proceeds from the Markov property satisfied by the coupling $(X | X_i, X_{i,j}) \sim (X | X_i)$. Finally, the last equivalence follows from observing that the event $\{X_{i,j} = i\}$ is contained in $\{X_i = i\}$ due equation (8).

Finally, to show the equivalence $(X | X_{i,j} = j) \sim X$, we have

$$\begin{aligned}
\Pr[X = \alpha | X_{i,j} = j] &= \sum_{\beta \in \mathcal{A}^{+i} \cup \{0\}} \Pr[X_i = \beta | X_{i,j} = j] \cdot \Pr[X = \alpha | X_i = \beta, X_{i,j} = j] \\
&= \sum_{\beta \in \mathcal{A}^{+i} \cup \{0\}} \Pr[X_i = \beta] \cdot \Pr[X = \alpha | X_i = \beta] \\
&= \Pr[X = \alpha] \ ,
\end{aligned}$$

where the second equality is due to the equivalence $(X_i | X_{i,j} = j) \sim X_i$ and the Markov property.

C.5 Proof of Claim 3.9 (continued)

We begin by establishing a technical claim, useful for the upcoming analysis, whose proof is deferred to the end of this section.

Claim C.4. *For any subset $S \subseteq [N]$ of cardinality at most $C - 1$ and any unit $i \in [N]$,*

$$f_{M-1}(S \cup \{i\}) - f_{M-1}(S) \leq f_M(S \cup \{i\}) - f_M(S) \ .$$

We proceed with the remaining two cases: $X_{i,j} = \alpha$ where $\alpha \in \{i, j\}$.

Conditional on the event $\{X_{i,j} = i\}$**.** When $X_{i,j} = i$, our coupling method entails that $X_i = i$ as well due to equation (8). As a result,

$$\begin{aligned}
\mathbb{E}\left[R_M\left(S^{+ij}\right) - R_M\left(S^{+i}\right)\middle| X_{i,j} = i\right] &= \mathbb{E}\left[R_{M-1}\left(S^{+j}\right) - R_{M-1}(S)\right] \\
&\leq \mathbb{E}\left[R_M\left(S^{+j}\right) - R_M(S)\right] \\
&= \mathbb{E}\left[R_M\left(S^{+j}\right) - R_M(S)\middle| X_{i,j} = i\right] ,
\end{aligned}$$

where the first equality follows from the decomposition (17) by observing that the terms $r_{(X_{i,j}|X_{i,j}=i)} = r_{(X_i|X_{i,j}=i)} = r_i$ cancel out, and the next inequality holds due to Claim C.4. The last equality holds since $X_j|X_{i,j} = i$ and $X|X_{i,j} = i$ have the same distribution as X_j and X, respectively, as shown in Claim 3.4. Now, by reordering the terms in the above inequality,

$$\mathbb{E}\left[R_M\left(S^{+ij}\right) - R_M\left(S^{+j}\right)|X_{i,j} = i\right] \leq \mathbb{E}\left[R_M\left(S^{+i}\right) - R_M(S)|X_{i,j} = i\right] . \tag{35}$$

Conditional on the event $\{X_{i,j} = j\}$**.** In this case, our coupling method entails that $X_j = j$ as well. Indeed, using Bayes rule, equation (6) along with the marginal distributions of X_j and $X_{i,j}$ (see Claim 3.3), imply that $\Pr[X_j = j|X_{i,j} = j] = 1$. Therefore,

$$\begin{aligned}
\mathbb{E}\left[R_M\left(S^{+ij}\right) - R_M\left(S^{+j}\right)\middle| X_{i,j} = j\right] &= \mathbb{E}\left[R_{M-1}\left(S^{+i}\right) - R_{M-1}(S)\right] \\
&\leq \mathbb{E}\left[R_M\left(S^{+i}\right) - R_M(S)\right] \\
&= \mathbb{E}\left[R_M\left(S^{+i}\right) - R_M(S)\middle| X_{i,j} = j\right] ,
\end{aligned} \tag{36}$$

where the first equality is a consequence of (17) by observing that the terms $r_{(X_{i,j}|X_{i,j}=j)} = r_{(X_j|X_{i,j}=j)} = r_j$ cancel out, the next inequality follows from Claim C.4, and the last equality holds since $X_i|X_{i,j} = j$ and $X|X_{i,j} = j$ have the same distribution as X_i and X, respectively, by Claim 3.4.

Proof of Claim C.4. To establish the desired claim, recall that the random residual subsets of units at the k-th arrival, obtained in the proof of Lemma 3.6, respectively denoted by S_k and T_k when initially stocking S_1 and T_1 with $S_1 \subseteq T_1$ and $|T_1 \setminus S_1| \leq 1$, satisfy $S_k \subseteq T_k$ for every realization. In addition, using a transformation similar to that of equation (15), with $S_1 = S$ and $T_1 = S \cup \{i\}$, we have

$$(f_M(S \cup \{i\}) - f_M(S)) - (f_{M-1}(S \cup \{i\}) - f_{M-1}(S)) = \mathbb{E}[f_1(T_M) - f_1(S_M)] \geq 0 .$$

To understand the latter inequality, note that since $S_M \subseteq T_M$ for every realization, and since these subsets have cardinality at most C, we have $\mathbb{E}[f_1(T_M) - f_1(S_M)] \geq 0$ due to f_1 being restricted-non-decreasing. ■

C.6 Proof of Claim 3.10

Suppose on the contrary that there exists a product $i \in \mathcal{A}^*$ with a selling price of $r_i < \text{OPT}_{\text{static}} = \mathbb{E}[\mathcal{R}_1(\mathcal{A}^*)]$, where $\mathcal{R}_1(\mathcal{A})$ stands for the random revenue generated by a single customer, when the set of stocked products is $\mathcal{A}$. By calculations identical to those leading to equation (14),

$$\mathbb{E}\left[\mathcal{R}_1(\mathcal{A}^*)\right] = \frac{w_i}{1+w(\mathcal{A}^*)} \cdot r_i + \left(1 - \frac{w_i}{1+w(\mathcal{A}^*)}\right) \cdot \mathbb{E}\left[\mathcal{R}_1\left(\mathcal{A}^* \setminus \{i\}\right)\right] \ .$$

In other words, $\mathbb{E}[\mathcal{R}_1(\mathcal{A}^*)]$ can be written as a convex combination of r_i and $\mathbb{E}[\mathcal{R}_1(\mathcal{A}^* \setminus \{i\})]$. Since $r_i < \mathbb{E}[\mathcal{R}_1(\mathcal{A}^*)]$, it follows that $\mathbb{E}[\mathcal{R}_1(\mathcal{A}^* \setminus \{i\})] > \mathbb{E}[\mathcal{R}_1(\mathcal{A}^*)]$, contradicting the optimality of $\mathcal{A}^*$.

C.7 Proof of Claim 3.12

The proof relies on the following technical claims regarding IFR distributions.

Lemma C.5 (Goyal et al. (2016))**.** *Let M be a non-negative integer-valued IFR random variable. For any $\alpha \in [0,1]$, the random variable $X \sim B(M,\alpha)$ also follows an IFR distribution.*

Lemma C.6 (Aouad et al. (2015))**.** *Let X be a non-negative IFR random variable, and for some constant C let $\bar{X} = \min\{X, C\}$. Suppose that $\mathbb{E}[\bar{X}] \leq \delta C$ for $\delta \in [0,1]$. Then, $\mathbb{E}[\bar{X}] \geq (1-\delta) \cdot \mathbb{E}[X]$.*

We argue that $\mathbb{E}[\bar{Y}_i(u_i^{\propto})] \geq \mathbb{E}[Y_i]/2$ whenever $\mathbb{E}[\bar{Y}_i(u_i^{\propto})] \leq u_i^{\propto}/2$. For this purpose, based on Lemma C.5, since the number of customers M is assumed to be IFR distributed, we know that $Y_i \sim B(M,\psi_i)$ follows an IFR distribution as well. As a result, by specializing Lemma C.6 with $\delta = 1/2$ and $C = u_i^{\propto}$, which is equivalent to assuming that $\mathbb{E}[\bar{Y}_i(u_i^{\propto})] \leq u_i^{\propto}/2$, we infer that $\mathbb{E}[\bar{Y}_i(u_i^{\propto})] \geq \mathbb{E}[Y_i]/2$. Therefore,

$$\mathbb{E}\left[\bar{Y}_i\left(u_i^{\propto}\right)\right] \quad \geq \quad \frac{1}{2} \cdot \min\left\{u_i^{\propto}, \mathbb{E}\left[Y_i\right]\right\} \ .$$

C.8 Proof of Lemma A.2

Let $\mathcal{H}^-$ be the set of heavy products whose selling price is less than $\epsilon^2 r_{i_{\max}}/(2n^2C)$, and $\mathcal{H}^+$ those with a selling price greater than $2n^2C \cdot r_{i_{\max}}/\epsilon^3$. Following the approach of Section 2.2, since the expected revenue function is subadditive (see Lemma 2.2), we have

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{L}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\tilde{\mathcal{H}}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^+}\right)\right] \geq \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{L}}\right)\right] + \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}}\right)\right] \geq \mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \ . \tag{37}$$

First, we observe that the contribution of any product $i \in \mathcal{H}^-$ toward the expected revenue of $U^*_{\mathcal{H}^-}$ is at most

$$\Pr\left[M \geq 1\right] \cdot C \cdot r_i \leq \Pr\left[M \geq 1\right] \cdot \frac{\epsilon^2 r_{i_{\max}}}{2n^2} \leq \frac{\epsilon}{n} \cdot \Pr\left[M \geq 1\right] \cdot \frac{r_{i_{\max}} w_{i_{\max}}}{1+w_{i_{\max}}} \leq \frac{\epsilon}{n} \cdot \mathbb{E}[\mathcal{R}(U^*)] \ ,$$

where the first inequality holds by definition of $\mathcal{H}^-$, and the second inequality holds since $i_{\max}$ is a heavy product. The last inequality is obtained by observing that the optimal expected revenue $\mathbb{E}[\mathcal{R}(U^*)]$ is lower bounded by the corresponding quantity with respect to the inventory vector that stocks a single unit of product $i_{\max}$ and nothing more, which is at least $\Pr[M \geq 1] \cdot r_{i_{\max}} w_{i_{\max}}/(1 + w_{i_{\max}})$. Consequently, by summing over all products $i \in \mathcal{H}^-$, we infer that

$$\mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}^-}\right)\right] \leq \epsilon \cdot \mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] \ .$$

Hence, when $\mathcal{H}^+ = \emptyset$, by inequality (37), it follows that $\mathbb{E}[\mathcal{R}(U^*_{\mathcal{L}})]+\mathbb{E}[\mathcal{R}(U^*_{\tilde{\mathcal{H}}})] \geq (1-\epsilon)\cdot\mathbb{E}[\mathcal{R}(U^*)]$.

In the opposite case, when $\mathcal{H}^+ \neq \emptyset$, consider some product $i \in \mathcal{H}^+$. As before, the optimal expected revenue $\mathbb{E}[\mathcal{R}(U^*)]$ is lower bounded by the expected revenue when stocking a single unit of product i, thus we obtain

$$\begin{aligned}
\mathbb{E}\left[\mathcal{R}\left(U^*\right)\right] &\geq \Pr\left[M \geq 1\right] \cdot \frac{r_i w_i}{1 + w_i} \\
&\geq \Pr\left[M \geq 1\right] \cdot \frac{nC}{\epsilon^2} \cdot r_{i_{\max}} \\
&\geq \Pr\left[M \geq 1\right] \cdot \frac{nC \cdot r_{i_1}}{\epsilon^2} \cdot \frac{w_{i_1}}{1 + w_{i_1}} \cdot \frac{1 + w_{i_{\max}}}{w_{i_{\max}}} \\
&\geq \Pr\left[M \geq 1\right] \cdot \frac{C \cdot r_{i_1}}{2\epsilon} \\
&\geq \frac{1}{2\epsilon} \cdot \mathbb{E}\left[\mathcal{R}\left(U^*_{\mathcal{H}}\right)\right] \ , \qquad (38)
\end{aligned}$$

where i_1 is the most expensive product stocked by $U^*_{\mathcal{H}}$. Here, the second inequality holds since $r_i \geq 2n^2 C \cdot r_{i_{\max}}/\epsilon^3$ and $w_i \geq \epsilon/n$, the third inequality follows by definition of $i_{\max}$ given that $r_{i_{\max}} w_{i_{\max}}/(1 + w_{i_{\max}}) \geq r_{i_1} w_{i_1}/(1 + w_{i_1})$, the fourth inequality holds since $w_{i_1} \geq \epsilon/n$, and the last inequality is due to the fact that r_{i_1} is the most expensive product on stock in $U^*_{\mathcal{H}}$. By combining inequality (37) with (38), we conclude that $\mathbb{E}[\mathcal{R}(U^*_{\mathcal{L}})] \geq (1 - 2\epsilon) \cdot \mathbb{E}[\mathcal{R}(U^*)]$.

C.9 Proof of Claim A.5

We first observe that using the formula of conditional expectation (relative to the value of M), we can restrict attention to a deterministic M. The desired inequality is proven inductively over C. For $C = 0$, we clearly have $\mathbb{E}[\bar{Y}] = \mathbb{E}[\bar{X}] = 0$.

For $C \geq 1$, by the induction hypothesis, $\bar{X}' = \min\{X, C-1\}$ and $\bar{Y}' = \min\{Y, C-1\}$ satisfy $\mathbb{E}[\bar{Y}'] \geq \theta \cdot \mathbb{E}[\bar{X}']$, and we wish to prove an analogous inequality between the expectations of $\bar{X} = \min\{X, C\}$ and $\bar{Y} = \min\{Y, C\}$. Each of the Binomial variables X and Y can be viewed as the terminating value of a Binomial process, counting the number of successes among M independent Bernoulli trials, with respective parameters α and $\theta\alpha$. We begin by defining the stopping time τ_X that corresponds to the first trial in which the Binomial process underlying the variable X, denoted by $X_1, \ldots X_M$, attains the value $C - 1$. If there are fewer than $C - 1$ successes among the M trials, then $\tau_X = M$. Next, observe that the expected value of X decomposes as follows:

$$\begin{aligned}
\mathbb{E}\left[\bar{X}\right] &= \mathbb{E}\left[\min\{X, C-1\} + \mathbb{I}\left[X > C-1\right]\right] \\
&= \mathbb{E}\left[\min\left\{X, C-1\right\}\right] + \Pr\left[X > C-1\right]
\end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}\left[\min\left\{X, C-1\right\}\right] + \sum_{\tau=0}^{M} \Pr\left[\tau_X = \tau\right] \cdot \Pr\left[X - X_\tau \geq 1 | \tau_X = \tau\right] \\
&= \mathbb{E}\left[\min\left\{X, C-1\right\}\right] + \sum_{\tau=0}^{M} \Pr\left[\tau_X = \tau\right] \cdot \Pr\left[X - X_\tau \geq 1\right] \\
&= \mathbb{E}\left[\min\left\{X, C-1\right\}\right] + \sum_{\tau=0}^{M} \Pr\left[\tau_X = \tau\right] \cdot \left(1 - (1-\alpha)^{M-\tau}\right) \ . \qquad (39)
\end{aligned}$$

The fourth equality follows from the independence of the Bernoulli trials, and the last equality holds since $X - X_\tau \sim B(M - \tau, \alpha)$. In an analogous way, τ_Y is defined as the first trial in which the Binomial process underlying the variable Y attains the value $C - 1$, with $\tau_Y = M$ when $Y < C - 1$. Based on the sequence of equations leading to (39),

$$\mathbb{E}\left[\bar{Y}\right] = \mathbb{E}\left[\min\left\{Y, C-1\right\}\right] + \sum_{\tau=0}^{M} \Pr\left[\tau_Y = \tau\right] \cdot \left(1 - (1-\theta\alpha)^{M-\tau}\right) \ . \qquad (40)$$

By the induction hypothesis, we already know that $\mathbb{E}[\min\{Y, C-1\}] \geq \theta \cdot \mathbb{E}[\min\{X, C-1\}]$. Thus, given (39) and (40) it remains to show that

$$\sum_{\tau=0}^{M} \Pr\left[\tau_Y = \tau\right] \cdot \left(1 - (1-\theta\alpha)^{M-\tau}\right) \geq \theta \cdot \sum_{\tau=0}^{M} \Pr\left[\tau_X = \tau\right] \cdot \left(1 - (1-\alpha)^{M-\tau}\right) \ . \qquad (41)$$

Note that since the function $\varphi_k : x \mapsto 1 - (1-x)^k$ is concave over the interval $[0, 1]$ for any $k \in \mathbb{N}$, we infer that $\varphi_k(\theta\alpha) \geq \theta \cdot \varphi_k(\alpha) + (1-\theta) \cdot \varphi_k(0) = \theta \cdot \varphi_k(\alpha)$, and therefore

$$1 - (1-\theta\alpha)^{M-\tau} \geq \theta \cdot \left(1 - (1-\alpha)^{M-\tau}\right) \ .$$

Hence, by observing that the right-hand side of the latter inequality is non-decreasing in τ, it is sufficient to prove that τ_X is stochastically smaller than τ_Y to derive the desired inequality (41). This property is easily derived by observing that the success parameter of the process $X_1, \ldots, X_M$ is lower-bounded by that of $Y_1, \ldots, Y_M$.

D Tested Heuristics

Local search. The algorithm iteratively improves the objective value, where in each step a single unit is transferred from one product to the other, until reaching a local minimum. Starting with an initial inventory vector, we iteratively implement the best swap between products, i.e., one that generates the largest incremental increase in the expected revenue, evaluated through our sampling-based oracle. Specifically, letting $U^{(k)}$ denote the inventory vector obtained at the beginning of step k, a swap is represented by an ordered pair of products (i, j), where the current inventory level $u_i^{(k)}$ of product i is strictly positive. The inventory vector $U_{i \to j}^{(k)}$ resulting from this swap is derived from $U^{(k)}$ through decreasing $u_i^{(k)}$ by one unit and augmenting $u_j^{(k)}$ by one unit. With this definition, we either proceed to step $k+1$ with the inventory vector $U_{i \to j}^{(k)}$ that maximizes $\mathbb{E}[\mathcal{R}(U_{i \to j}^{(k)})]$ over all swaps (i, j), or terminate the algorithm when none of these swaps improves the expected revenue by a factor greater than 1%. To alleviate the risk of 'bad starts', the vector $U^{(1)}$ is defined by initially stocking C units of the product that maximizes $r_i w_i$, similar to Goyal et al. (2016).

Gradient-descent approach. We consider a suitable adaptation of the stochastic gradient-descent algorithm of Mahajan and van Ryzin (2001) to the MNL-based dynamic assortment planning problem. In contrast to the latter paper, here the revenue function is defined only for integer-valued inventory vectors. Hence, similar to the approach of Goyal et al. (2016), we utilize a continuous relaxation of the revenue function, defined through the Lovász extension of a discrete function. Letting $f : \mathbb{Z}^n \to \mathbb{R}$ denote the expected revenue function, its Lovász extension $\hat{f} : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\hat{f}(U) = f(\lfloor U \rfloor) + \sum_{i=1}^{n} \left(u_{\pi(i)} - u_{\pi(i-1)} \right) \cdot \left[f\left(\lfloor U \rfloor + \sum_{k=1}^{i} e_{\pi(k)} \right) - f\left(\lfloor U \rfloor + \sum_{k=1}^{i-1} e_{\pi(k)} \right) \right] ,$$

where the permutation π sorts products by the increasing fractional part of their inventory, namely, $u_{\pi(1)} - \lfloor u_{\pi(1)} \rfloor \leq \cdots \leq u_{\pi(n)} - \lfloor u_{\pi(n)} \rfloor$. The Lovász extension is piecewise linear, and its gradient can be approximately computed using the sampling-based oracle given in Appendix A.1.

Starting with the initial solution $U^{(0)} = 0$, and letting $U^{(k)}$ denote the solution obtained at the end of step k, each iteration consists of computing $U^{(k+1)} = \max\{0, U^{(k)} + \epsilon \nabla f(U^{(k)})\}$, where ϵ is the step size. When the latter vector does not lie in the feasible region $\{U \in \mathbb{R}^n : \|U\|_1 \leq C\}$, it is projected onto the boundary by linear rescaling. Through trial and error, we picked a step size of $\epsilon_k = \max\{0.05 \cdot C, \frac{C - \|U_k\|_1}{2}\}$. The algorithm terminates when $U^{(k+1)}$ hits the boundary (i.e., $\|U^{(k+1)}\|_1 = C$) and the objective value does not improve by a factor greater than 0.5%. Since the gradient-descent algorithm is particularly slow, we force termination after 250 iterations. Finally, it remains to 'round' the resulting inventory vector to an integral one. Suppose that $U^{(k+1)}$ is the inventory vector obtained following the gradient-descent algorithm; then $\lfloor U^{(k+1)} \rfloor$ is augmented greedily, by stocking at each step a unit of the product with maximal marginal expected revenue, until reaching C units.

Dynamic programming. With some similarities to our setting, Topaloglu (2013) studied a joint assortment and inventory problem, where the demand is formed by a Poisson arrival process. However, the problem considered is incomparable to our setting, since his formulation does not take into account stock-out substitution effects. Instead of being governed by stock-outs, the assortment dynamics is at the discretion of the retailer, who can vary the offered assortment over time to better balance stocking constraints. Still, the algorithm devised by Topaloglu (2013) is a reasonable alternative to our approach, especially since the optimal policy in his model was proven to have a compact structure, being a mixture over at most n assortments under a Poisson demand process and a single assortment under a suitable normal approximation.

In the above-mentioned model, the problem formulation is given by:

$$\begin{aligned} \max_{U,y} \quad & \sum_{i \in [n]} \left(r_i \cdot \mathbb{E}\left[\min\left\{ U_i, \text{Poisson}\left(\mathbb{E}[M] \cdot \sum_{S: i \in S} y(S) \cdot \frac{w_i}{1 + w(S)} \right) \right\} \right] - c \cdot U_i \right) \\ \text{s.t.} \quad & \sum_{S \subseteq [n]} y(S) = 1 \end{aligned}$$

Here, U is the offered inventory vector, and for each possible assortment $S \subseteq [n]$ there is a corresponding decision variable $y(S)$ that describes its probability to be offered. In addition, the parameter c stands for the per-unit cost of any product. This parameter can be thought of as the Lagrangian multiplier associated with the cardinality constraint; in our setting, it can be determined through a bisection search. Now, since the objective function above is separable with respect to the products, one can cast this problem in dynamic programming terms. Specifically, we introduce the change of variable $\alpha_i = \sum_{S: i \in S} y(S) \cdot \frac{w_i}{1+w(S)}$, where α_i is the consumption rate of product i, and incorporate simple compatibility constraints between different products: $\alpha_0 + \sum_{i \in [n]} \alpha_i = 1$ and $\alpha_i \leq \frac{w_i}{w_0} \cdot \alpha_0$. At each step of the recursion, corresponding to some product $i \in [n]$, we approximately guess the consumption rate α_i, which immediately implies an optimal stocking level U_i to balance between marginal revenue and cost. We also implement the simplified recursion developed by Topaloglu (2013) under a normal approximation of the demand process. For a detailed description of these algorithms, we refer the reader to Sections 5 and 7 of his paper.

Deterministic relaxation. An additional approach that deals with stock-out substitution is the continuous-time deterministic relaxation developed by Honhon et al. (2010) and later on studied by Honhon and Seshadri (2013). Here, the stochastic nature of the choice process is overlooked. Given the initial inventory vector U and its corresponding assortment S, one assumes that each product $i \in S$ is consumed at a constant rate of $\alpha_i = w_i/(1 + w(S))$, until one of the products in S is depleted. Specifically, the first stock-out occurs at time $\min_{i \in S}(U_i/\alpha_i)$. Similarly, at the beginning of each subsequent epoch, the consumption rates are updated to reflect the changes of assortment, and the current epoch terminates at the next stock-out event. In this setting, the total consumption of products is indeed deterministic with respect to the initial stocking decisions. To optimize the latter, Honhon et al. (2010) devised a dynamic programming approach that exploits the special structure of epochs and runs in time $O(8^n)$. Due to the exponential dependency on the number of products, this approach is not applicable in our experimental setting, with $n = 20$ products. Instead, we cast the resulting deterministic model as a mixed integer program and use a state-of-the-art commercial solver (Gurobi Optimization 2015). To obtain faster convergence, the solver is given access to a warm-start solution, using the same initial inventory vector as the local search heuristic described earlier. In most cases, the solver indeed returns close-to-optimal solutions (to the relaxation) within the allotted time limit of 2000 seconds. This benchmark is informative from a modeling perspective, since it sheds light on the relative merits of using a deterministic demand process rather than the actual stochastic one.

Discrete-greedy. The discrete-greedy algorithm starts with zero inventory levels for all products, and iteratively augments the current inventory vector by a single unit of the product that incurs the largest increase in the expected revenue, until reaching C units. The expected revenue is evaluated using our sampling-based procedure. It is worth mentioning that this approach is the closest in spirit to the way our algorithm operates on heavy-expensive products, where a restricted-non-decreasing and restricted-submodular set function is approximately maximized through a greedy procedure (see Section 3.2).

References

Aouad, A., Levi, R. and Segev, D. (2015), Approximation algorithms for dynamic assortment optimization models. Working paper, available online as SSRN report #2641268.

Ben-Akiva, M. E. and Lerman, S. R. (1985), *Discrete Choice Analysis: Theory and Application to Travel Demand*, MIT Press.

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

Bront, J. J. M., Méndez-Díaz, I. and Vulcano, G. (2009), 'A column generation algorithm for choice-based network revenue management', *Operations Research* **57**(3), 769–784.

Chandukala, S. R., Kim, J., Allenby, G. M. and Otter, T. (2008), *Choice Models in Marketing: Economic Assumptions, Challenges and Trends*, Now Publishers Inc.

Davis, J., Gallego, G. and Topaloglu, H. (2013), Assortment planning under the Multinomial Logit model with totally unimodular constraint structures. Working paper.

Désir, A., Goyal, V., Segev, D. and Ye, C. (2015), Capacity constrained assortment optimization under the Markov Chain based choice model. Working paper, available online as SSRN report #2626484.

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

Feldman, J. and Topaloglu, H. (2017), 'Revenue management under the Markov chain choice model', Operations Research, forthcoming.

Ford, L. R. (1957), 'Solution of a ranking problem from binary comparisons', *American Mathematical Monthly* **64**(8), 28–33.

Gaur, V. and Honhon, D. (2006), 'Assortment planning and inventory decisions under a locational choice model', *Management Science* **52**(10), 1528–1543.

Goyal, V., Levi, R. and Segev, D. (2016), 'Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand', *Operations Research* **64**(1), 219–235.

Grover, R. and Vriens, M. (2006), *The Handbook of Marketing Research: Uses, Misuses, and Future Advances*, Sage Publications.

Guadagni, P. M. and Little, J. D. (1983), 'A Logit model of brand choice calibrated on scanner data', *Marketing Science* **2**(3), 203–238.

Gurobi Optimization, I. (2015), 'Gurobi optimizer reference manual'.
URL: *http://www.gurobi.com*

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

Honhon, D. and Seshadri, S. (2013), 'Fixed vs. random proportions demand models for the assortment planning problem under stockout-based substitution', *Manufacturing & Service Operations Management* **15**(3), 378–386.

IHL (2015), Retailers and the ghost economy: $1.75 trillion reasons to be afraid, Technical report, IHL Group.

Luce, R. D. (1959), *Individual Choice Behavior a Theoretical Analysis*, John Wiley & Sons.

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

Maystre, L. and Grossglauser, M. (2015), Fast and accurate inference of Plackett–Luce models, *in* 'Advances in Neural Information Processing Systems', pp. 172–180.

McFadden, D. (1973), 'Conditional Logit analysis of qualitative choice behavior', *Frontiers in Econometrics* pp. 105–142.

McFadden, D. (1980), 'Econometric models for probabilistic choice among products', *Journal of Business* **53**(3), S13–S29.

Megiddo, N. (1979), 'Combinatorial optimization with rational objective functions', *Mathematics of Operations Research* **4**(4), 414–424.

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

Muckstadt, J. A. and Sapra, A. (2010), *Principles of Inventory Management: When You Are Down to Four, Order More*, Springer Science & Business Media.

Nagarajan, M. and Rajagopalan, S. (2008), 'Inventory models for substitutable products: optimal policies and heuristics', *Management Science* **54**(8), 1453–1466.

Negahban, S., Oh, S. and Shah, D. (2012), Iterative ranking from pair-wise comparisons, *in* 'Advances in Neural Information Processing Systems', pp. 2474–2482.

Nemhauser, G. L., Wolsey, L. A. and Fisher, M. L. (1978), 'An analysis of approximations for maximizing submodular set functions – I', *Mathematical Programming* **14**(1), 265–294.

Plackett, R. L. (1975), 'The analysis of permutations', *Applied Statistics* **24**(2), 193–202.

Rusmevichientong, P., Shen, Z.-J. M. and Shmoys, D. B. (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. and Topaloglu, H. (2014), 'Assortment optimization under the Multinomial Logit model with random choice parameters', *Production and Operations Management* **23**(11), 2023–2039.

Rusmevichientong, P. and Topaloglu, H. (2012), 'Robust assortment optimization in revenue management under the Multinomial Logit choice model', *Operations Research* **60**(4), 865–882.

Segev, D. (2015), Assortment planning with nested preferences: Dynamic programming with distributions as states? Working paper, available as SSRN report #2587440.

Shaked, M. and Shanthikumar, J. (1994), *Stochastic Orders and Their Applications*, Academic Press, New York.

Smith, S. A. and Agrawal, N. (2000), 'Management of multi-item retail inventory systems with demand substitution', *Operations Research* **48**(1), 50–64.

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

Topaloglu, H. (2013), 'Joint stocking and product offer decisions under the Multinomial Logit model', *Production and Operations Management* **22**(5), 1182–1199.