# Simplifying Revenue Management

Harsh Tarak Sheth

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2024

# Abstract

Simplifying Revenue Management

Harsh Tarak Sheth

In this thesis, we study three revenue management problems where we propose simple algorithms with provable guarantees. While online marketplaces provide retailers with tremendous flexibility, they are often large, noisy, have multiple stakeholders, and could be more challenging to characterize. These complexities give rise to a preference for simple, interpretable policies. Further, traditional marketplaces such as brick-and-mortar stores cannot always leverage tools designed for online environments due to physical constraints, higher latency, etc. With these motivations in mind, we develop algorithms for assortment optimization and pricing that are easy to implement in practice and have theoretical justifications for their performance.

In Chapter 1, we consider a dynamic assortment optimization problem where the seller has a fixed inventory of multiple substitutable products to sell over a fixed time horizon. We consider two modifications to the traditional problem. First, we simplify the assortment planning by restricting assortment changes to "product retirements". When a product is retired, it becomes unavailable to all future customers. Second, we assume the seller has flexibility regarding which customers to approach. In each period, the seller chooses which subset of products to retire and selects a customer to visit. The selected customer then receives an option to purchase one of the available products, i.e., non-retired products with positive remaining inventory. We provide two policies for this problem. Our first policy guarantees a constant fraction of the best possible revenue. Our second policy is near-optimal but requires the problem to have a specific structure.

In Chapter 2, we study the fundamental joint pricing and inventory management problem. The optimal policy for the model we consider is known to be an $(s, S, \boldsymbol{p})$ policy: when the inventory level drops to $s$ units, the seller immediately places an order to replenish the inventory to $S$ units. Specifically, the optimal pricing policy $\boldsymbol{p}$ has a different price for every inventory state. We proposed simple policies requiring no more than three prices and prove that these policies are near-optimal compared to optimal policies which require more prices and are less robust. In particular, when orders cannot be backlogged, we show that a single price is sufficient for good performance.

In Chapter 3, we analyze assortment optimization and pricing with opaque products. An opaque product is one for which only partial information is available to the buyer at the time of purchase. When a customer selects the opaque product, the seller can fulfill the purchase using any of the offered products. Opaque products can help sellers boost total sales. We propose simple policies for assortment optimization with provable constant factor guarantees, which are near-optimal in numerical experiments. We also provide upper bounds for the advantage of selling opaque products.

# Table of Contents

# List of Figures

# List of Tables

# Acknowledgements

I would like to start by thanking my parents and grandparents for their love, support and trust through the years. My advisors, Adam and Vineet, for their feedback and patience. My collaborators, Yeqing, Omar and Jiaqi, for indulging my half-baked ideas. First year course instructors, Yuri, Karl and Ton, for giving me the confidence and skills to say I got this. Manan, for understanding me like no other. All the friends I made in the department: Jacob, Prusty, Apurv, Goutam, Raghav, Mali, Rachitesh, Sudeep, Ruizhe, Lingyi, Oussama, Michael, Akshit, Priyank, Madhu, Darshan, Matias, Abdellah, Aapeli, who were sometimes the reason I came to office, and always the reason I stayed. Liz, Winsor, Kristen and everyone else on the staff, for always bearing with my last-minute requests. Yashaswini, Yug for giving me company through my late nights. Yashraj, Nihar, Sanat, Vishwa, Aakriti, Keni, Damani, and Kritika, for the amazing vacations, birthdays, and pandemic games. Parth, for the daily Instagram content. Mitali, for making the last year fun, my sister, Manasi, for being the grown up in the room. And finally Xuan, for always giving me more than I could have asked for, and sharing Boba. It is not a stretch to say that this document would have been impossible without the love and support of all these people.

## Dedication

To Mummy, Papa, and Manasi.

# Introduction

Revenue management and pricing problems study what products to offer, when, to which customer, and at what prices. The field took shape in the 1980s with the airline industry when American Airlines segmented customers into two fare classes, successfully fending off competitors offering discounted fares [1, 2]. It has since blossomed into an area of specialty that is useful across many industries, including but not limited to online retail, ride-sharing, hospitality, and healthcare. Since the digital era has made it possible to collect and store data on massive scales, it has opened up doorways to model and capture consumer behavior in a way that was not possible earlier. Businesses have increasingly used this newfound understanding of their customers' preferences to improve profit margins.

The main goal of the topics I research in this thesis is to make revenue management tools more accessible for all businesses and enable widespread adoption. In particular, my research focuses on the design and analysis of *simple* policies for revenue management. Online marketplaces provide tremendous flexibility to retailers to make customer experience more scalable, personalizable, and dynamic. However, online marketplaces are often large, complex, and noisy, and what might be "optimal" on paper might not be feasible in practice. Accordingly, this thesis considers relaxing the optimality requirement in favor of ease of implementation and robustness. Moreover, the policies considered in this document are also more accessible for traditional marketplaces such as brick-and-mortar stores, which cannot always leverage tools designed for online environments due to physical constraints, higher latency, etc. My work focuses on two crucial revenue management settings: assortment optimization and pricing. In each chapter, we introduce a policy set that is easy to implement and then provide theoretical and practical justifications for their performance.

In Chapter 1, we consider a multi-period, multi-product revenue management problem where the seller has a fixed inventory of multiple substitutable products to sell over a fixed time horizon. We study the dynamic assortment optimization problem when changes in the assortment are limited to product retirements. In each period, the seller chooses which subset of products to retire and selects a customer to visit. When a product is retired, it becomes unavailable to all future customers. When a customer is selected, they are offered all available products, i.e., non-retired products with positive remaining inventory. The objective of the seller is to dynamically retire products and select customers to maximize the total expected revenue over a fixed time horizon. Such product retirement decisions are essential when the seller cannot personalize the set of products offered to each customer.

For this problem, we show that the objective can be bounded above with a linear program. Using this linear program, we design a policy that guarantees a constant fraction of the optimal dynamic retirement-selection policy. Further, given an optimal solution to the linear program with specific structure, we give an asymptotically optimal policy. We show that such a solution can always be found when there are only two products or only one customer type. Finally, we show that our policies perform well in numerical experiments calibrated to models from real-world experiments when compared to natural benchmarks.

In Chapter 2, we study a fundamental joint pricing and inventory management problem. We consider an infinite horizon, continuous review, stochastic inventory system. Demand arrives according to a Poisson process, where the customer makes a purchase if their valuation exceeds the current price. The seller can replenish inventory for a fixed cost and also incurs holding and backlogging costs. The optimal policy to maximize profit under this setting is an $(s, S, \boldsymbol{p})$ policy: when the inventory level drops to $s$ units, the seller instantaneously orders to replenish the inventory to $S$ units. This optimal pricing policy $\boldsymbol{p}$ has a different price for every inventory state. However, such inventory-based pricing policies are challenging to implement for sellers and may lead to strategic customer behavior or perceptions of unfairness. In this work, we prove that simple policies where only a small number of prices are offered are near-optimal compared to the optimal policy, which requires $S - s$ prices.

We prove theoretical guarantees for our proposed simple policies when customer valuations follow a monotone hazard rate (MHR) distribution. When no backlogging is allowed (lost sales model), we show that there exists a static, single-price policy that garners at least as much revenue as the optimal dynamic policy while incurring at most $\sqrt{1 + \ln(S)}$ times the cost of the optimal dynamic policy. Furthermore, for the classic uniform valuation (linear demand) model, we improve the cost ratio to 1.225. When backlogging is allowed, we show that a three-price policy achieves similar guarantees. In numerical experiments, we show that our simple policies capture a significant fraction of the profit in most cases, with the performance improving as the profit margin increases.

Finally, in Chapter 3, we consider assortment optimization and pricing with opaque products. An opaque product is one for which only partial information is available to the buyer at the time of purchase. We consider a seller offering substitutable products, which are similar but differentiated in some features. A standard example here would be the set of all flights with the same origin and destination on a given day. When a customer arrives, the seller chooses to offer a subset of these products to the customer. The seller also offers the customer the option of purchasing an opaque product. If the customer chooses to purchase the opaque product, the seller can fulfill the purchase with any of the offered products of the seller's choosing. Such opaque selling mechanisms help sellers boost sales and give them more flexibility for inventory management while rewarding customers with a discount for their flexibility. We study how adding an opaque option modifies the assortment optimization problem for the multinomial logit choice model.

For our results, we assume that customers are risk-averse when evaluating the opaque option and value it as much as the lowest valuation product in the assortment. We find that the nested-by-revenue heuristic is no longer optimal with an opaque product added to the assortment. Instead, we show that offering the better of a nested-by-revenue assortment and the best singleton assortment guarantees a half approximation. Furthermore, a nested-by-revenue-and-valuation heuristic has the same guarantees and is near-optimal in numerical experiments. When considering the joint pricing and assortment optimization problem, we find that offering the opaque product does not improve revenue, which we prove when there are two products or when pricing is uniform.

# Chapter 1: Revenue Management with Product Retirement and Customer Selection

## 1.1 Introduction

Revenue management with customer choice focuses on maximizing revenue by influencing the purchase decisions of customers via the assortment of products that is offered [1]. The problem of *dynamic assortment optimization* has recently gained much traction as it arises in many applications, including online retail and recommendations [3, 4]. In the classical setting of dynamic assortment optimization, the seller has the full flexibility of personalizing the assortments for every customer as they sequentially arrive to the platform. Although powerful for online marketplaces, such dynamic policies are not suitable or even feasible for many other settings such as brick-and-mortar stores (e.g., fast-fashion stores, department stores), electronic marketplaces (e.g., iPhones), and business-to-business (B2B) settings. In the former case, all customers entering the store see the same products; and for the latter two, available products offered by the sellers are common knowledge and can often be obtained from the firm's webpage. Hence, in all cases, assortment personalization is not an option for revenue maximization. Motivated by such real-world contexts, in this work, we study two novel practices that are implementable for the aforementioned settings: *product retirement* and *customer selection*. We provide policies to implement these practices with provable performance guarantees.

Product retirements are a simple way to control the set of available products, where once a product is retired, it becomes unavailable for purchase for all future customers. They arise naturally in many settings: for brick-and-mortar stores, retiring a product can be interpreted as either sending the item to a lower-tier department store, or moving the item to the clearance section; for electronics and B2B sellers, product retirement can be interpreted as discontinuing older services, models, or generations of the products. Optimally retiring products is essential to revenue maximization because it helps balance the following trade-off: retiring a product too early can lead to loss of demand, whereas retiring too late might cannibalize the sales of newer and more profitable products.

In some settings, product retirements can be coupled with customer selection to enhance their impact. For instance, in many B2B settings, sellers select potential customers every day, and sales representatives then offer these customers all the available products via a sales call [5]. In such settings, the sellers can jointly make decisions on when to retire older generations and which customers to approach each day to balance the trade-off of customer preference for older products with the higher revenue of newer products. In fact, this chapter is a direct result of work done with a B2B hardware provider who operates in exactly this manner, and we use their proprietary data to evaluate our algorithms.

Formally, we study the problem where a seller has a fixed inventory of multiple substitutable products to be sold over a finite time horizon. Customers are differentiated by types, where each customer type chooses from the available products according to the *multinomial logit* (MNL) choice model. The seller can decide to retire a product at any time, after which it can no longer be offered to any future customers. The seller also dynamically decides on the sequence in which to visit the customers. The seller cannot optimize the assortment shown to the visited customer. A visited customer is shown all available products, i.e., all products that are not retired and have inventory remaining. The goal of the seller is to maximize the total expected revenue.

Our main contributions are the following:

1. We construct a policy that achieves at least $1/4 - o(1)$ of the expected revenue of the optimal dynamic policy. Our algorithm determines a static sequence of the customers and assortments by constructing an approximate greedy solution to a linear program, whose objective value is an upper bound for the total expected revenue. Our analysis is tight as the inventory grows large.

2. When there is an optimal solution to the LP upper bound with a particular nested structure,

we construct a policy that mimics the optimal LP solution. We show that as the time horizon and inventories of products grow large, our policy achieves asymptotic optimality. We also show this nested structure always exists and can be found efficiently when there are two products or one customer type.

3. We conduct numerical experiments using a proprietary B2B dataset where multiple technology products are sold to multiple customer types. We observe that our proposed policies outperform benchmarks under most scenarios. We also observe that the particular nested structure required for an asymptotically optimal policy can be found in an overwhelming majority of cases, and for these cases, our policy results in the best performance.

We note that our theoretical guarantees also apply to the general stochastic dynamic assortment optimization problem with customer selection, since the performance guarantees are with respect to an LP upper bound that does not account for product retirement constraints and can offer an arbitrary assortment to any customer.

### 1.1.1 Literature review

Our work falls under the larger umbrella of dynamic assortment optimization. Gallego et al. [6] and Liu and Van Ryzin [7] showed that the optimal revenue for the choice-based network revenue management problem can be approximated by a deterministic linear program. This LP benchmark, which is a fluid relaxation of the original problem, has become a popular tool to analyze algorithm performance for dynamic assortment optimization problems. When this fluid relaxation LP can be solved, the optimal solution can be used to construct an asymptotically optimal policy for the dynamic assortment optimization problem. Ma et al. [8] design static policies using the optimal LP solution and show provable guarantees even in the non-asymptotic setting. The LP in general has exponentially many decision variables and requires approximation schemes to solve. However, Gallego et al. [9] showed that the LP can be simplified to a sales-based version for a general class of choice models that includes the MNL choice model as a special case. Feldman and Topaloglu [10] similarly show the existence of a simplified LP for the Markov chain choice model. Both works also show that the optimal solution to their respective simplified LPs has the property that the assortments offered by the optimal solution can be ordered such that one assortment is nested in another. This structural property is useful when thinking of assortment decisions in terms of product retirements, and we use it when developing our asymptotically optimal policy. In our work, we use the sales-based version of this LP benchmark [9] for the MNL choice model, modifying it to accommodate customer selection. Liang et al. [11] study the structure of the sales-based LP for the MNL choice model to propose policies for joint inventory and assortment planning. Our work differs from the ones mentioned above due to the unique challenges introduced by limiting ourselves to product retirements in the presence of multiple customer types.

Our work is the first to consider product retirements in the context of revenue management, however, several other modifications and extensions to the dynamic assortment optimization problem have been considered. Davis et al. [12] consider a problem where products are *introduced* over time. While in our work the primary constraints are finite inventory and product retirement, the primary constraint in Davis et al. [12] is that at most one product can be introduced in each period. Other extensions considered in the literature include add-ons [13], multiple stages [14, 15] and reusable resources [16, 17, 18, 19].

While the customer arrivals in our model are endogenous and controlled by the seller, a majority of the literature focuses on exogenous customer arrivals. One popular framework is that of stochastic customer arrivals where it is assumed that some distributional knowledge about customer types is known beforehand. For stochastic customer arrivals, Bernstein et al. [4] characterized the optimal policy for the setting with two identically priced products and two customer types. Further, there are several works that use optimal solutions of the LP benchmark to construct policies that are asymptotically optimal [7, 20, 3, 8].

The other popular customer arrival framework is that of adversarial customer arrivals. These results analyze algorithms with respect to the worst-case customer arrival sequence under the competitive ratio framework. Chan and Farias [21] showed that myopic policies are $1/2$-competitive for stochastic depletion problems, which includes dynamic assortment optimization. Golrezaei et al. [3] proposed inventory-balancing algorithms that are $(1 - 1/e)$-competitive for a large class of customer choice models, and also showed that no policy could do better. Ma and Simchi-Levi [22] extended the inventory-balancing algorithm to a setting that allows different customer types to have different prices for each product and showed a price-dependent competitive ratio.

While several works of literature have studied stochastic and adversarial customer arrivals, few works in revenue management literature have studied the problem of customer selection. One work to do so is from Cominetti et al. [23], who propose algorithms to select the optimal subset of customers for offering discounts. Outside revenue management, the customer selection aspect of our setting is related to the selection problem that is dealt with in context of the prize-collecting traveling salesman problem [24, 25], the prize-collecting Steiner tree problem [26, 27], facility location with penalties [28, 29, 30], and supply chain planning with market choice [31]. This set of problems do offline customer selection, as the set of potential customers (or nodes) is available beforehand. This is in contrast to supply chain management with online customer selection [32, 33], where customer selection decisions are made without knowledge of the future customer arrivals. In our model, while the set of customers is known beforehand, it may be suboptimal to perform the customer selection step offline due to the stochasticity of the customer purchase decisions, which makes the customer selection aspect of our setting different from both lines of work described above.

### 1.1.2 Organization

We describe the model and LP upper bound in Section 1.2. In Section 1.3, we describe our policy which gives a constant factor approximation. In Section 1.4, we describe our policy which is asymptotically optimal when an optimal LP solution with a certain structure is available. In Section 1.5, we compare our policies' performance against some natural benchmarks using numerical experiments based on real data from a B2B setting.

## 1.2 Model

We consider a seller with a finite inventory of substitutable products which are sold to a sequence of customers over a finite time horizon. The seller has access to a pool of customers, possibly of different types, and their purchase behavior is governed by known, type-specific choice models. In every period, the seller decides which type of customer to visit, and all products that have not been retired and have inventory remaining are offered to that customer. The seller may decide to retire a product at any time, at which point the product will no longer be available to future customers. The goal is to design near-optimal policies for product retirement and customer selection that maximize the seller's expected revenue.

We denote the set of products by $\mathcal{N} := \{1, \ldots, n\}$ and the number of selling periods by $T$. Each product $i$ has a revenue $r_i$ and finite, non-replenishable inventory (capacity) $c_i$. Let $c_{\min} := \min_{i \in \mathcal{N}} c_i$. We denote the set of customer types by $\mathcal{L} := \{1, \ldots, m\}$, and let $b_j$ be the number of type-$j$ customers that the seller has access to. For each type $j$, we let $\pi_i^j(S)$ denote the probability that a type-$j$ customer purchases product $i$ when the set of available products is $S$. The customer can also choose to not purchase at all, which we may refer to as the no-purchase option or product 0 for convenience. Thus, $\pi_0^j(S) = 1 - \sum_{i \in S} \pi_i^j(S)$ for all $j \in \mathcal{L}$ and $S \subseteq \mathcal{N}$. Any product that is not available cannot be purchased by a customer, i.e., $\pi_i^j(S) = 0$ if $i \notin S \cup \{0\}$. In this work, we focus on the MNL choice model and provide additional results on the independent demand model in Appendix A.5. We use $R_\pi(S)$ to denote the expected revenue from a customer with choice model $\boldsymbol{\pi} := \{\pi_i(S) : i \in \mathcal{N},\ S \subseteq \mathcal{N}\}$ when offered $S$, i.e.,

$$R_\pi(S) := \sum_{i \in S} r_i \pi_i(S). \tag{1.1}$$

We now describe the sequence of events for any policy $\mathcal{A}$. In period 0, the seller selects the products they will be selling over the selling season, and retires the rest. In each period $t \in [T]$, let $S_t(\mathcal{A})$ denote the assortment of *available* products – products with positive remaining inventory that have not been retired. The seller first selects a new customer, $J_t(\mathcal{A})$, from the remaining pool of customers. The customer then purchases product $i$ among $S_t(\mathcal{A}) \cup \{0\}$ with probability $\pi_i^{J_t(\mathcal{A})}(S_t(\mathcal{A}))$. If product $i$ runs out of inventory at this point, it leaves the assortment. Further, the seller decides which products, if any, to retire before selecting the next customer. In other words, $S_{t+1}(\mathcal{A}) \subseteq S_t(\mathcal{A})$, where $S_t(\mathcal{A}) \setminus S_{t+1}(\mathcal{A})$ is the set of products that either ran out of inventory or were retired at the end of period $t$. The last period that a product was offered in is referred to as its retirement time, which is denoted as $d_i$ for product $i$. We denote the product choice of customer $J_t(\mathcal{A})$ by $X_t(\mathcal{A})$. The total sales of product $i$ and total number of type-$j$ customers selected are denoted by

$$Z_i(\mathcal{A}) := \sum_{t=1}^{T} \mathbb{I}[X_t(\mathcal{A}) = i] \text{ and } Y_j(\mathcal{A}) := \sum_{t=1}^{T} \mathbb{I}[J_t(\mathcal{A}) = j],$$

respectively. The objective is to find a product retirement and customer selection policy $\mathcal{A}$ that maximizes expected revenue, i.e.,

$$\textsc{Rev}(\mathcal{A}) := \mathbb{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{n} r_i \pi_i^{J_t(\mathcal{A})}(S_t(\mathcal{A}))\right] = \sum_{i=1}^{n} r_i \mathbb{E}[Z_i(\mathcal{A})], \tag{1.2}$$

under demand and supply constraints. Here, by demand constraints we mean that the policy cannot select more type-$j$ customers that there are available, i.e., $Y_j(\mathcal{A}) \leq b_j$; whereas supply constraints refer to the constraint that the total sales of product $i$ cannot exceed its initial inventory, i.e., $Z_i(\mathcal{A}) \leq c_i$.

### 1.2.1 Multinomial logit choice model

We assume that purchase probabilities for each customer type are governed by a *multinomial logit* (MNL) choice model. The MNL choice model [34, 35, 36] is one of the most widely used customer choice models, both in theory and practice. Under the MNL model, the probability of a type-$j$ customer purchasing $i$ among the set of available products $S$ is,

$$\pi_i^j(S) = \frac{u_i^j}{1 + \sum_{i' \in S} u_{i'}^j}, \tag{1.3}$$

where $u_i^j$ is known as the attractiveness parameter of product $i$ for customer type $j$. The attractiveness parameter for the no-purchase option, $u_0^j$, is normalized to 1. We note some important properties of the MNL choice model, which we will use repeatedly throughout this chapter. First, the MNL choice model has the *substitutability property*, i.e., for any MNL choice model $\boldsymbol{\pi}$, assortment of products $S$ and products $i, j \in \mathcal{N}$ such that $i \in S$ and $j \notin S$, $\pi_i(S) \geq \pi_i(S \cup \{j\})$. Also, the product sets to consider when computing the unconstrained revenue-maximizing assortment for (1.1) are *nested-by-revenue* [37], meaning that the revenue-maximizing assortment would be all products with revenue above a certain threshold.

### 1.2.2 Upper bound

We now present a dynamic assortment and customer selection based upper bound for the expected revenue that can be achieved by any policy. Specifically, we consider the following sales based LP [9]:

$$
\begin{aligned}
z^*_{SLP} := \max_{q_i^j} \quad & \sum_{j=1}^{m} \sum_{i=1}^{n} r_i q_i^j \\
\text{s.t.} \quad & \sum_{j=1}^{m} \left( \sum_{i=1}^{n} q_i^j + q_0^j \right) \leq T, \\
& \sum_{i=1}^{n} q_i^j + q_0^j \leq b_j && \forall j \in \mathcal{L}, \\
& \sum_{j=1}^{m} q_i^j \leq c_i && \forall i \in \mathcal{N}, \\
& q_i^j \leq u_i^j q_0^j && \forall i \in \mathcal{N}, j \in \mathcal{L}, \\
& q_i^j \geq 0 && \forall i \in \mathcal{N}, j \in \mathcal{L}.
\end{aligned}
\tag{SLP}
$$

Here, $q_i^j$ represents the total expected sales of product $i$ to type-$j$ customers under the LP relaxation. The first three constraints are equivalent to the time horizon, demand, and supply constraints respectively. For any set of products $S$, the MNL choice probability of purchasing product $i$ can be described as $\pi_i(S) = \mathbb{I}[i \in S] u_i \pi_0(S)$, which implies that the expected sales of product $i$ under any policy can be at most $u_i$ times the expected number of no-purchases. The fourth constraint follows directly from this observation. Note that the above LP does not include any product retirement constraints and, in particular, allows the assortments to change arbitrarily. Therefore, it in fact serves as a deterministic fluid relaxation for the dynamic assortment optimization problem with customer selection under the MNL choice model, and is an upper bound as stated in Lemma 1.1.We provide a proof for Lemma 1.1 in Appendix A.1.

**Lemma 1.1** (LP upper bound with customer selection [9])**.** *Suppose all customers make purchase decisions according to a type-specific MNL model. For any product retirement and customer selection policy* $\mathcal{A}$*, the expected revenue* $\text{REV}(\mathcal{A})$ *is at most* $z^*_{SLP}$.

In Sections 1.3 and 1.4, we cover two policies for product retirement and customer selection with provable performance guarantees for MNL customers. Both our policies are non-adaptive and rely on constructing solutions to (SLP). We also separately consider the case where customer demands for products are independent of the assortment offered (i.e., the independent demand choice model) in Appendix A.5. When the demands for products are independent, it is optimal to never retire a product, and only customer selection decisions need to be made. We show that an adaptive myopic policy guarantees a $1 - 1/e$ fraction of the optimal expected revenue under the independent demand model.

## 1.3 Constant Factor Policy

In this section, we present policy A1 that gives a constant factor approximation. The main idea of the policy is to construct a good feasible solution of (SLP) by selecting customers one at a time, and then sequencing the customers to accommodate product retirements.

### 1.3.1 Structure of policy A1

The policy has two main steps. The first step is customer selection, where we select customers and corresponding assortments to construct a feasible solution to the LP upper bound (SLP). The second step determines the sequence of customers and retirement times for the products. In Theorem 1.2, we show that this policy guarantees at least $1/4 - o(1)$ fraction of the optimal expected revenue. The details of A1 are provided in Algorithm 1. Moreover, the analysis of our algorithm is tight as $c_{\min}$ goes to infinity. The tight example is deferred to Appendix A.2.4.

**Theorem 1.2** (Approximation guarantee for A1)**.** *Recall* $c_{\min} := \min_i c_i$. *Then,*

$$\textsc{Rev}(\text{A1}) \geq \left(\frac{c_{\min} - 1}{4c_{\min} - 2}\right) z^*_{SLP}.$$

*Moreover, this analysis is tight as* $c_{\min} \to \infty$.

*Customer selection (Step 1 of Algorithm 1).* The policy selects a subset of the customers based on a greedy solution of the (SLP). To construct this solution, in each round for a total of $T$ rounds, the policy selects the customer and corresponding assortment from the remaining products and customer types which gives the maximum expected revenue. Once the selection is done, the available customer and product inventory is updated in a fluid sense. Specifically, suppose $(\widetilde{J}_t, \widetilde{S}_t)$ are the customer type and assortment selected in round $t$, then the remaining inventory for product $i$ is decreased by $\pi_i^{\widetilde{J}_t}(\widetilde{S}_t)$, and the number type-$J_t$ customers available is reduced by 1 unit. A customer type is no longer selected when no more customers of that type remain, whereas a product is no longer offered when its expected consumption exceeds $c_i - 1$, i.e., less than one unit of inventory remains. This guarantees that the total consumption of customer types and products satisfy the demand and supply constraints. Let $\widetilde{q}_i^j$ denote the total consumption of product $i$ by type-$j$ customers, i.e.,

$$\widetilde{q}_i^j := \sum_{t=1}^{T} \mathbb{I}[\widetilde{J}_t = j] \pi_i^{\widetilde{J}_t}(\widetilde{S}_t). \tag{1.4}$$

We show in Lemma 1.3 below that this greedy solution will have an objective value that is at least a $1/2 - o(1)$ fraction of the optimal value $z^*_{SLP}$. The proof idea for the approximation guarantee is to use dual-fitting. A detailed proof is provided in Appendix A.2.2.

**Lemma 1.3.** *The solution to* (SLP) *constructed in the customer selection step (described in* (1.4)*)*

**Algorithm 1** A1 Policy

```
// Step 1: Customer selection
initialize $\widetilde{Y}_j \leftarrow 0\ \forall j \in \mathcal{L}, \widetilde{Z}_i \leftarrow 0\ \forall i \in \mathcal{N}, \mathcal{S} \leftarrow \emptyset$
for $t = 1$ to $T$ do
    $\widetilde{\mathcal{N}}_t \leftarrow \{i \in \mathcal{N} \mid \widetilde{Z}_i \leq c_i - 1\}, \widetilde{\mathcal{L}}_t \leftarrow \{j \in \mathcal{L} \mid \widetilde{Y}_j \leq b_j - 1\}$
    $\widetilde{R}_t \leftarrow \max_{j \in \widetilde{\mathcal{L}}_t, S \subseteq \widetilde{\mathcal{N}}_t} R_{\pi^j}(S)$
    $(\widetilde{J}_t, \widetilde{S}_t) \leftarrow \arg\max_{j \in \widetilde{\mathcal{L}}_t, S \subseteq \widetilde{\mathcal{N}}_t} R_{\pi^j}(S)$          // ties broken arbitrarily
    $\widetilde{Y}_{\widetilde{J}_t} \leftarrow \widetilde{Y}_{\widetilde{J}_t} + 1, \widetilde{Z}_i \leftarrow \widetilde{Z}_i + \pi_i^{\widetilde{J}_t}(\widetilde{S}_t) \quad \forall i \in \widetilde{S}_t$
    $\mathcal{S} \leftarrow \mathcal{S} \cup (\widetilde{J}_t, \widetilde{R}_t)$
end for
// Step 2: Construct customer sequence and retirement times
for $t = 1$ to $T$ do
    $(J_t, R_t) \leftarrow \arg\min_{(j,R) \in \mathcal{S}} R$          // ties broken arbitrarily
    $\mathcal{S} \leftarrow \mathcal{S} \setminus (J_t, R_t)$
end for
$d_i \leftarrow \max\{t = 1, \ldots, T \mid r_i \geq R_t\}$          // retirement times
// Step 3: Implement policy with static retirement times
$Z_i \leftarrow 0\ \forall i \in \mathcal{N}$
for $t = 1$ to $T$ do
    $S_t \leftarrow \{i \in \mathcal{N} \mid Z_i \leq c_i - 1$ and $t \leq d_i\}$
    $X_t \leftarrow$ Product purchased by type-$J_t$ customer when offered $S_t$
    $Z_{X_t} \leftarrow Z_{X_t} + 1$
end for
```

*is feasible and guarantees a* $(c_{\min} - 1)/(2c_{\min} - 1)$ *approximation of the optimal LP value, i.e.,*

$$\sum_{i=1}^{n} \sum_{j=1}^{n} r_i \widetilde{q}_i^j \geq \frac{c_{\min} - 1}{2c_{\min} - 1} z_{SLP}^*.$$

*Customer sequencing and product retirements (Step 2 of Algorithm 1).* In this step the policy sequences the customers selected in the first step to accommodate product retirements. Specifically, the policy arranges the customers in increasing order of $\widetilde{R}_t := R_{\pi^{\widetilde{J}_t}}(\widetilde{S}_t)$ ($R_\pi(S)$ is defined in (1.1)). We let $(J_t, \hat{S}_t)_{t=1}^T$ denote the new sequence of customers and assortments, with $R_t := R_{\pi^{J_t}}(\hat{S}_t)$. The product retirement policy is as follows: at the end of time period $t$, all products $i$ with $r_i < R_{t+1}$ are retired. In other words, customer $J_t$ is shown all products with inventory remaining such that $r_i \geq R_t$. The intuition behind this step is that under the MNL choice model, adding products with prices higher than the expected revenue of the assortment, only increases the expected revenue

Figure 1.1: Customer sequencing step in A1

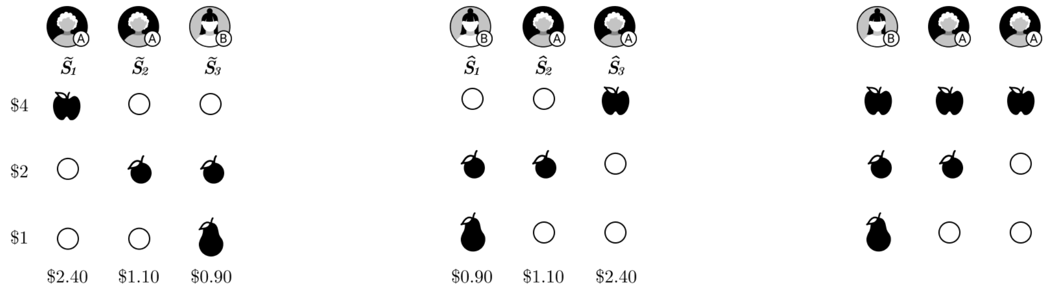

(a) Select customers (b) Order by revenue (c) Fill assortment gaps

*Note:* In this example there are three products ($n = 3$) and two customer types ($m = 2$). The prices of apples, oranges and pears are \$4, \$2 and \$1 respectively. Figure 1.1a shows the order in which customers are selected along with the corresponding revenue-maximizing assortment, and its expected revenue. Figure 1.1b shows the final sequence of customers, after arranging them in increasing order of the expected revenue. In Figure 1.1c, the gaps in the assortments are filled to make them revenue-ordered. It illustrates how the product retirement policy will work. Pears will be retired at $t = 1$, and oranges will be retired at $t = 2$.

of the new assortment. Moreover, this step makes all assortments revenue-ordered and therefore, nested, giving a natural product retirement policy candidate. Figure 1.1 illustrates the working of this step using a toy example. We show in Lemma 1.4 below that with this customer sequence and product retirement schedule, our policy earns at least half of the objective value of the feasible solution to (SLP) constructed in the customer selection step. The exact proof details can be found in Appendix A.2.3. Note that both the sequence and retirement times in this policy are static, and determined at the start of the algorithm. The policy A1 is therefore a non-adaptive policy.

**Lemma 1.4.** *The expected revenue of* A1 *is at least half of the objective value of* (SLP) *at* $\widetilde{\boldsymbol{q}}$*, i.e.,*

$$\textsc{Rev}(\mathrm{A1}) \geq \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} r_i \widetilde{q}_i^j .$$

The proof of Theorem 1.2 follows directly from combining the chain of inequalities in Lemmas 1.3 and 1.4. Therefore, the expected revenue of A1 is guaranteed to be at least $1/4 - o(1)$ times the LP upper bound. Details of the proof are deferred to Appendix A.2.1.

## 1.4 Asymptotically Optimal Policy

Policy A1 operates by constructing a good feasible solution to (SLP) and extracting a policy from it. In this section, we show that when an optimal (SLP) solution has a certain structure, which we call *consistent nested assortments*, we can construct a policy that mimics the optimal solution. We prove this policy is asymptotically optimal. We show that if there is only one customer type, i.e., $m = 1$, or if there are only two products, i.e., $n = 2$, then we can always obtain an optimal solution with consistent nested assortments. Interestingly, in our numerical section, we found that the optimal (SLP) solution we get from a standard solver has consistent nested assortments for a majority of instances, even for many products and customer types.

### 1.4.1 Consistent nested assortments

Even though product retirements are not explicitly captured among the constraints of (SLP), there is a connection between the two. For any feasible solution $\boldsymbol{q}$ of (SLP) and for every customer type $j$, there exists an ordered sequence of assortments such that (i) each assortment is a subset of the previous one and (ii) the quantities $\{q_i^j : i \in \mathcal{N}\}$ can be recovered using these assortments [9, 10]. In particular, there exists a collection of assortments $\mathcal{O}^j(\boldsymbol{q}) = \{O_1, \ldots, O_{|\mathcal{O}^j(\boldsymbol{q})|}\}$ such that $O_1^j \supset O_2^j \supset \cdots \supset O^j_{|\mathcal{O}^j(\boldsymbol{q})|}$, which we call *nested assortments*, and positive real numbers $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}) = \{\widetilde{\tau}_S^j \in \mathbb{R}_{>0} \mid S \in \mathcal{O}^j(\boldsymbol{q})\}$ where $\widetilde{\tau}_S^j$ is the amount of time for which assortment $S$ should be offered such that,

$$\sum_{S \in \mathcal{O}^j(\boldsymbol{q})} \widetilde{\tau}_S^j \cdot \pi_i^j(S) = q_i^j, \quad \forall i \in \mathcal{N} \cup \{0\}. \tag{1.5}$$

Consider a deterministic fluid setting, where customers are infinitesimal and visited at a constant rate. and a type-$j$ customer purchases product $i$ when offered an assortment $S$ at a deterministic rate of $\pi_i^j(S)$. In this setting, the product sales corresponding to (SLP) solution to customer type $j$ can be recovered as follows: offer assortment $O_1^j$ for $\widetilde{\tau}_1^j$ time periods, and then remove products in $O_1^j \setminus O_2^j$; offer the resulting assortment $O_2^j$ for $\widetilde{\tau}_2^j$ time periods and then remove products in $O_2^j \setminus O_3^j$; and so on and so forth.

**Procedure 1** NESTEDASSORTMENT$(\boldsymbol{q}, j)$

```
Input: Feasible solution $\boldsymbol{q}$ of (SLP), customer type $j$
Output: Collection of nested assortments $O^j(\boldsymbol{q})$ and offering times $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q})$
  initialize $S \leftarrow \{i \in \mathcal{N} \mid q_i^j > 0\}$, $\hat{q}_i \leftarrow q_i^j \ \ \forall i \in \mathcal{N} \cup \{0\}$, $O^j(\boldsymbol{q}) \leftarrow \{\}$
  while $S \neq \emptyset$ do
      $O^j(\boldsymbol{q}) \leftarrow O^j(\boldsymbol{q}) \cup \{S\}$
      $\Delta \leftarrow \min_{i \in S} \left\{ \frac{\hat{q}_i}{\pi_i^j(S)} \right\}$
      $\widetilde{\tau}_S^j(\boldsymbol{q}) \leftarrow \Delta$
      $\hat{q}_i \leftarrow \hat{q}_i - \Delta \pi_i^j(S) \quad \forall i \in S \cup \{0\}$
      $S \leftarrow \{i \in S \mid \hat{q}_i > 0\}$
  end while
  if $\hat{q}_0 > 0$ then
      $O^j(\boldsymbol{q}) \leftarrow O^j(\boldsymbol{q}) \cup \{\emptyset\}$
      $\widetilde{\tau}_\emptyset^j(\boldsymbol{q}) \leftarrow \hat{q}_0$
  end if
  return $O^j(\boldsymbol{q})$, $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}) = \{\widetilde{\tau}_S^j(\boldsymbol{q}) \mid S \in O^j(\boldsymbol{q})\}$
```

Given a feasible solution $\boldsymbol{q}$ and a customer type $j$, to construct this collection of nested assortments and offering times, we use an iterative algorithm based on the idea of deterministic fluid consumption. We start by offering all products in $O_1^j := \{i \in \mathcal{N} \mid q_i^j > 0\}$ and update fluid consumption at rate $\pi_i^j(O_1^j)$ for each product $i$ until the consumption of some product hits $q_i^j$. The process is repeated until sales of all products are recovered. Details are provided in Procedure 1. The algorithm runs efficiently and its output satisfies (1.5) as stated in Lemma 1.5. The algorithm is in a similar vein to the one provided by Feldman and Topaloglu [10] for the Markov chain choice model. We show that the tuple $(O^j(\boldsymbol{q}), \widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}))$ is in fact unique for each feasible solution $\boldsymbol{q}$ and customer type $j$.

**Lemma 1.5** (Correctness of Procedure 1)**.** *Procedure 1 terminates in at most n iterations, and returns nested assortments $O^j(\boldsymbol{q})$ and positive offering times $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q})$ which satisfy* (1.5)*. Further, $(O^j(\boldsymbol{q}), \widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}))$ are uniquely defined for each feasible solution $\boldsymbol{q}$ and customer type $j$.*

In the deterministic fluid setting, we see that the sales corresponding to a particular customer type can be recovered by sequentially removing products from an assortment, which is reminiscent of product retirements. However, note that the collection of nested assortments can be different for different customer types. Consequently, the order in which products are removed from the assortment can be different for different customer types. If the nested assortments corresponding to an (SLP) solution have this issue, then one can infer from the uniqueness of the nested assortments, that even in the deterministic fluid setting, no product retirement policy can mimic the sales corresponding to the solution. It follows that in order to be able to recover all sales using a product retirement policy, the order in which products are removed should be consistent across customer types. This condition is equivalent to $\mathcal{O}(\boldsymbol{q}) := \bigcup_{j=1}^{m} \mathcal{O}^j(\boldsymbol{q})$ being a collection of nested assortments, i.e., $\mathcal{O}(\boldsymbol{q}) = \{O_1, \ldots, O_{|\mathcal{O}|}\}$ such that $O_1 \supset O_2 \supset \cdots \supset O_{|\mathcal{O}|}$. We say an (SLP) solution has consistent nested assortments if it satisfies this condition.

**Definition 1.6** (Consistent nested assortments)**.** A feasible solution $\boldsymbol{q}$ of (SLP) has consistent nested assortments if there exists a collection of nested assortments $\mathcal{O}(\boldsymbol{q})$ and offering times $\boldsymbol{\tau}^j(\boldsymbol{q}) = \{\tau_S^j \in \mathbb{R}_{\geq 0} \mid S \in \mathcal{O}\}$ for all $j \in \mathcal{L}$ such that,

$$q_i^j = \sum_{S \in \mathcal{O}} \tau_S^j \pi_i^j(S) \quad i \in \mathcal{N} \cup \{0\}, \forall j \in \mathcal{L}. \tag{1.6}$$

In particular, given nested assortments $\mathcal{O}^j(\boldsymbol{q})$ and offering times $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q})$ for every customer type $j$, $\boldsymbol{q}$ has consistent nested assortments if $\mathcal{O}(\boldsymbol{q})$ and $\boldsymbol{\tau}^j(\boldsymbol{q})$ defined as

$$\mathcal{O}(\boldsymbol{q}) := \bigcup_{j=1}^{m} \mathcal{O}^j(\boldsymbol{q}), \quad \tau_S^j(\boldsymbol{q}) := \begin{cases} \widetilde{\tau}_S^j(\boldsymbol{q}) & S \in \mathcal{O}^j(\boldsymbol{q}) \\ 0 & S \in \mathcal{O}(\boldsymbol{q}) \setminus \mathcal{O}^j(\boldsymbol{q}) \end{cases} \tag{1.7}$$

satisfy (1.6) and $\mathcal{O}(\boldsymbol{q})$ is a collection of nested assortments.

If $\boldsymbol{q}$ has consistent nested assortments, then the sales of all products can be recovered as follows: offer assortment $O_1$ for $\tau_1^j$ time to each customer type $j \in \mathcal{L}$, then permanently retire products in $O_1 \setminus O_2$; offer assortment $O_2$ for $\tau_2^j$ time to each customer type $j \in \mathcal{L}$, then permanently retire products in $O_2 \setminus O_3$; and so on and so forth. Therefore, product retirements are sufficient to recover all sales. Moreover, if $\boldsymbol{q}$ is an optimal solution of (SLP), this policy maximizes revenue in the deterministic fluid setting. We show that this policy idea is also asymptotically optimal for our dynamic, stochastic problem. We illustrate this connection between an (SLP) solution which has consistent nested assortments and product retirements in the following example.

**Example 1.7.** Consider the following instance: $n = 2$, $m = 3$, $T = 45$, $\boldsymbol{r} = \{5, 2\}$, $\boldsymbol{c} = \{20, 30\}$, $\boldsymbol{b} = \{15, 18, 30\}$, $\boldsymbol{u}^1 = \{2, 5\}$, $\boldsymbol{u}^2 = \{1, 2\}$, $\boldsymbol{u}^3 = \{0.5, 0.5\}$. Table 1.1 shows optimal solution $\boldsymbol{q}$ and corresponding nested assortments for this instance. It can be seen that the optimal solution has consistent nested assortments with $\mathcal{O} = \{\{1, 2\}, \{1\}\}$. A natural candidate policy for product retirements and customer selections can mimic this optimal solution: visit all customers that are to be shown $\{1, 2\}$, then retire product 2, and then visit all customers that are to be shown $\{1\}$.

Table 1.1: Optimal solution and nested assortments for example

| $j$ | **1** | **2** | **3** |
|---|---|---|---|
| $q_1^j$ | 10 | 7 | 3 |
| $q_2^j$ | 0 | 4 | 3 |
| $q_0^j$ | 5 | 7 | 6 |

(a) Optimal solution

| $j$ | **1** | **2** | **3** |
|---|---|---|---|
| $\mathcal{O}^j(\boldsymbol{q})$ | $\{\{1\}\}$ | $\{\{1, 2\}, \{1\}\}$ | $\{\{1, 2\}\}$ |
| $\widetilde{\tau}^j(\boldsymbol{q})$ | $\{15\}$ | $\{8, 10\}$ | $\{12\}$ |

(b) Nested assortments and offering times from $\boldsymbol{q}$

- **Step 1:** Visit 12 type-3 customers and 8 type-2 customers.
- **Step 2:** Retire product 2.
- **Step 3:** Visit 10 type-2 customers and 15 type-1 customers.
- **Remark:** Customers are offered all products that have not been retired and have inventory left. ■

As seen from the example, if an optimal solution with consistent nested assortments can be obtained for an instance, a product retirement and customer selection policy can be constructed using the optimal solution. Generally, the offering times and total customer selections of each type may not be integral, and thus it requires some nuance when moving from the deterministic fluid setting of the LP to the stochastic and discrete assortment planning setting. Nonetheless, the key idea is to mimic the nested assortments obtained from the optimal LP solution, as shown in the example. We show in Theorem 1.8 that this policy, referred to as A2, is asymptotically optimal as $T$ and $c_{\min}$ grow large. Details about the exact structure of A2 and proof idea for Theorem 1.8 are deferred to Section 1.4.3.

**Theorem 1.8** ((1 − $\epsilon$)-Guarantee for A2)**.** *Consider $\epsilon$ such that $T \geq \frac{2m}{\epsilon}$ and $c_{\min} \geq \frac{2}{\pi\epsilon^2}$. If $\boldsymbol{q}$ is an optimal solution of* (SLP) *with consistent nested assortments, then* $\textsc{Rev}(A2) \geq (1-\epsilon)z^*_{SLP}$.

When there is only one customer type, i.e., $m = 1$, it is easy to see that every feasible solution has consistent nested assortments with $O(\boldsymbol{q}) = O^1(\boldsymbol{q})$. However, when $m \geq 2$, it is unclear whether a solution will have consistent nested assortments even in the simplest case where there are only two products, i.e., $n = 2$. For instance, in the two product case, the nested assortments corresponding to two customer types $j_1$ and $j_2$ may be $O^{j_1}(\boldsymbol{q}) = \{\{1,2\},\{1\}\}$ and $O^{j_2}(\boldsymbol{q}) = \{\{1,2\},\{2\}\}$ respectively. Since both collections have different assortments of size 1, no uniform order exists. However, we show in Section 1.4.2 that in the two product case, there always exists an optimal solution with consistent nested assortments, and it can be found efficiently. Moreover, in our numerical experiments in Section 1.5, our computed optimal solutions of the (SLP) had consistent nested assortments on 89% of the instances.

### 1.4.2 Two product case

When there are only two products, we show that there is always an optimal solution with consistent nested assortments (Theorem 1.9). Moreover, we provide an algorithm to find this solution efficiently. Details of the algorithm and a proof of Theorem 1.9 are provided in Appendix A.4.

**Theorem 1.9.** *For any instance with $n = 2$, there exists an optimal solution of* (SLP) *with consistent nested assortments and this can be found efficiently.*

In the two product case, the nested assortments outputted for a customer type with non-zero selections by the NESTEDASSORTMENT (Procedure 1) procedure can either be $\{\{1,2\}\}$, $\{\{1,2\},\{1\}\}$

or $\{\{1, 2\}, \{2\}\}$. The key step in the proof is to show that, for any feasible (SLP) solution, given two customer types such that the corresponding nested assortments are $\{\{1, 2\}, \{1\}\}$ and $\{\{1, 2\}, \{2\}\}$, a new feasible solution with the same objective value can be found such that the nested assortments corresponding to the two customer types are $\{\{1, 2\}, \{1\}\}$ and $\{\{1, 2\}\}$, or $\{\{1, 2\}, \{2\}\}$ and $\{\{1, 2\}\}$. In order to prove this, we describe two linear programs corresponding to the potential consistent nested assortments, with constraints reflecting that the consumption of products must be the same without requiring more customers of either type. We then look at the duals of these linear programs and prove that one of them must be feasible and bounded, implying that when there are two customer types we can find consistent nested assortments. Then, we provide a careful algorithm to deal with many customer types and prove that we only need to run our procedure for two customer types at most $m$ times (avoiding an infinite cycle that alternates between the two potential nested assortments). Overall, we are able to efficiently transform any optimal solution of (SLP) to an optimal solution with consistent nested assortments. Figure 1.2 gives an idea of how this transformation works.

We leave it as an open question whether an optimal solution to (SLP) with consistent nested assortments always exists for $n \geq 3$. If one is guaranteed to exist, it would also be necessary to show that such a solution can be found efficiently when there are multiple optimal solutions to (SLP). In our numerical experiments in Section 1.5, we show that the solutions returned by a solver usually have consistent nested assortments.

### 1.4.3 Structure of policy A2

In this section, we elaborate on the policy idea discussed in Section 1.4.1. We assume that we are given an optimal (SLP) solution with consistent nested assortments. The key idea is to mimic the sales of this optimal solution, while ensuring that constraints are satisfied sample path-wise. We saw that in a deterministic fluid setting, an optimal product retirement and customer selection policy can be constructed using consistent nested assortments and offering times. When the policy idea is translated from the deterministic fluid setting to the stochastic and discrete setting, there are

Figure 1.2: Transforming the optimal solution in the two product case

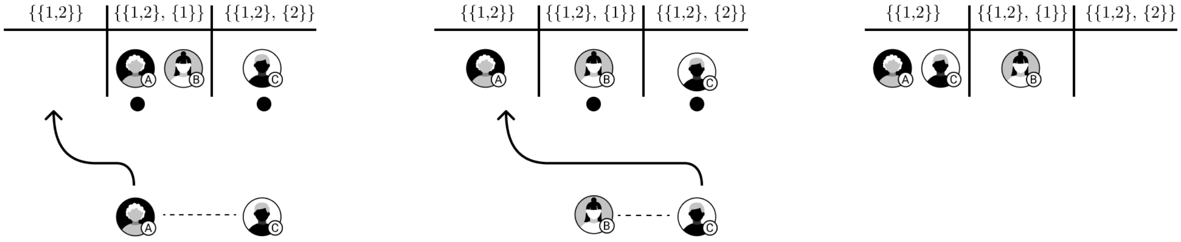

(a) (b) (c)

*Note:* Each time, we pick two customer types, one each from column 2 (nested assortments $\{\{1,2\},\{1\}\}$) and column 3 (nested assortments $\{\{1,2\},\{2\}\}$). We then transform the solution for these two customer types such that one of the customer types moves to column 1 (nested assortments $\{\{1,2\}\}$). In this example, we have three customer types. We first pick customer type $A$ and $C$ from columns 2 and 3 respectively, and after the transformation, $A$ moves to column 1, as depicted in Figure 1.2a. We then pick customer type $B$ and $C$ from columns 2 and 3 respectively, and after the transformation, $C$ moves to column 1, as depicted in Figure 1.2b. The final solution, as shown in Figure 1.2c, has consistent nested assortments $\{\{1,2\},\{1\}\}$.

two sources of losses. One, since customer selections of each type need to be integers that satisfy the demand and time horizon constraints; we use a rounding scheme to modify the offering times to achieve this, which introduces a rounding loss. Two, due to limited inventory of products, some sales will be lost, which will also incur a revenue loss. We construct A2 such that these losses can be bounded above as a function of $T$ and $c_{\min}$, and they disappear as $T$ and $c_{\min}$ grow large.

We describe A2 in two stages. The first stage is the customer and assortment selection stage. In this stage, the policy first calculates the consistent nested assortments and offering times using the optimal solution. The policy then implements a simple rounding procedure to get integer selections of each customer type, as well as a probability distribution over the nested assortments for each selected customer. The second stage is implementation, where the policy draws an assortment for each selected customer according to their probability distribution, and then visits the customers in decreasing order of the size of their assortment, showing them all products with available inventory from their assortment.

**Stage I: Customer and assortment selection**

- **Step 0:** Given an optimal solution $\boldsymbol{q}$ with consistent nested assortments, compute $O(\boldsymbol{q})$ and $\boldsymbol{\tau}^j(\boldsymbol{q})$ according to (1.7).

- **Step 1:** Add dummy customers in each customer type, that are offered no products, so that the total number of customers selected, denoted by $Y_j$, is integer. Denote the $k$-th type-$j$ customer by the tuple $(j, k)$. Since the total customers available, $(b_j)_{j \in \mathcal{L}}$, are integers, this step does not violate the demand constraints, however, it may violate the time horizon constraint.

- **Step 2:** For each customer type $j$, divide the offering times (which add up to $Y_j$) into $Y_j$ probability vectors, where the probability vector for customer $(j, k)$ is denoted by $\boldsymbol{a}_k^j = \{a_{k,S}^j \mid S \in O(\boldsymbol{q})\}$. It is not important for the analysis how this step is carried out, but for completeness, we specify a procedure in Algorithm 2.

- **Step 3:** Remove customers with the lowest expected revenue, until the total number of customers left is $T$.

We also provide pseudocode for this stage with implementation details in Algorithm 2. Lemma 1.10 states that at the end of the stage, the total customer selections are integer and satisfy all constraints, $(\boldsymbol{a}_k^j)_{j \in \mathcal{L}, k \in [Y_j]}$ are probability vectors, and the optimality loss is at most by a factor of $m/(T + m)$.

**Algorithm 2** A2: Stage I

```
Input: Optimal solution of (SLP) $\boldsymbol{q}$ with consistent nested assortments
Output: $\mathcal{O}$, $(Y_j)_{j\in\mathcal{L}}$ $(\boldsymbol{a}^j_k)_{j\in\mathcal{L},k\in[Y_j]}$
  // Step 0: Consistent nested assortments
  $\mathcal{O}^j, \widetilde{\boldsymbol{\tau}}^j \leftarrow \textsc{NestedAssortment}(\boldsymbol{q}, j)\ \forall j \in \mathcal{L}$
  $\mathcal{O} \leftarrow \bigcup_{j=1}^m \mathcal{O}^j$
  $\tau^j_S \leftarrow$ if $S \in \mathcal{O}^j$ then $\widetilde{\tau}^j_S$ else $0 \quad \forall S \in \mathcal{O}, j \in \mathcal{L}$
  // Step 1: Add dummy customers
  $\mathcal{O} \leftarrow \mathcal{O} \cup \{\emptyset\}$
  for $j = 1$ to $m$ do
    $\tau^j_\emptyset \leftarrow \left\lceil \sum_{S\in\mathcal{O}\setminus\{\emptyset\}} \tau^j_S \right\rceil - \sum_{S\in\mathcal{O}\setminus\{\emptyset\}} \tau^j_S$
    $Y_j \leftarrow \sum_{S\in\mathcal{O}} \tau^j_S$
  end for
  // Step 2: Construct probability vectors
  initialize $a^j_{k,S} \leftarrow 0\ \forall j \in \mathcal{L}, k \in [Y_j], S \in \mathcal{O}$
  for $j = 1$ to $m$ do
    initialize $k \leftarrow 1$, $l \leftarrow 1$, $\hat{\boldsymbol{\tau}}^j \leftarrow \boldsymbol{\tau}^j$
    while $k \le Y_j$ and $l \le |\mathcal{O}|$ do
      $\Delta \leftarrow \min\left\{1 - \sum_{S\in\mathcal{O}} a^j_{k,S}, \hat{\tau}^j_{O_l}\right\}$
      $a^j_{k,O_l} \leftarrow a^j_{k,O_l} + \Delta$
      $\hat{\tau}^j_{O_l} \leftarrow \hat{\tau}^j_{O_l} - \Delta$
      if $\sum_{S\in\mathcal{O}} a^j_{k,S} = 1$ then $k \leftarrow k + 1$
      end if
      if $\hat{\tau}^j_{O_l} = 0$ then $l \leftarrow l + 1$
      end if
    end while
  end for
  // Step 3: Reduce total customer selections to $T$
  while $\sum_{j\in\mathcal{L}} Y_j > T$ do
    $(\widetilde{j}, \widetilde{k}) \leftarrow \operatorname*{arg\,min}_{j\in\mathcal{L},k\in[Y_j]} \sum_{S\in\mathcal{O}} \sum_{i\in\mathcal{N}} a^j_{S,k} r_i \pi^j_i(S)$
    $Y_{\widetilde{j}} \leftarrow Y_{\widetilde{j}} - 1$
    $\boldsymbol{a}^{\widetilde{j}}_k \leftarrow \boldsymbol{a}^{\widetilde{j}}_{k+1}\ \forall k = \widetilde{k}, \ldots, Y_{\widetilde{j}}$
  end while
  return $\mathcal{O}$, $(Y_j)_{j\in\mathcal{L}}$ $(\boldsymbol{a}^j_k)_{j\in\mathcal{L},k\in[Y_j]}$
```

**Lemma 1.10.** *The output of Algorithm 2 satisfies the following.*

*(i)* $Y_j \in \mathbb{Z}$, $Y_j \leq b_j$ *for all* $j \in \mathcal{L}$ *and* $\sum_{j=1}^{m} Y_j \leq T$.

*(ii)* $\sum_{S \in O(\boldsymbol{q})} a_{k,S}^{j} = 1$ *for all* $j \in \mathcal{L}, k \in [Y_j]$ *and* $\sum_{k=1}^{Y_j} a_{k,S}^{j} \leq \tau_S^j$ *for all* $j \in \mathcal{L}, S \in O(\boldsymbol{q}) \setminus \{\emptyset\}$.

*(iii)* $\sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \sum_{S \in O(\boldsymbol{q})} a_{k,S}^{j} r_i \pi_i^j(S) \geq \frac{T}{T+m} z_{SLP}^{*}$.

**Stage II: Implementation of randomized policy**

- **Step 0:** From the previous stage, get a collection of nested assortments $O = \{O_1 \ldots, O_{|O|}\}$ such that $O_1 \supset O_2 \supset \ldots \supset O_{|O|}$, the number of customers of each type to visit $(Y_j)_{j \in \mathcal{L}}$, and probability vectors $(\boldsymbol{a}_k^j)_{j \in \mathcal{L}, k \in [Y_j]}$ for each customer denoting distributions over the nested assortments.
- **Step 1:** For every selected customer $(j, k)$, where $j \in \mathcal{L}, k \in [Y_j]$, draw assortment $D_k^j \sim Categorical((a_{k,S}^j)_{S \in O})$.
- **Step 2:** Construct customer sequence by sorting customers in decreasing order of $|D_k^j|$.
- **Step 3:** Visit customers in order of the customer sequence. Show customer $(j, k)$ all products in $D_k^j$ with inventory remaining when visited.

The second stage represents a product retirement policy. This is because $(D_k^j)_{j \in \mathcal{L}, k \in [Y_j]}$ are all sampled from a set of nested assortments, and if a customer $(j', k')$ is before customer $(j'', k'')$ in the sequence, then $D_{k'}^{j'} \supseteq D_{k''}^{j''}$. Moreover, it asymptotically achieves the expected revenue of the selected customers as the initial inventories increase, specifically at the rate of $O(1/\sqrt{c_{\min}})$. Define

$$\gamma(k) := 1 - e^{-k} \frac{k^k}{k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}, \tag{1.8}$$

where $\gamma \to 1$ as $k \to \infty$ at the rate of $O(1/\sqrt{k})$. In Lemma 1.11, we bound the revenue earned by the second stage of A2 by $\gamma(c_{\min})$ times the objective value corresponding to the rounded (SLP) solution from Stage I.

**Lemma 1.11.** *Given inputs* $(\mathcal{O}, (Y_j)_{j\in\mathcal{L}}, (\boldsymbol{a}_k^j)_{j\in\mathcal{L},k\in[Y_j]})$ *from Stage I, the expected revenue earned by the Stage II policy with the given inputs is at least*

$$\textsc{Rev}(\text{A2}) \geq \gamma(c_{\min}) \sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \sum_{S\in\mathcal{O}(\boldsymbol{q})} a_{k,S}^j r_i \pi_i^j(S) \ .$$

The proof of this lemma is provided in Appendix A.3.4. Combining guarantees of Lemmas 1.10 and 1.11 with the theorem conditions completes the proof of Theorem 1.8. Details of the proof are deferred to Appendix A.3.2.

## 1.5 Numerical experiments

In this section, we numerically evaluate the performance of the policies for multiple customer types described in Sections 1.3 and 1.4. We conduct these experiments on a dataset obtained from a large hardware provider. The dataset provides information on 8 different products and 2208 customers divided into 24 different customer types. For each customer type, we are provided with an MNL choice model fitted over the 8 products.

### 1.5.1 Simulation parameters

We test our policies under different scenarios by varying supply (i.e., the total number of products) and demand (i.e., the time horizon). Each scenario is defined by a *selection rate* and *load factor*. Here, the *selection rate* is the ratio of the time horizon and number of available customers, and the *load factor* is the ratio of the time horizon and the total inventory. To ensure robustness of our policy evaluation, we construct multiple problem *instances* for each scenario. These instances are generated by changing the distribution of product inventories and the number of customers available of each type, which we achieve by multiplying each of $(c_i)_{i=1}^n$ and $(b_j)_{j=1}^m$ with independent multiplicative noise. The multiplication factors are sampled from a scaled and translated beta distribution.

For each scenario, 100 instances are generated which are simulated 50 times for every policy.

We consider selection rates of 0.25, 0.5, 0.75 and 1.0. A selection rate of 1.0 indicates that all customers must be selected, and the selection problem purely involves constructing a customer sequence, whereas a selection rate of 0.25 lays high importance on selecting the right set of customers. We look at load factors ranging from 0.2 to 1.0, where a small load factor indicates that there is supply far exceeds the demand and product retirements are important, whereas a load factor of 1.0 is the case where supply is equal to demand.

### 1.5.2 Policies considered

**Benchmark policies.** To benchmark the performance of our policy, we compare its performance to some naive product retirement and customer selection policies. The simplest policy is to never retire products and select customers at random in each time period (NR-R). To make the benchmark smarter, we also consider another strategy that does not retire products, but chooses customers greedily at each step (NR-G). Our third benchmark policy (ER-G) aggressively retires products, by retiring all products with revenue less than the expected revenue of the best revenue-maximizing assortment at time 0, and selects customers greedily.

**Our policies.** We evaluate the performance of both policies for multiple customer types proposed in this chapter, i.e., A1 from Section 1.3 and A2 from Section 1.4.3. We also add a third policy to the consideration set, which is a slight modification of A1. A1 has the tendency to retire products aggressively, due to the fact that it is a static policy where the revenue thresholds are chosen greedily. For instance, A1 will retire all products whose revenue is less than the revenue threshold, even if that means that no more products will be left to offer to the remaining customers. To mitigate this, we introduce a parameter $x \in (0, 1]$ in the original A1 policy, which is multiplied to the revenue threshold $R_t$ to get a new, lower threshold and thus push the retirement times for products further. We refer to the policy parameterized by $x$ as $\mathrm{A}_{px}$. The policy $\mathrm{A}_{px}$ still retains the theoretical guarantees of A1, but with additional losses. In particular, it guarantees a $\left(\frac{x}{2(1+x)} - o(1)\right)$ fraction of the LP optimal value. In our simulations, we find that using $x = 0.75$ works reasonably

Table 1.2: Average revenue for policies under different scenarios as % of the upper bound.

| Selection | Load | NR-R | | NR-G | | ER-G | | A1 | | $A_{p0.75}$ | | A2 | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Rate | Factor | Worst | Avg | Worst | Avg | Worst | Avg | Worst | Avg | Worst | Avg | Worst | Avg | Success |
| 0.250 | 0.200 | 0.699 | 0.735 | 0.840 | 0.862 | 0.738 | 0.989† | 0.913 | 0.989† | 0.925 | 0.971 | 0.992 | 0.999* | 100% |
| | 0.400 | 0.708 | 0.760 | 0.839 | 0.884 | 0.525 | 0.897 | 0.834 | 0.970† | 0.902 | 0.965 | 0.986 | 0.997* | 100% |
| | 0.600 | 0.712 | 0.804 | 0.853 | 0.928 | 0.313 | 0.699 | 0.834 | 0.935 | 0.910 | 0.964† | 0.971 | 0.992* | 95% |
| | 0.800 | 0.750 | 0.857 | 0.896 | 0.967† | 0.278 | 0.553 | 0.796 | 0.900 | 0.872 | 0.956 | 0.975 | 0.986* | 75% |
| | 1.000 | 0.776 | 0.904 | 0.938 | 0.988*† | 0.224 | 0.504 | 0.786 | 0.875 | 0.864 | 0.942 | 0.976 | 0.984 | 38% |
| 0.500 | 0.200 | 0.733 | 0.766 | 0.847 | 0.863 | 0.794 | 0.988† | 0.928 | 0.987 | 0.934 | 0.965 | 0.994 | 1.000* | 100% |
| | 0.400 | 0.740 | 0.788 | 0.847 | 0.885 | 0.507 | 0.887 | 0.858 | 0.967† | 0.944 | 0.967† | 0.992 | 0.998* | 97% |
| | 0.600 | 0.745 | 0.825 | 0.856 | 0.921 | 0.281 | 0.700 | 0.796 | 0.933 | 0.900 | 0.965† | 0.984 | 0.995* | 94% |
| | 0.800 | 0.766 | 0.860 | 0.873 | 0.954 | 0.253 | 0.589 | 0.806 | 0.901 | 0.884 | 0.963† | 0.983 | 0.991* | 88% |
| | 1.000 | 0.798 | 0.912 | 0.913 | 0.985† | 0.166 | 0.505 | 0.776 | 0.866 | 0.863 | 0.945 | 0.981 | 0.989* | 63% |
| 0.750 | 0.200 | 0.762 | 0.805 | 0.860 | 0.874 | 0.715 | 0.965 | 0.935 | 0.990† | 0.933 | 0.965 | 0.996 | 1.000* | 100% |
| | 0.400 | 0.770 | 0.819 | 0.860 | 0.886 | 0.431 | 0.885 | 0.856 | 0.977† | 0.906 | 0.963 | 0.993 | 0.999* | 100% |
| | 0.600 | 0.784 | 0.848 | 0.861 | 0.911 | 0.285 | 0.735 | 0.812 | 0.945 | 0.902 | 0.967† | 0.987 | 0.996* | 99% |
| | 0.800 | 0.781 | 0.879 | 0.865 | 0.936 | 0.274 | 0.607 | 0.802 | 0.906 | 0.871 | 0.968† | 0.986 | 0.994* | 89% |
| | 1.000 | 0.818 | 0.921 | 0.897 | 0.970† | 0.316 | 0.515 | 0.763 | 0.869 | 0.851 | 0.946 | 0.985 | 0.992* | 63% |
| 1.000 | 0.200 | 0.808 | 0.875 | 0.872 | 0.888 | 0.566 | 0.907 | 0.860 | 0.988† | 0.902 | 0.965 | 0.995 | 1.000* | 100% |
| | 0.400 | 0.815 | 0.886 | 0.869 | 0.896 | 0.372 | 0.836 | 0.814 | 0.976† | 0.905 | 0.968 | 0.995 | 0.999* | 100% |
| | 0.600 | 0.826 | 0.895 | 0.866 | 0.903 | 0.334 | 0.748 | 0.825 | 0.953 | 0.877 | 0.967† | 0.991 | 0.998* | 99% |
| | 0.800 | 0.812 | 0.917 | 0.872 | 0.922 | 0.270 | 0.634 | 0.781 | 0.918 | 0.862 | 0.968† | 0.988 | 0.996* | 95% |
| | 1.000 | 0.859 | 0.947 | 0.850 | 0.943 | 0.207 | 0.523 | 0.765 | 0.886 | 0.839 | 0.955† | 0.987 | 0.993* | 83% |

* indicates best average performance

† indicates best average performance excluding A2

*Policy Guide*: NR-R - No retirement policy with random customer selection, NR-G - No retirement policy with greedy customer selection, ER-G - Early retirement policy with greedy customer selection, A1 - Our approximation policy (Section 1.3), $A_{p0.75}$ - Less aggressive A1, A2 - Our asymptotically optimal policy (Section 1.4.3).

well for our dataset. Parameterized policy $\mathrm{A}_{p0.75}$ guarantees a $3/14$ factor of the optimal revenue. Implementations for these policies can be found online.

### 1.5.3 Performance evaluation

The policies are evaluated based on the average revenue gained by the policy compared to the optimal value of the LP benchmark. For each scenario and policy, we calculate the ratio of the average revenue earned and the optimal LP value for each instance. In Table 1.2 we report the average ratio and minimum ratio over all the instances of that scenario.

For A2, we compute an optimal solution of (SLP); if the optimal solution has consistent nested assortments, then we proceed with the algorithm, else we discard it. In the “Success" column, we report the percentage of instances that were not discarded for the given scenario. The overall success rate across all scenarios comes out to be 89%. The number of discarded instances goes up with load factor. We believe this is due to the fact that at high load factors, multiple optimal solutions exist since all products can be sold off. Among the non-discarded instances, the performance of this algorithm is expectedly good and also less volatile compared to all other algorithms. The performance of the algorithm improves with lower load factors (larger initial inventories) and higher selection rates (larger time horizon) which is consistent with the assumptions of Theorem 1.8. Since this policy does not guarantee a solution, we also look at what policies perform the best after excluding A2.

In all supply-demand scenarios, our policy A1 and parametrized variant $\mathrm{A}_{p0.75}$consistently outperform the benchmarks and their theoretical performance guarantees. The policy A1 works very well when the load factors are low but suffers from early retirements when the load factor is high, however $\mathrm{A}_{p0.75}$ makes up for this and has more balanced performance. The no retirement policies NR-R and NR-G do well when the load factor and selection rate are high, as these are low supply, high demand scenarios where a majority of the products are purchased over the time horizon. These policies do not perform well in scenarios when product retirements are necessary. It is also worth noting that NR-G consistently outperforms NR-R, especially at low selection rates, emphasizing the importance of the customer selection decisions. The performance of the early retirement policy ER-G is in contrast to that of the no retirement policies. ER-G has its best performance when the load factor and selection rate are low, i.e., high supply, low demand scenarios, where offering only the highest revenue product suffices. However, ER-G performs dismally when demand is high and is also very volatile in its performance.

In conclusion, we note that there is a visible trade-off between retiring products early and retiring products late. Simple benchmark policies which do not attempt to deal with this trade-off do well in certain scenarios, but fail to perform consistently. Our policies consistently perform well in a wide range of scenarios, as they make an effort to correctly manage this trade-off.

## 1.6 Conclusion

In this chapter, we consider a novel framework for revenue management where sellers control the sequence of customers and can only retire products. Under this framework, we design policies for the multinomial logit choice model. We observed that a sales based LP can be used to upper bound the performance of any dynamic policy. We first propose a static policy for both product retirement and customer selection which guarantees at least a $(1/4 - o(1))$ fraction of the optimal expected revenue. This policy first selects customers and a revenue threshold greedily to construct a feasible solution to the LP upper bound, and then visits the selected in order of the revenue threshold, offering all products with revenues higher than the threshold. We construct policy that is asymptotically optimal as the time horizon and inventory grow large and gave lower bounds on the convergence rate when an optimal solution to the LP upper bound has specific structure, i.e., consistent nested assortments. When all customers' purchase decisions can be captured by a single multinomial logit model, this nested structure exists naturally in every feasible solution of the LP upper bound. When there are multiple customer types, we present an efficient algorithm to construct an optimal solution with consistent nested assortments when there are only two products. We conduct numerical experiments on a proprietary B2B dataset that consists of multiple customer types being sold multiple technology products, and observe that our proposed policies outperform several benchmarks, particularly when product retirements are important for increasing revenue.

# Chapter 2: Simple Policies for Joint Pricing and Inventory Management

## 2.1 Introduction

In this chapter, we study a fundamental joint pricing and inventory management problem. We focus on an infinite horizon, single product system with continuous review. Customers arrive with unit demand according to a Poisson process and make a purchase if their valuation for the product exceeds the current price. The seller incurs a holding costs for excess inventory and shortage costs for backlogged demand (when allowed). Moreover, there is a fixed cost for ordering inventory and orders are fulfilled immediately. The goal for the seller is to maximize the long run average profit, i.e., the combined revenue from the customers minus the ordering, holding, and backlogging costs. In this setting, an optimal policy has many parameters and is highly dynamic, with the price changing every time a unit is sold. In this work, we analyze the effectiveness of simple pricing policies where only one or a few prices are ever offered, and provide strong performance guarantees against the optimal policy.

Dynamic pricing, where the price keeps changing with the inventory level, requires that sellers be able to query the current inventory level in real-time and have a mechanism to update the price frequently. Thus, the implementation of a fully dynamic pricing policy may not be feasible due to business and logistical constraints. Moreover, inventory systems are known to suffer from inventory record inaccuracy [38, 39], which may result in selecting the wrong prices. The trajectory of prices in dynamic pricing policies is also highly dependent on an accurate demand model and thus suffers when there is model mis-specification [40]. A simple policy is less susceptible to these issues. For instance, a static policy where the same price is offered throughout requires is entirely unaffected by inventory record inaccuracy. Moreover, a static policy is also easier to learn (for e.g. [41, 42, 43]) and is less sensitive to the choice of demand model, in contrast to a high dimensional optimal policy [44].

From the customer perspective, dynamic pricing may not be appealing to customers, who do not usually have access to the underlying inventory state. In fact, customers often are strategic and wait for favorable prices using price trackers [45, 46], which is difficult to account for in a model and results in less profit than expected (refer to [47] for a recent survey). In contrast, a simple pricing policy is more strategy-proof [48, 49] and requires less burden on customers to get favorable prices, which is often a luxury for the well-informed. In this work, we assume that the model does not suffer from inventory inaccuracy, model mis-specification, strategic customer behavior which all clearly disadvantage the optimal dynamic pricing strategy. We do show that even under these ideal conditions, a simple price strategy with one or a few prices can achieve both revenue and costs that are very close to the optimal policy.

### 2.1.1 Contributions

The optimal dynamic pricing policy in the model we consider is known to be an $(s, S, \boldsymbol{p})$ policy [50]: when the inventory level drops to $s$ units, the seller immediately places an order to replenish the inventory to $S$ units. The optimal pricing policy $\boldsymbol{p}$ has a different price for all $S - s$ inventory states. We provide multiplicative performance guarantees for simple policies compared to the optimal dynamic pricing policy. Since profit is a mixed objective of revenue minus costs, it is not possible to give useful worst-case multiplicative guarantees on profit. Consequently, we divide the profit into two components, revenue and costs, and compare policies individually on the two metrics. Our results are as follows:

- We first consider the case where backlogging orders is not allowed, which is equivalent to a lost sales model, and focus on the class of monotone hazard rate (MHR) valuation distributions which include most standard demand models. The optimal policy in this setting is an $(0, S^*, \boldsymbol{p}^*)$. We show that keeping the inventory policy the same and opting for a particular *static* price policy, one can obtain more revenue than the optimal dynamic policy, maintain the same ordering costs, and lose a factor of at most $1 + \ln(S^*)$ on holding costs. We then show that the inventory policy can be adjusted to obtain a better bound of $\sqrt{1 + \ln(S^*)}$ for the total costs while still earning at least as much revenue as the optimal dynamic policy.

- We show that when valuations are uniform, which corresponds to the classic linear demand model, our static price policy guarantees at least as much revenue as the optimal policy while increasing costs by at most a factor of 1.225.
- We construct examples showing that our analyses are nearly tight. In particular, the worst-case valuation distribution corresponds to an exponential demand model.
- When backlogging orders is allowed, we show that the holding cost ratio can be arbitrarily bad for any static pricing policy. For this setting, we propose a *3-price policy*, where there is a different price when the on-hand inventory level is positive, zero, or negative. Suppose $(s^*, S^*, \boldsymbol{p}^*)$ is the optimal dynamic policy. For any MHR value distribution, our policy earns at least as much revenue as the optimal dynamic policy. Our policy also incurs the same ordering costs as the optimal dynamic policy while the holding cost ratio is at most $1 + \ln(S^* - s^*)$.
- In numerical experiments, we compare the profits for our proposed policy as well as the optimal static policy, with the optimal dynamic profit in the lost sales setting. We find that static policies capture most of the profit in a majority of cases. Moreover, for the instances where the profit ratio is low, the profit margin for the optimal dynamic policy is close to zero.

### 2.1.2 Literature Review

The benefit of jointly optimizing inventory and pricing has been noted since the early development of periodic review inventory theory by Whitin [51]. A mainstream of the literature focuses on the structure of the profit-optimal control policies. In a multi-period setting, Federgruen and Heching [52] prove that a list-price-base-stock policy is optimal for both the non-stationary finite-horizon and stationary infinite-horizon problems when considering the holding and variable ordering costs. With fixed ordering costs, Chen and Simchi-Levi [53] demonstrated that the $(s, S, p)$

policy is optimal for the finite horizon problem: the optimal inventory policy follows the $(s, S)$ policy while the optimal pricing strategy is a list-price policy. Moreover, they extended their earlier results to the infinite horizon case for both the average and discounted profit criterion. Huh and Janakiraman [54] provide an approach for generalizing many of the early results for both back-order and lost sales settings. In continuous-review models such as the one we study, the joint pricing-inventory problem has been investigated under Poisson demand, where the arrival rate of the Poisson process depends on the price. In an infinite-horizon setting, Chen and Simchi-Levi [50] proved the existence of a stationary $(s, S, p)$ inventory policy that maximizes expected profit in both the discounted and average cases. Chao and Zhou [55] obtained the closed-form solutions for the optimal policy and developed efficient algorithms to compute them within the same Poisson demand model with joint pricing-inventory decisions. Our model is built upon the optimal structure results in the literature and compares the performance of a static price to the optimal dynamic pricing policy. The value of dynamic pricing in inventory systems has only been investigated numerically in the literature (see e.g. [52, 56, 57]).

Another stream of literature examines the efficacy of static pricing policies in various settings where the optimal policy is inherently dynamic. Gallego and Van Ryzin [58] proposed the use of a simple fixed price policy in a single-product revenue management setting and provide performance guarantees. Recently, Chen et al. [48] studied a multi-product revenue management problem over a fixed horizon with strategic customers and demonstrated similar guarantees. Ma et al. [8] considered the pricing and assortment control problem and provided a performance guarantee on a deterministic pricing policy in the form of a price calendar. A recent paper by Besbes et al. [59] analyzes the efficacy of static pricing policies in the reusable resources setting. market share, and service level from the optimal dynamic pricing policy. Recently, Lei et al. [60] consider a joint inventory and pricing problem in a one-warehouse multi-store setting with lost sales and Poisson demand. They prove a theoretical lower bound on the performance of any policy that does not dynamically adjust prices, regardless of how sophisticated the replenishment policy is.

The rest of this chapter is organized as follows. In Section 2.2, we describe the model. In

Section 2.3, we describe the main policy idea and results, which hold for the lost sales version of our model. In Section 2.4, we consider the extension with backlogging costs. In Section 2.5, we present numerical examples of the performance of simple pricing policies. We conclude our chapter in Section 2.6.

## 2.2 Model and Preliminaries

We consider a joint pricing and inventory management problem for a single product in a continuous review inventory system over an infinite time horizon. The customers arrive randomly with unit demand according to a Poisson process with an arrival rate $\Lambda$. Each customer has a random valuation $V$ for the product drawn independently from a known distribution with cumulative distribution function $F(\cdot)$ and density function $f(\cdot)$. The probability that a customer purchases a product if it is priced at $p$ is thus $\bar{F}(p) := \mathbb{P}(V \geq p) = 1 - F(p)$. Customers with valuations lower than the price $p$ leave the system without purchasing anything. The arrival rate of purchasing customers when the price is $p$ is denoted by $\lambda(p) := \Lambda\bar{F}(p)$. Since $\lambda(p)$ is a monotonically decreasing function of $p$, it has a unique inverse function $p(\lambda)$, which is the price offered to maintain a demand arrival rate of $\lambda$. We assume that the revenue rate $p(\lambda)\lambda$ is concave in $\lambda$ which is equivalent to the standard assumption that valuation distribution is regular.

We focus our analysis on a special class of regular distributions with monotone hazard rates (MHRs). For an MHR distribution, the hazard rate, i.e., $f(p)/\bar{F}(p)$, is non-decreasing in $p$. Common examples of MHR valuation distributions lead to the following demand rate functions:

- **Linear Demand**: When $V \sim Uniform(0, 1/b)$ and $\Lambda = a$, $a, b > 0$, then the demand rate takes the form

$$\lambda(p) = a(1 - bp).$$

- **Exponential Demand**: When $V \sim Exp(b)$ and $\Lambda = a$, $a, b > 0$, then the demand rate takes the form

$$\lambda(p) = ae^{-bp}.$$

- **Logistic Demand**: When $V \sim Logistic(p_\circ, 1/b)$ truncated to positive reals, $\Lambda = a, a, b > 0$, then the demand rate takes the form

$$\lambda(p) = \frac{a(1 + e^{-bp_\circ})}{1 + e^{b(p-p_\circ)}}.$$

Since demand follows a Poisson process, the seller only makes replenishment and pricing decisions based on the current inventory level. We assume the supply lead time is zero as in [55] and [50]. The costs we consider are a fixed ordering cost, holding cost per unit per period, and backlogging cost per unit per period. The seller pays a fixed ordering cost of $K$ for each order. The holding cost is linear in the on-hand inventory level, expressed as $h$ per unit product per time period. The unsatisfied demand is assumed to be backlogged, with a linear backordering cost of $\pi$ unit product per time period. We separately consider settings with and without backlogging. In our model, the setting where demand cannot be backlogged is equivalent to a lost sales model since there are no lead times and customers only purchase one unit of the product at a time.

The objective is to maximize the long-run average profit over an infinite time horizon. [50] proved that optimal policy under these conditions is an $(s, S, \boldsymbol{p})$ policy. The inventory replenishment follows an $(s, S)$ policy: when the inventory level drops to $s$, the seller immediately places an order that increases the inventory level to $S$. Here, $s$ and $S$ are referred to as the reorder point and order-up-to level, respectively. The optimal price depends on the on-hand inventory: $\boldsymbol{p}$ denotes the vector that contains the optimal price for each inventory state. When backlogging is not allowed, the optimal policy becomes an $(S, \boldsymbol{p})$ policy, where $s = 0$. Since each price is associated with a unique arrival rate, we use notation $(s, S, \boldsymbol{\lambda})$ and $(s, S, \boldsymbol{p})$ interchangeably to denote a joint pricing and inventory control policy.

In an $(s, S, \boldsymbol{\lambda})$ policy, the inventory control policy is a $(s, S)$ policy whereby $S - s$ units are ordered every time the inventory state reaches $s$. Accordingly, the inventory state cycles between states $S, S-1, \ldots, s+2, s+1$. The price (and correspondingly demand arrival rate) is only a function of the inventory state, and $\lambda_i$ is the demand arrival rate at state $i$.

### 2.2.1 Performance metrics

The long-term average profit $\textsc{Profit}(s, S, \boldsymbol{\lambda})$ is composed of three components: revenue ($\textsc{Rev}(s, S, \boldsymbol{\lambda})$), ordering costs ($\textsc{Order}(s, S, \boldsymbol{\lambda})$) and holding + backordering costs ($\textsc{Hold}(s, S, \boldsymbol{\lambda})$),

$$\textsc{Profit}(s, S, \boldsymbol{\lambda}) := \textsc{Rev}(s, S, \boldsymbol{\lambda}) - \textsc{Order}(s, S, \boldsymbol{\lambda}) - \textsc{Hold}(s, S, \boldsymbol{\lambda}). \tag{2.1}$$

For convenience, we use holding costs to refer to holding + backordering costs in the rest of the chapter. In our theoretical analysis, we give multiplicative guarantees for the performance our policies with respect to the optimal dynamic policy on each of these components separately.

**Formulations.** All metrics depend on the steady-state distribution of the inventory state and the inter-arrival time between orders. Let $T(s, S, \boldsymbol{\lambda})$ be the expected inter-arrival time between orders and $\pi_i(s, S, \boldsymbol{\lambda})$ be the steady-state probability that the system is in inventory state $i$. Since the time taken for demand to arrive in inventory state $i$ is $Exp(\lambda_i)$,

$$T(s, S, \boldsymbol{\lambda}) = \sum_{i=s+1}^{S} \frac{1}{\lambda_i}. \tag{2.2}$$

It follows from the elementary renewal theorem that the steady-state probability for state $i$ is

$$\pi_i(s, S, \boldsymbol{\lambda}) = \frac{\frac{1}{\lambda_i}}{T(s, S, \boldsymbol{\lambda})}. \tag{2.3}$$

Using the above definitions, the metrics can be expressed as,

$$\textsc{Rev}(s, S, \boldsymbol{\lambda}) = \frac{\sum_{i=s+1}^{S} p(\lambda_i)}{T(s, S, \boldsymbol{\lambda})}, \ \textsc{Order}(s, S, \boldsymbol{\lambda}) = \frac{K}{T(s, S, \boldsymbol{\lambda})}, \ \textsc{Hold}(s, S, \boldsymbol{\lambda}) = \sum_{i=s+1}^{S} f(i)\pi_i(s, S, \boldsymbol{\lambda}), \tag{2.4}$$

where $f(i) := hi\mathbb{I}[i > 0] - \pi i\mathbb{I}[i < 0]$. The first two follow from the elementary renewal theorem while the holding cost formulation uses the steady state distribution of the inventory state.

### 2.2.2 Optimal dynamic policy structure

In this section, we review results on the structure of the optimal prices that were discovered in [55], as they will be useful throughout the chapter. We leverage the structure of the optimal dynamic policy to prove our results for holding costs as well as tightness.

Define $g_{i,\gamma}(\lambda)$ as,

$$g_{i,\gamma}(\lambda) := p(\lambda) - \frac{ih + \gamma}{\lambda}, \tag{2.5}$$

where $\gamma \in \mathbb{R}$ is a constant that is a proxy for the long-term average profit while $i$ is an integer denoting the inventory state. We state the following lemma, which is directly from Theorems 1 and 2 from [55].

**Lemma 2.1** (Theorems 1 and 2 [55]). *The optimal dynamic policy* $(s^*, S^*, \boldsymbol{\lambda}^*)$ *with optimal long-term average profit* $\gamma^*$ *satisfies the following.*

*(i) Given optimal profit* $\gamma^*$*, the optimal rate* $\lambda_i^*$ *is the maximizer of* $g_{i,\gamma^*}(\lambda)$*. By first order conditions:*

$$g'_{i,\gamma^*}(\lambda_i^*) = 0 \Rightarrow p'(\lambda_i^*) = -\frac{hi + \gamma^*}{(\lambda_i^*)^2}.$$

*(ii) Given optimal profit* $\gamma^*$*, the optimal inventory policy* $(s^*, S^*)$ *is given by:*

$$s^* = \max\{i \le 0 \mid g_{i,\gamma^*}(\lambda_i^*) \le 0\}, \quad S^* = \begin{cases} 0 & g_{1,\gamma^*}(\lambda_1^*) < 0 \\ \max\{i > 0 \mid g_{i,\gamma^*}(\lambda_i^*) \ge 0\} & \textit{otherwise} \end{cases}.$$

*(iii) Given policy* $(s, S, \boldsymbol{\lambda})$*, the corresponding long-term average profit* $\gamma$ *is the unique solution to the equation*

$$\sum_{i=s+1}^{S} g_{i,\gamma}(\lambda_i) = K.$$

*(iv) The optimal rates are decreasing from state* $s + 1$ *to* $0$*, i.e.,* $\lambda_{s+1}^* \le \ldots \le \lambda_0^*$ *and increasing from* $0$ *to* $S^*$*, i.e.,* $\lambda_0^* \le \ldots \le \lambda_{S^*}^*$*.*

Lemma 2.1 gives the structure of the optimal dynamic policy $(s^*, S^*, \boldsymbol{\lambda}^*)$ as a function of the optimal profit $\gamma^*$. Interestingly, note that the optimal rates in Lemma 2.1(i) do not depend on the inventory policy or any of the other arrival rates outside of being a function of the optimal profit, $\gamma^*$. Chao and Zhou [55] show that root-finding schemes can be used to find $\gamma^*$ that satisfies Lemma 2.1(iii) for policy $(s^*, S^*, \boldsymbol{\lambda}^*)$.

## 2.3 Lost Sales Model

We first analyze the setting where backlogging orders is not allowed, i.e., demand is lost if it cannot be fulfilled with on-hand inventory. For our model, it implies that an order is placed immediately when there is no on-hand inventory left, and the optimal dynamic policy is an $(S, \boldsymbol{p})$ policy. In this setting, we show that a static price policy is near-optimal. We first describe the static price that we use in our main theoretical results in this section. Note that this is not the profit-maximizing price, but a proxy that is more amenable to analysis. In our numerical experiments section, we compare the performance of the profit-maximizing price with our static price. After giving our static price construction, we state our main results along with a proof and tight examples.

### 2.3.1 Static price construction

The static price that we use for our analysis is, in fact, derived from the optimal dynamic pricing policy itself. In particular, we choose the static price such that the corresponding arrival rate is equal to the expected arrival rate of the optimal dynamic pricing policy. A related idea was used by Besbes et al. [59], who studied pricing in an Erlang loss system. However, they used a modification of the expected arrival rate under the optimal policy and their analysis does not require studying a mixed objective. Suppose $(S^*, \boldsymbol{\lambda}^*)$ is the optimal dynamic policy where $\boldsymbol{\lambda}^* = \{\lambda_1^*, \ldots, \lambda_{S^*}^*\}$. Then our constructed static rate $\lambda_{static}$ is given by

$$\lambda_{static} := \mathbb{E}_i[\lambda_i^*] = \sum_{i=1}^{S^*} \pi_i(S^*, \boldsymbol{\lambda}^*)\lambda_i^* = \frac{\sum_{i=1}^{S^*} \lambda_i^* \frac{1}{\lambda_i^*}}{\sum_{i=1}^{S^*} \frac{1}{\lambda_i^*}} = \frac{S^*}{\sum_{i=1}^{S^*} \frac{1}{\lambda_i^*}}, \tag{2.6}$$

where the third equality follows from (2.2) and (2.3). We denote the static pricing policy using the notation $\lambda_{static} \cdot \mathbf{1}$, where $\mathbf{1}$ is the vector of all ones. The choice of this static rate allows us to prove strong performance guarantees directly comparing to the optimal policy, in contrast to methods that rely on using a deterministic relaxation of the problem.

### 2.3.2 Main results

Our main results for this section are captured in Theorems 2.2, 2.3, and 2.4. We denote revenue, ordering costs and holding costs for policy $(S, \boldsymbol{\lambda})$ as $\textsc{Rev}(S, \boldsymbol{\lambda})$, $\textsc{Order}(S, \boldsymbol{\lambda})$, and $\textsc{Hold}(S, \boldsymbol{\lambda})$, respectively. For our first set of results, we modify the pricing policy without changing the inventory policy, i.e., the static policy uses the same value of $S$ as the optimal dynamic policy. The guarantees under this setting are given in Theorem 2.2.

**Theorem 2.2** (Static Price Guarantees). *Let $(S^*, \boldsymbol{\lambda}^*)$ be the optimal dynamic policy for an instance where customer valuations follow an MHR distribution. Then the static policy $(S^*, \lambda_{static} \cdot \mathbf{1})$ has the following guarantees:*

- *Revenue ratio:*
$$\frac{\textsc{Rev}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Rev}(S^*, \boldsymbol{\lambda}^*)} \geq 1.$$
- *Ordering cost ratio:*
$$\frac{\textsc{Order}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)} = 1.$$
- *Holding cost ratio:*
$$\frac{\textsc{Hold}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} \leq 1 + \ln(S^*).$$

For the next set of results, we show that modifying the inventory policy improves our guarantees. As can be observed in Theorem 2.2, using order-up-to level $S^*$ with our static pricing policy achieves an ordering cost ratio of 1 while the holding cost ratio can be greater than 1. Decreasing the order-up-to level reduces the holding cost ratio while increasing the ordering cost ratio.

Thus, we can adjust the order-up-to level in order to minimize the maximum of the two cost ratios. Theorem 2.3 specifies the guarantees with our static rate and the optimized inventory policy.

**Theorem 2.3** (Guarantees for Optimized Inventory Policy)**.** *Let* $(S^*, \lambda^*)$ *be the optimal dynamic policy for an instance where customer valuations follow an MHR distribution. Then there exists* $S_{static} \leq S^*$ *such that the static policy* $(S_{static}, \lambda_{static} \cdot \mathbf{1})$ *has the following guarantees:*

- *Revenue ratio:*
$$\frac{\textsc{Rev}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Rev}(S^*, \lambda^*)} \geq 1.$$
- *Total cost ratio:*
$$\frac{\textsc{Costs}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(S^*, \lambda^*)} \leq \sqrt{1 + \ln(S^*)},$$
*where* $\textsc{Costs}(S, \lambda) = \textsc{Order}(S, \lambda) + \textsc{Hold}(S, \lambda)$.

Theorems 2.2 and 2.3 give performance guarantees for any MHR valuation distribution, and in Section 2.3.4 we show that there exist MHR distributions for which the cost guarantees are tight up to a constant factor. The proof for both theorems is provided in Section 2.3.3.

We give improved bounds for the case of linear demand, which corresponds to customer valuations being drawn from a uniform distribution. When demand is linear, we obtain small constant bounds for the cost ratios, as stated in Theorem 2.4 below, while still guaranteeing at least as much revenue as the optimal policy.

**Theorem 2.4** (Constant Factor Guarantees for Linear Demand)**.** *Let* $(S^*, \lambda^*)$ *be the optimal dynamic policy for an instance where demand is linear. The static policies we construct have the following improved guarantees for the costs.*

*(i) When the inventory policy is unchanged, the holding cost ratio for policy* $(S^*, \lambda^*)$ *is*

$$\frac{\textsc{Hold}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \lambda^*)} \leq 1.5.$$

*(ii) When the inventory policy is adjusted to optimize the overall cost ratio, there exists* $S_{static} \leq S^*$ *such that the total cost ratio for policy* $(S_{static}, \lambda_{static} \cdot \mathbf{1})$ *is*

$$\frac{\textsc{Costs}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(S^*, \boldsymbol{\lambda}^*)} \leq \sqrt{1.5} \leq 1.225,$$

*where* $\textsc{Costs}(S, \boldsymbol{\lambda}) = \textsc{Order}(S, \boldsymbol{\lambda}) + \textsc{Hold}(S, \boldsymbol{\lambda})$.

We provide a proof of Theorem 2.4 in Appendix B.1.3. The main idea is that when demand is linear, there is limited variation in the arrival rates of the optimal dynamic policy, so a static price is not far off from the optimal dynamic policy's behavior.

### 2.3.3 Proof of Theorems 2.2 and 2.3

For the proof of Theorems 2.2 and 2.3, we first determine the ratios of the performance metrics when comparing our static policy with a general $S$ to the optimal dynamic policy. Then we conclude the proof by substituting $S = S^*$ for Theorem 2.2 and $S = S_{static}$ for Theorem 2.3.

**Revenue ratio.** MHR distributions are a special case of regular distributions. Moreover, for any regular distribution, the revenue rate function $\lambda p(\lambda)$ is concave in $\lambda$. Using this fact, we show that the revenue rate for the static policy $(S, \lambda_{static} \cdot \mathbf{1})$ is greater than or equal to the revenue rate for the optimal dynamic policy $(S^*, \boldsymbol{\lambda}^*)$, for any value of $S$.

$$\begin{aligned}\textsc{Rev}(S, \lambda_{static} \cdot \mathbf{1}) &= \frac{\sum_{i=1}^{S} p(\lambda_{static})}{T(S, \lambda_{static} \cdot \mathbf{1})} = \frac{S p(\lambda_{static})}{\frac{S}{\lambda_{static}}} = \lambda_{static} p(\lambda_{static}) = \mathbb{E}_i[\lambda_i^*] p(\mathbb{E}_i[\lambda_i^*]) \\ &\geq \mathbb{E}_i[\lambda_i^* p(\lambda_i^*)] = \sum_{i=1}^{S^*} \pi_i(S^*, \boldsymbol{\lambda}^*) \lambda_i^* p(\lambda_i^*) = \frac{\sum_{i=1}^{S^*} p(\lambda_i^*)}{T(S^*, \boldsymbol{\lambda}^*)} = \textsc{Rev}(S^*, \boldsymbol{\lambda}^*),\end{aligned}$$

where the first and last equality both follow from the definition of revenue rate in (2.4), second equality follows from (2.2) and substituting $\lambda_i = \lambda_{static}$ for $i = 1, \ldots, S$, the fourth equality follows from the definition of $\lambda_{static}$ in (2.6), the inequality follows from Jensen's inequality and the concavity of $\lambda p(\lambda)$, and the sixth equality follows from (2.3). Therefore, for any MHR distribution and any order-up-to level $S$,

$$\frac{\textsc{Rev}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Rev}(S^*, \boldsymbol{\lambda}^*)} \geq 1. \tag{2.7}$$

**Ordering cost ratio.** From the definition of $\lambda_{static}$ in (2.6) and (2.2), it follows that

$$\lambda_{static} = \frac{S^*}{T(S^*, \boldsymbol{\lambda}^*)}.$$

Therefore, for the static policy $(S, \lambda_{static} \cdot \mathbf{1})$, the ordering cost rate is given by,

$$\textsc{Order}(S, \lambda_{static} \cdot \mathbf{1}) = \frac{K}{\frac{S}{\lambda_{static}}} = \frac{S^* K}{S T(S^*, \boldsymbol{\lambda}^*)}.$$

Therefore, for our construction of the static rate and (2.4), we have

$$\frac{\textsc{Order}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)} = \frac{S^*}{S}. \tag{2.8}$$

**Holding cost ratio.** We start with the holding cost rate for the static policy. Note that since static policies spend the same amount of time in each inventory state, the holding cost rate does not depend on the static arrival rate and only depends on the mean inventory state, which for our case is $(S+1)/2$ when the order-up-to level is $S$. Thus, for the static policy $(S, \lambda_{static} \cdot \mathbf{1})$, the holding cost rate is given by

$$\textsc{Hold}(S, \lambda_{static} \cdot \mathbf{1}) = \frac{h(S+1)}{2}. \tag{2.9}$$

The holding cost rate for the optimal dynamic policy depends on the demand model. However, for MHR distributions, there is a lower bound on the holding cost rate for the optimal dynamic pricing policy as described in Lemma 2.5.

**Lemma 2.5** (Lower Bound on Optimal Holding Costs)**.** *For any MHR distribution, the holding cost rate for the optimal dynamic policy* $(S^*, \boldsymbol{\lambda}^*)$ *has a lower bound*

$$\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) \geq \frac{hS^*}{\ln(eS^*)}.$$

The key observation required to prove this result is that there is a lower bound on the time that the system much spend in each inventory state for MHR distributions. A formal proof for Lemma 2.5 is provided in Appendix B.1.1. It follows from (2.9) and Lemma 2.5 that for any order-up-to level $S$, the holding cost ratio is

$$\frac{\textsc{Hold}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} \leq \frac{(1 + \ln(S^*))(S + 1)}{2S^*}. \tag{2.10}$$

**Proof of Theorem 2.2.** In Theorem 2.2, we analyze the static policy when $S = S^*$. Substituting $S = S^*$ in (2.7), (2.8) and (2.10) we get,

- Revenue ratio:

$$\frac{\textsc{Rev}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Rev}(S^*, \boldsymbol{\lambda}^*)} \geq 1.$$

- Ordering cost ratio:

$$\frac{\textsc{Order}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)} = \frac{S^*}{S^*} = 1.$$

- Holding cost ratio:

$$\frac{\textsc{Hold}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} \leq \frac{(1 + \ln(S^*))(S^* + 1)}{2S^*} \leq 1 + \ln(S^*),$$

  where the last inequality uses the fact that $S^* \geq 1$. □

**Proof of Theorem 2.3.** We prove this theorem by construction. First, note from (2.7) that the revenue ratio is still at least 1, irrespective of the new value of $S_{static}$. Consider the value of $S_{static}$ to be

$$S_{static} = \left\lceil \frac{S^*}{\sqrt{1 + \ln(S^*)}} \right\rceil. \tag{2.11}$$

Note that the total cost ratio is at most as much as the larger of the holding and ordering cost ratios. We show that both ratios are bounded above by $\sqrt{1 + \ln(S^*)}$ at the chosen value for $S_{static}$ in the following lemma. The proof of the lemma is deferred to Appendix B.1.2.

**Lemma 2.6** (Upper bound on cost ratios for $(S_{static}, \lambda_{static} \cdot \mathbf{1})$)**.** *The ordering and holding cost ratios for the static policy* $(S_{static}, \lambda_{static} \cdot \mathbf{1})$ *have the upper bound*

$$\max\left\{\frac{\textsc{Order}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)}, \frac{\textsc{Hold}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)}\right\} \leq \sqrt{1 + \ln(S^*)}.$$

From Lemma 2.6 it follows that there exists $S_{static}$ for which total cost ratio is bounded above by $\sqrt{1 + \ln(S^*)}$.

□

### 2.3.4 Tightness Analysis

This section gives examples when bounds in Theorems 2.2, 2.3, and 2.4 are tight, up to a constant factor when the bound depends on $S^*$. In Example 2.7, we show tightness of the revenue ratio bound in Theorems 2.2 and 2.3.

**Example 2.7** (Tight example for revenue ratio bound)**.** Our static policies guarantee that the revenue ratio is at least 1. This ratio is tight when there are no ordering costs, i.e., $K = 0$. The optimal inventory policy is to order one unit of product after every demand arrival, i.e., $S^* = 1$. When $S^* = 1$, the revenue ratio is exactly 1 since the optimal policy is static. ■

Next, we cover the tightness of the three cost ratios stated in our results for MHR distributions, i.e., ordering and holding cost ratios for policy $(S^*, \lambda_{static} \cdot \mathbf{1})$ (Theorem 2.2), and total cost ratio for policy $(S_{static}, \lambda_{static} \cdot \mathbf{1})$ (Theorem 2.3). For the latter two bounds, we show rightness by constructing examples whose cost ratios are within a constant factor of the stated upper bound.

**Example 2.8** (Tight example for cost ratio bounds in Theorem 2.2)**.** Since the bound for the ordering cost ratio for $(S^*, \lambda_{static} \cdot \mathbf{1})$, (2.8), is an equality, it is trivially always tight.

The holding cost ratio bound is tight when demand is exponential and the optimal profit rate is zero. Suppose the demand follows an exponential demand model and the demand function is given by $\lambda = ae^{-bp}$, with unique inverse function

$$p(\lambda) = \frac{1}{b}\ln\left(\frac{a}{\lambda}\right). \tag{2.12}$$

The optimal arrival rates are given by

$$p'(\lambda_i^*) = -\frac{1}{b\lambda} = -\frac{hi+\gamma}{(\lambda_i^*)^2} \Rightarrow \lambda_i^* = bhi + b\gamma, \tag{2.13}$$

where $\gamma$ is the optimal profit rate and $h$ is the holding cost parameter. The first equality follows from differentiating (2.12) and the second equality follows from Lemma 2.1(i). To model the costs in our tight example, we fix a positive integer $k$ and consider the following cost parameters:

$$h_k := \frac{a}{ebk}, \qquad K_k := \sum_{i=1}^{k} \frac{1}{b}\left(1 - \ln\left(\frac{i}{k}\right)\right). \tag{2.14}$$

We first show that the optimal profit rate and order-up-to level for instance $k$ are 0 and $k$, respectively. In particular, substituting values $\gamma^* = 0$ and $S^* = k$ along with the optimal rate formula from (2.13) satisfies the conditions stated in Lemma 2.1(ii) and (iii). Note that if $\gamma^* = 0$, then the optimal prices for inventory state $i \in [k]$ for this instance are given by

$$\lambda_i^* = bh_k i + b \cdot 0 = \frac{ai}{ek}, \tag{2.15}$$

where the first equality follows from (2.13) and the second equality follows from substituting $h_k$ defined in (2.14). Now, to see Lemma 2.1 (ii), note that,

$$g_{k,0}(\lambda_k^*) = p(\lambda_k^*) - \frac{h_k k}{\lambda_k^*} = \frac{1}{b}\ln\left(\frac{a}{\lambda_k^*}\right) - \frac{h_k k}{\lambda_k^*} = \frac{1}{b}\ln\left(\frac{a}{\frac{ak}{ek}}\right) - \frac{h_k k}{bh_k k} = \frac{1}{b} - \frac{1}{b} = 0,$$

where the first equality follows from (2.5), the second equality follows from (2.12), and the third equality follows from (2.15). Combining the above with the fact that $g_{i,0}(\lambda_i^*)$ is a decreasing function of $i$, it follows that $k$ is the largest integer with $g_{i,0}(\lambda_i^*) \geq 0$. Therefore, $S^* = k$ and $\gamma^* = 0$

satisfy Lemma 2.1(ii). To show Lemma 2.1(iii), note that

$$\sum_{i=1}^{k} g_{k,0}(\lambda_i^*) = \sum_{i=1}^{k}\left(\frac{1}{b}\ln\left(\frac{a}{\lambda_i^*}\right) - \frac{h_k i}{\lambda_i^*}\right) = \sum_{i=1}^{k}\frac{1}{b}\left(1 - \ln\left(\frac{i}{k}\right)\right) = K_k.$$

where the first equality follows from (2.12) and (2.5), the second equality follows from (2.15), and the last equality follows from (2.14). Thus, the optimal policy is $(k, \boldsymbol{\lambda}^*)$, and the optimal profit rate is 0 for our instance. It follows that the holding cost rate for the optimal dynamic policy is,

$$\textsc{Hold}(k, \boldsymbol{\lambda}^*) = \frac{\sum_{i=1}^{k} \frac{h_k i}{\lambda_i^*}}{\sum_{i=1}^{k} \frac{1}{\lambda_i^*}} = \frac{\frac{k}{b}}{\sum_{i=1}^{k} \frac{1}{b h_k i}} \le \frac{h_k k}{\int_1^{k+1} \frac{1}{x} dx} = \frac{h_k k}{\ln(k+1)}, \tag{2.16}$$

where the first equality follows from substituting (2.3) and (2.2) into the definition of the holding cost ratio from (2.4) and the second equality uses the fact that $\lambda_i^* = b h_k i$ for our instance. On the other hand, the holding cost ratio for static policy $(k, \lambda_{static} \cdot \mathbf{1})$ is $h_k(k+1)/2$. Therefore, the holding cost ratio for policy $(k, \lambda_{static} \cdot \mathbf{1})$ can be bounded below as,

$$\frac{\textsc{Hold}(k, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(k, \boldsymbol{\lambda}^*)} \ge \frac{k+1}{2k}\ln(k+1) \ge \frac{\ln(k)}{2}, \tag{2.17}$$

which is within a constant factor of $1 + \ln(k)$. ■

**Example 2.9** (Tightness of total cost ratio bound in Theorem 2.3)**.** We continue using the same instance as Example 2.8 and show that the best total cost ratio achievable by adjusting order-up-to level $S_{static}$ is within a multiplicative factor of the upper bound. Note that the ordering and holding cost rates decrease and increase, respectively, when the order-up-to level is increased for a static policy. It follows that the minimum possible total cost ratio occurs when both ratios are equal. Since the total cost ratio is bounded below by the minimum of holding and ordering costs, we know that for any order-up-to level $S = 1, \ldots, k$,

$$\frac{\textsc{Costs}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(k, \boldsymbol{\lambda}^*)} \ge \min\left\{\frac{k}{S}, \frac{S+1}{2k}(\ln(k+1))\right\},$$

which follows from (2.8), (2.9), and (2.16). The right-hand side is minimized when both its terms are equal, which occurs at

$$\widetilde{S} := -\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2k^2}{\ln(k+1)}}.$$

Using $\widetilde{S}$, we can lower bound the total cost ratio by

$$\begin{aligned}\frac{\textsc{Costs}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(k, \boldsymbol{\lambda}^*)} &\geq \frac{\widetilde{S}+1}{2k}(\ln(k+1)) = \frac{\ln(k+1)}{2k}\left(\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2k^2}{\ln(k+1)}}\right) \\ &\geq \frac{\ln(k+1)}{2k}\sqrt{\frac{2k^2}{\ln(k+1)}} = \sqrt{\frac{\ln(k+1)}{2}} \geq \sqrt{\frac{\ln(k)}{2}},\end{aligned}$$

which is within a constant factor of the total cost ratio derived in Theorem 2.3. ■

We next show that the analysis for linear demand is also tight. In particular, we construct a sequence of instances whose cost ratios approach the constants stated in Theorem 2.4.

**Example 2.10** (Tightness of cost ratio bounds in Theorem 2.4)**.** Suppose the demand function is linear and $\lambda(p) = a(1-bp)$. We construct our sequence of instances as follows. For every positive integer $k$, instance $k$ has cost parameters

$$h_k := \frac{a}{4bk}, \qquad K_k := \sum_{i=1}^{k} \frac{1}{b}\left(1 - \sqrt{\frac{i}{k}}\right). \tag{2.18}$$

Following the same line of analysis that we perform in Examples 2.8 and 2.9 for the exponential demand, we prove the following facts for instance $k$ when demand is linear.

**Lemma 2.11.** *Instance k with cost parameters* $(h_k, K_k)$ *has the following properties when demand is linear.*

*(i) The optimal order-up-to level and profit for instance k are k and* 0*, respectively.*

*(ii) The holding cost ratio for policy* $(k, \lambda_{static} \cdot \mathbf{1})$ *can be bounded below as*

$$\frac{\textsc{Hold}(k, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(k, \boldsymbol{\lambda}^*)} \geq \frac{3}{2}\left(1 - \frac{1}{\sqrt{k+1}}\right).$$

*(iii) The total cost ratio for policy* $(S, \lambda_{static} \cdot \mathbf{1})$ *for* $S = 1, \ldots, k$ *is bounded below as*

$$\frac{\textsc{Costs}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(k, \boldsymbol{\lambda}^*)} \geq \sqrt{\frac{3k}{2(k+1)}\left(1 - \frac{1}{\sqrt{k+1}}\right)},$$

The proof for Lemma 2.11 is deferred to Appendix B.1.4. As $k \to \infty$, it follows from Lemma 2.11(ii) and (iii) that holding and total cost ratio lower bounds approach 1.5 and $\sqrt{1.5}$, respectively. Therefore, the cost ratios from Theorem 2.4 are tight for this sequence of instances. ■

## 2.4 Backlogging Model

In this section, we consider the setting where unsatisfied demand from purchasing customers are backlogged. We assume a shortage cost of $\pi$ per unit of backlogged demand per unit time. The optimal inventory policy shifts from a ZIO policy to an $(s, S)$ policy where $s \leq 0$. The price is still only a function of the inventory state, and $\lambda_i$ is the (chosen) demand arrival rate at inventory state $i$, for $i = s+1, \ldots, S$. We first show that the cost ratio can be arbitrarily bad in this setting for a static pricing policy, even when the demand is linear. Later, we give the formulation for a 3-price policy with theoretical guarantees.

**Example 2.12** (Failure of static pricing)**.** Consider a setting where the demand function is linear with the form $\lambda(p) = a(1 - bp)$. Let the holding and backlogging costs be symmetric, i.e., $\pi = h$. We consider a sequence of instances $(\pi_k, h_k, K_k)$ with $\pi_k = h_k = \frac{a}{4b} - \gamma_k$ and optimal profit rate $\gamma_k$ for a sequence of positive numbers $(\gamma_k)_{k \in \mathbb{N}}$ such that $\gamma_k \to 0$. From Lemma 2.1, one can check that setting $K_k = \sum_{i=-1}^{1} g_{i,\gamma_k}(\lambda_i^*)$ achieves the required optimal profit rate.

The optimal inventory policy for every instance is the same, with $S^* = 1$ and $s^* = -2$, i.e., there are only three inventory positions which are 1, 0, and -1. Moreover, the optimal arrival rates are given by

$$\lambda_{-1}^k = \frac{a}{2}, \quad \lambda_0^k = \sqrt{ab\gamma_k}, \quad \lambda_1^k = \frac{a}{2}.$$

Therefore, the holding cost rate for the optimal dynamic pricing policy is

$$\textsc{Hold}(s^*, S^*, \boldsymbol{\lambda}^*) = \frac{\frac{2\pi}{a} + \frac{2h}{a}}{\frac{2}{a} + \frac{2}{a} + \frac{1}{\sqrt{ab\gamma_k}}} = \frac{4h\sqrt{b\gamma_k}}{4\sqrt{b\gamma_k} + \sqrt{a}},$$

which goes to 0 as $k$ increases since $\gamma_k \longrightarrow 0$. On the other hand, the holding cost rate for any static policy with arrival rate $\lambda_{static}$ is

$$\textsc{Hold}(s^*, S^*, \lambda_{static} \cdot \mathbf{1}) = \frac{\frac{\pi}{\lambda_{static}} + \frac{h}{\lambda_{static}}}{\frac{3}{\lambda_{static}}} = \frac{2h}{3},$$

which is a constant independent of the choice of $\lambda_{static}$. We conclude that the holding/backordering cost ratio diverges to infinity for static policies in the backlogging setting when the profit rate approaches 0. ■

Intuitively, when the profit rate is low, it is optimal to offer very high prices when there is no available inventory, i.e., in state 0, since the holding/backordering costs incurred in this state are 0. Consequently, the optimal dynamic policy spends large amounts of time in state 0. However, any static policy, by construction, will spend the same amount of time in every inventory state and is thus at a disadvantage. This problem does not arise in the lost sales model, since it was not possible to sell the product in state 0, and thus a static price performs reasonably well.

### 2.4.1 Simple pricing policy for backlogging

While static pricing may not be a good idea in the backlogging setting, we still want to find a simple pricing policy with a few prices and strong performance guarantees. In this section, we construct a policy that uses *three prices*, which provides theoretical guarantees in the backlogging paradigm. In this policy, we offer a different price for each of the following three regimes: negative inventory, zero inventory, and positive inventory. We denote the three corresponding arrival rates by $\lambda_-$, $\lambda_\circ$, and $\lambda_+$, respectively. The policy construction is similar to the static policy in Section

Figure 2.1: Three-price policy for backlogging ($s^* = -10, S^* = 200$)

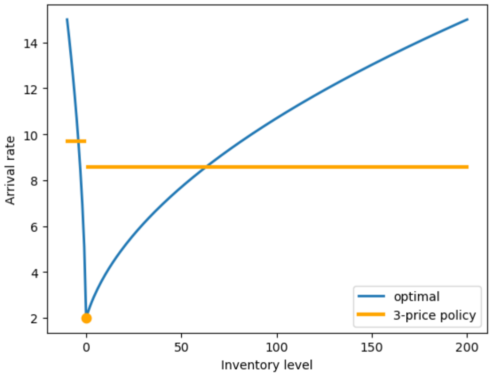

2.3.1. Suppose $(s^*, S^*, \lambda^*)$ is the optimal dynamic policy, then the static rates are given by

$$\lambda_+ = \mathbb{E}[\lambda_i^* \mid i > 0] = \frac{S^*}{\sum_{i=1}^{S^*} \frac{1}{\lambda_i^*}}, \qquad \lambda_\circ = \lambda_0^*, \qquad \lambda_- = \mathbb{E}[\lambda_i^* \mid i < 0] = \frac{-s^* - 1}{\sum_{i=s^*+1}^{-1} \frac{1}{\lambda_i^*}}.$$

We refer to the policy as $\lambda_{3p}$. Figure 2.1 shows an example of the arrival rates under policy $\lambda_{3p}$. The following theorem states the guarantees for policy $\lambda_{3p}$.

**Theorem 2.13.** *Let $(s^*, S^*, \lambda^*)$ be the optimal dynamic policy for an instance where customer valuations follow an MHR distribution. Then the three price policy $(s^*, S^*, \lambda_{3p})$ has the following guarantees:*

- *Revenue ratio:*

$$\frac{\textsc{Rev}(s^*, S^*, \lambda_{3p})}{\textsc{Rev}(s^*, S^*, \lambda^*)} \geq 1.$$

- *Ordering cost ratio:*

$$\frac{\textsc{Order}(s^*, S^*, \lambda_{3p})}{\textsc{Order}(s^*, S^*, \lambda^*)} = 1.$$

- *Combined holding and backlogging cost ratio:*

$$\frac{\textsc{Hold}(s^*, S^*, \lambda_{3p})}{\textsc{Hold}(s^*, S^*, \lambda^*)} \leq 1 + \ln(S^* - s^*).$$

We defer the proof of Theorem 2.13 to Appendix B.1.5. With just 3 prices compared to $S^* - s^*$

prices, we can achieve more revenue, same ordering costs, and at most a factor of $1 + \ln(S^* - s^*)$ in combined holding and backlogging costs.

## 2.5 Numerical Experiments

In this section, we numerically show the effectiveness of static policies for joint inventory and pricing problems. In this numerical study, we analyze the performance of *static policies* in the *lost-sales model* setting for two specific examples of MHR distributions: *linear* and *exponential* demand. We consider three policies to compare to the optimal policy.

- **Our proposed policy for proving guarantees:** We consider the policy from Theorem 2.2, $(S^*, \lambda_{static} \cdot \mathbf{1})$. While this static arrival rate make analysis more amenable, for practical purposes one is more likely to optimize the static arrival rate (price).
- **Profit-maximizing static policy with order-up-to level $S^*$:** This policy keeps the inventory policy to be the same as the optimal dynamic policy, and uses the static rate that maximizes profit. We let $\lambda^*_{static}(S)$ denote the profit-maximizing rate for order-up-to level $S$. Therefore, this policy is referred to as the $(S^*, \lambda^*_{static}(S^*) \cdot \mathbf{1})$ policy.
- **Profit maximizing static policy:** This policy optimizes both the static rate, as well as the order-up-to level to maximize profit. we refer to the optimal order-up-to level as $S^*_{static}$. Therefore, this policy is referred to as $(S^*_{static}, \lambda^*_{static}(S^*_{static}) \cdot \mathbf{1})$.

### 2.5.1 Methodology

For the test bed, we fix the demand function and consider 100 instances with different values of the ordering cost $K$ and holding cost $h$ in a $10 \times 10$ grid where $K = 500, 1000, \ldots, 5000$ and $h = 0.5, 1.0, \ldots, 5.0$. For the demand functions, we let the maximum arrival rate be 10, i.e., $\Lambda = 10$, and we let the expected valuation of the customer be 100, i.e., $\mathbb{E}[V] = 100$. Accordingly, the demand functions for linear and exponential demand come out to be $\lambda(p) = 10 - p/20$, and $\lambda(p) = 10e^{-p/100}$, respectively.

Figure 2.2: Variation of profit ratio and connection to optimal dynamic profit (linear demand)

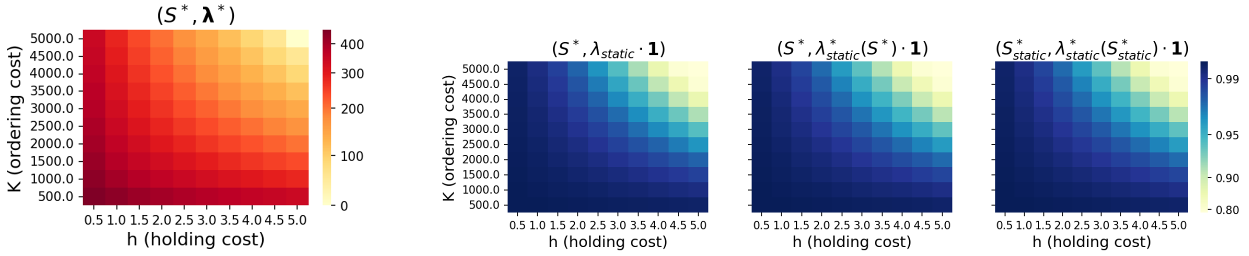

(a) Optimal dynamic profit

(b) Profit ratio

*Note:* Figure 2.2a shows the optimal dynamic profit for each instance, whereas Figure 2.2b shows the profit ratio for the different static policies we consider. For the policies from left to right, we observe average profit ratios of 0.9689, 0.9693, and 0.9703, respectively.

Figure 2.3: Variation of profit ratio and connection to optimal dynamic profit (exponential demand)

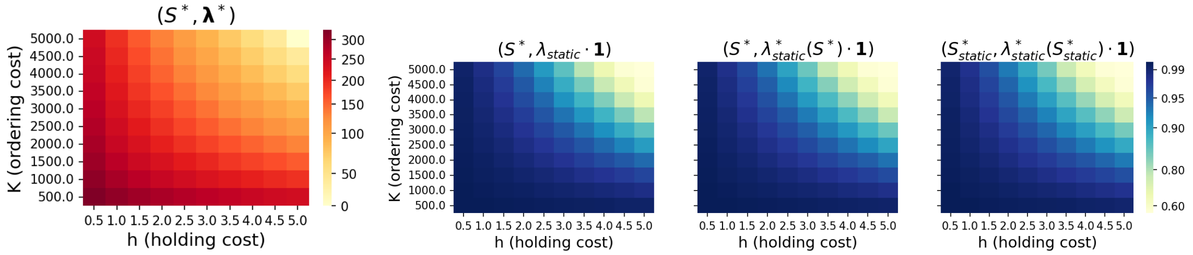

(a) Optimal dynamic profit

(b) Profit ratio

*Note:* Figure 2.3a shows the optimal dynamic profit for each instance, whereas Figure 2.2b shows the profit ratio for the different static policies we consider. For the policies from left to right, we observe average profit ratios of 0.9128, 0.9171, and 0.9249, respectively.

Figure 2.4: Profit ratio vs. profit margin

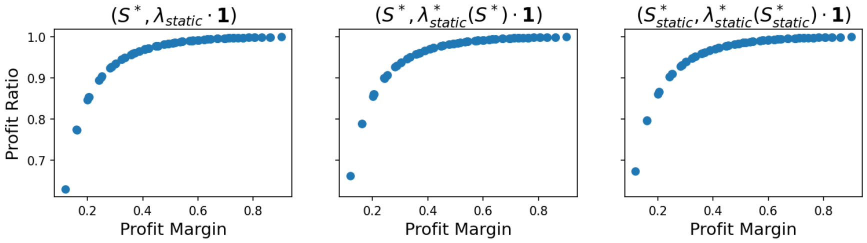

(a) Linear demand

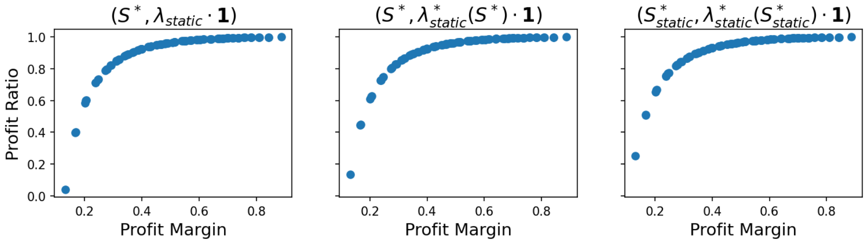

(b) Exponential demand

### 2.5.2 Results and analysis

Our primary goal from the numerical experiments is to compare the performance of the static policies directly in terms of long-term average profit. With that in mind, the metric that we mainly focus on is *profit ratio*. Figures 2.2b and 2.3b show how the profit ratio varies with instance parameters for linear and exponential demand, respectively. It can be seen that all three static policies capture a large fraction of the optimal dynamic profit in most instances. In particular, the profit ratio is much greater than 0.9 for a large majority of the linear demand instances and 0.8 for some exponential demand instances. Figures 2.2a and 2.3a give the heatmap for the optimal dynamic profit rates for these instances, and one can see that the instances where the profit ratio is low coincide with the instances where the profit is low. This observation is illustrated better in Figures 2.4a and 2.4b, where the profit ratio for the different static policies is plotted against the profit margin (i.e., profit divided by revenue) for the optimal dynamic policy.

In terms of comparison of performance across the different static policies, all policies perform similarly to each other in terms of profits. However, in terms of individual metrics such as revenue, holding costs, and ordering costs, the policies perform differently due to the differences in their constructions. These differences are illustrated in Figures 2.5 and 2.6.

Figure 2.5: Variation in revenue and cost ratios for linear demand

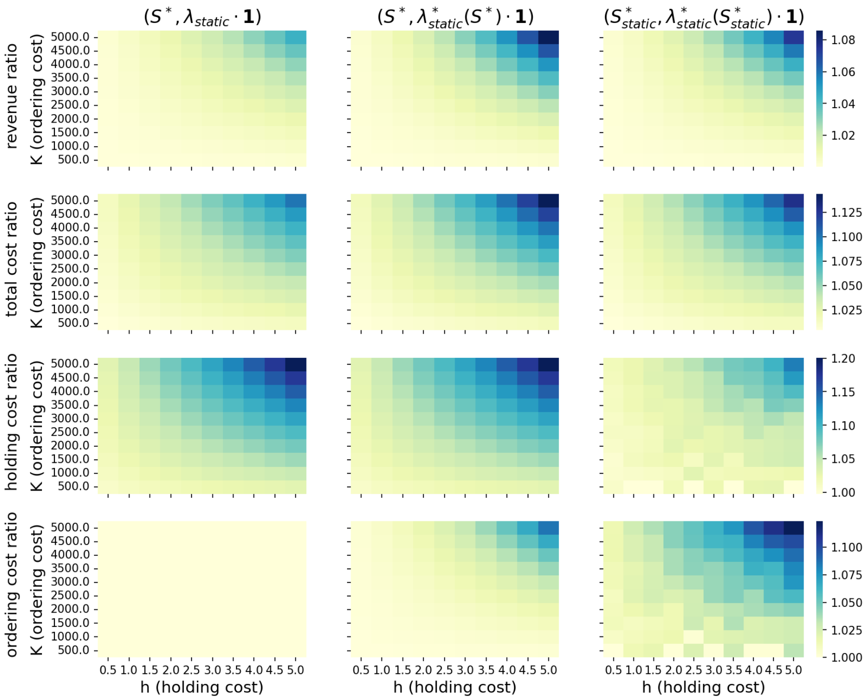

Figure 2.6: Variation in revenue and cost ratios for exponential demand

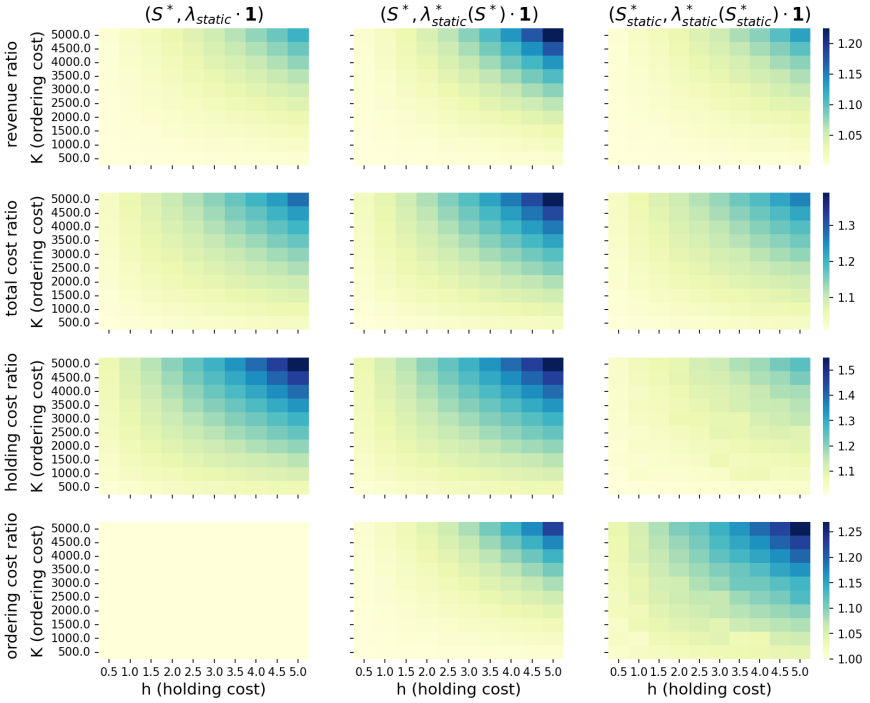

## 2.6 Conclusion

In this work, we study the performance of simple policies for joint pricing and inventory control problems with continuous review and a general class of distributions that are MHR. For our analysis, we break the profit rate down into revenue and costs, and give multiplicative guarantees for each component. Our proofs rely on a particular construction for our simple policies that directly use the optimal dynamic policy, in contrast to heuristics that rely on fluid limits. When there is no backlogging, we prove that a static policy is sufficient for good performance guarantees. We prove that revenue actually increases, while the total costs are at most $\sqrt{1 + \ln(S^*)}$ factor from optimal and at most 1.225 factor when demand is linear. When backlogging is allowed, a policy that uses only three prices has similar performance guarantees. In numerical experiments, we explicitly compare the profit rates of simple policies with the optimal dynamic policies. Simple policies capture a large fraction of the profit in most cases. For the instances where the profit ratio is low, the profit margin is low as well.

This work provides a foundation for potentially many avenues for results in more general settings. Future research directions include studying whether simple, low-dimensional policies can be provably near-optimal for periodic review settings, multiple products, multiple echelons, and problems with lead times.

# Chapter 3: Pricing and Assortment Optimization Considerations with Opaque Products

## 3.1 Introduction

In this chapter, we study the classic problems of assortment optimization and pricing from revenue management literature under the opaque selling mechanism. We use the term opaque selling to refer to instances where sellers providing substitutable products that differ on some features (for instance, shirts of different sizes, flights at different times of day) offer a "surprise bag" style of product (i.e., an opaque product) for their more flexible customers. Such mechanisms have been employed in the airline, hospitality, and retail industries [61, 62]. Offering opaque products allows sellers to boost their sales using differential pricing for their flexible and inflexible customer bases. It also allows the seller to manage their product inventories more efficiently. Both these aspects have been studied through several lenses in literature. This chapter studies the pricing and assortment optimization problem for the multinomial logit (MNL) choice model.

We consider a seller offering substitutable products, which are similar but differentiated along some features. When a customer arrives, the seller chooses an assortment of these products to offer to the customer. The product prices may or may not be exogenous. The customer may choose to buy one of the offered products outright or pay a discounted price for the opaque product option. If the customer chooses the opaque product, the seller may fulfill the purchase with any of the offered products as chosen by the seller. To differentiate the two different product types a customer can choose from, we call them traditional and opaque products. An opaque product purchase is fulfilled by one of the traditional products in the offered assortment.

A key question to address when modeling the opaque mechanism is determining a customer's value for the opaque product. Several early works have studied the idea of "probabilistic selling" where the customer has prior knowledge of the probability of receiving every traditional product when purchasing the opaque product, and thus the valuation of the opaque product being the expected value over the possible alternatives. The other approach to modeling customer valuations is to assume the customer has no information beforehand on what product they may receive with the opaque option. If customers are *risk-averse*, they are assumed to value the opaque product as if it were their least valued product in the assortment. On the other hand, if they are risk-neutral, they are assumed to expect to receive a product uniformly from the assortment, which is the same idea as probabilistic selling with a uniform distribution. This chapter focuses on the risk-averse customer model.

### 3.1.1 Our contributions

- For risk-averse MNL customers, we provide closed-form solutions for purchase probabilities. The purchase probability for a product under opaque selling can be described as a linear combination of traditional MNL purchase probabilities corresponding to exponentially many assortments. In each assortment, while the same products are offered each time, the prices of a subset are changed to the opaque price.

- We show that the revenue-maximizing assortment may no longer have the nested-by-revenue structure that is observed for MNL customers under traditional selling. Instead, we propose heuristics with constant factor approximation guarantees. We first show that offering the better of the optimal traditional assortment and the best singleton assortment guarantees half the optimal revenue for risk-averse customers. We then extend this idea to a nested-by-revenue-and-valuation heuristic, which retains the performance guarantee and is near-optimal in practice.

- For the price optimization problem, we empirically observe that opaque selling does not do better than traditional selling. We prove these results for risk-averse MNL customers in two cases: (1) when prices are uniform and (2) when there are two products.

- In our numerical experiments, we find that assortment optimization heuristics significantly outperform the proven guarantees. In particular, the nested-by-revenue-and-valuation heuristic did not give the optimal assortment in only 5/11350 instances.

### 3.1.2 Literature review

The work by Elmachtoub and Hamilton [62] is most similar to ours in the literature. In this work, the authors analyze pricing under the opaque selling mechanism for exchangeable valuation distributions, while our work focuses on the multinomial logit choice model. They consider both risk-averse and risk-neutral customers in their analyses. Opaque selling for risk-neutral customers has also been studied under the name of probabilistic selling in the literature [63, 64, 65, 66]. These papers focus on settings with two or three products with customer valuations drawn from the Hotelling or Salop's circle choice models and show that there is value in offering opaque products when different customer segments value products differently. Strauss et al. [67] also consider the pricing problem with opaque products, particularly in the setting of flexible time slots. They assumed a nested logit model for customer valuations, with traditional and opaque products belonging in separate nests.

Another line of literature studies the power of opaque selling for capacity and inventory management. Gallego and Phillips [61] first considered revenue management with "flexible products" in their work. They studied a model with two products where the seller can defer the decision of allocating a traditional product to customers purchasing a flexible product to a predefined time after the purchase. Gallego et al. [68] extend the concept of flexible products to networks. They also consider choice models in their work and provide algorithms for assortment and allocation planning with asymptotic guarantees when the customer valuations are drawn from general "attraction" models (which encapsulate the MNL choice model). Their work differs from ours as they do not assume the opaque or flexible product's valuation to be a deterministic function of the traditional products and instead allow for it to have a random valuation. Similarly, Liu et al. [69] also provide an algorithm for assortment and allocation planning based on the fluid relaxation of the problem.

In their model, the allocation of a traditional product to customers purchasing an opaque product is done at the time of purchase. Their policy relies on the fluid relaxation linear program being solvable. Elmachtoub et al. [70] and Elmachtoub et al. [71] also study the value of opaque selling in inventory management settings.

## 3.2 Preliminaries

We consider a seller with a set $\mathcal{N} := \{1, \ldots, n\}$ of substitutable products, where $\boldsymbol{r} = \{r_i \mid i \in \mathcal{N}\}$ denotes the vector of product prices. These products are referred to as traditional products to contrast them from the opaque product mechanism. The prices are assumed to be exogenously given or endogenous depending on whether we talk about assortment or price optimization respectively. We model customer behavior using random utility models. Accordingly, the notation $\pi(i \mid \textsc{Mech})$ is used to denote the probability that the customer purchases product $i$ when shown assortment $S$ under mechanism $\textsc{Mech}$. The expected revenue earned by the mechanism is denoted by,

$$\textsc{Rev}(\textsc{Mech}) := \sum_{i=1}^{n} r_i \pi(i \mid \textsc{Mech}).$$

We consider two mechanisms in this work, the traditional selling mechanism, denoted by $\textsc{Trad}$ and the opaque selling mechanism, denoted by $\textsc{Opq}$. In the most general sense, a mechanism is parameterized using an assortment $S$ and a vector of prices for the traditional products $\boldsymbol{r}_S := \{r_i \mid i \in S\}$, while also including some auxiliary parameters, such as the opaque price for the opaque selling mechanism. For instance, the traditional selling mechanism where assortment $S$ and prices $\boldsymbol{r}_S$ are offered is denoted by $\textsc{Trad}(S, \boldsymbol{r}_S)$. However, to keep notation concise, we drop parameters that remain unchanged as applicable, For instance, for the price optimization problem, since the assortment is assumed to be $\mathcal{N}$, we refer to the traditional selling mechanism as $\textsc{Trad}(\boldsymbol{r})$.

### 3.2.1 Traditional selling mechanism

We assume that customer valuations for traditional products are Logistic random variables. The vector of random variables $(V_i)_{i=1}^n$ denotes a customer's valuation for the traditional products with $V_i \sim Logistic(v_i, 1)$, modeled as follows:

$$V_i = v_i + \epsilon_i - \epsilon_0.$$

Here, $(\epsilon_i)_{i=0}^n$ are i.i.d. random variables with distribution $Gumbel(0, 1)$. Note that due to the assumption of independence and the Gumbel distribution, $\epsilon_i - \epsilon_0 \sim Logistic(0, 1)$, which implies that the valuations are Logistic random variables with mean $v_i$. These modeling assumptions are in line with the well-known multinomial logit (MNL) choice model for modeling customer purchase decisions. The vector, $(U_i)_{i=1}^n$ denotes the utilities of the traditional products. Since the utility is (valuation - price), its distribution is given as $U_i \sim Logistic(v_i - r_i, 1)$. A customer also always has the option to not make a purchase. This no-purchase option is denoted as product 0, with utility $U_0 = 0$. In the traditional assortment selling mechanism, the seller selects a subset $S \subseteq \mathcal{N}$ of products to offer to the customer, and then the customer purchases the product that maximizes their utility. For the MNL choice model, the probability that a customer purchases product $i$ under mechanism $\textsc{Trad}(S, \boldsymbol{r}_S)$ is given by:

$$\pi(i \mid \textsc{Trad}(S, \boldsymbol{r}_S)) = \mathbb{P}\big[U_i > U_j \forall j \in S \cup \{0\} \setminus \{i\}\big] = \frac{e^{v_i - r_i}}{1 + \sum_{j \in S} e^{v_j - r_j}}. \tag{3.1}$$

Here, we assume without loss of generality that $i \in S \cup \{0\}$. If $i \notin S$, $\pi(i \mid \textsc{Trad}(S, \boldsymbol{r}_S)) = 0$. Further, the expected revenue earned by the traditional mechanism is,

$$\textsc{Rev}(\textsc{Trad}(S, \boldsymbol{r}_S)) = \sum_{i \in S} r_i \pi(i \mid \textsc{Trad}(S, \boldsymbol{r}_S)) = \frac{\sum_{i \in S} r_i e^{v_i - r_i}}{1 + \sum_{i \in S} e^{v_i - r_i}} \tag{3.2}$$

### 3.2.2 Opaque selling mechanism

The seller may also choose to offer an opaque product, denoted by $q$, at price $\rho$. If the customer purchases the opaque product, the seller can choose to satisfy this purchase with any of the traditional products that was offered to the customer. In their work, Elmachtoub and Hamilton [62] consider two approaches to model the customer's valuation of the opaque product, risk-averse and risk-neutral. If $S$ is the assortment of traditional products that the customer is offered, the valuation of the opaque product based on their definitions is

1. **Risk Averse:** $V_q = \min_{i \in S} V_i$.
2. **Risk Neutral:** $V_q = \frac{1}{|S|} \sum_{i \in S} V_i$.

In this work, we assume customers are risk-averse, i.e., they assume the worst case when deciding to purchase the opaque product. The customer purchases the opaque product if its utility $U_q := V_q - \rho$ is larger than that of all traditional products and the no-purchase option. We refer to the mechanism of offering assortment $S$ of traditional products at prices $\boldsymbol{r}_S$ along with the opaque product option at price $\rho$ as $\textsc{Opq}(S, \boldsymbol{r}_S, \rho)$. The mechanism where the opaque price is endogenously optimized is referred to as $\textsc{Opq}^*(S, \boldsymbol{r}_S)$, where the opaque price is $\rho^*(S, \boldsymbol{r}_S)$.

**Risk-averse customers.**

If customers are risk-averse, they assume the worst-case scenario, i.e., the opaque product purchase will be satisfied with their least preferred product. We assume that $\rho \leq \min_{i \in S} r_i$. We can make this assumption without loss of generality since the valuation of the opaque product is the lowest among all products by definition always, and consequently, the opaque product maximizes a customer's utility only if its price is lower than the price of all products.

When $\rho \leq \min_{i \in S} r_i$, the purchase probabilities of products admit closed form solutions under the risk-averse customer assumption. Fix traditional product $i \in S$, the purchase

$$\begin{aligned}
\pi(i \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) &= \mathbb{P}\big[U_i > U_j \forall j \in S \setminus \{i\}, U_i > U_0, U_i > U_q\big] \\
&\overset{(a)}{=} \mathbb{P}\left[U_i > U_j \forall j \in S \setminus \{i\}, U_i > U_0, U_i > \min_{k \in S} V_k - \rho\right] \\
&\overset{(b)}{=} \mathbb{P}\left[U_i > U_j \forall j \in S \setminus \{i\}, U_i > U_0, \bigcup_{k=1}^{n} U_i > V_k - \rho\right] \\
&\overset{(c)}{=} \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \mathbb{P}\big[U_i > U_j \forall j \in S \setminus \{i\}, U_i > U_0, U_i > V_k - \rho \forall k \in I\big] \\
&\overset{(d)}{=} \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \mathbb{P}\big[U_i > U_j \forall j \in S \setminus (I \cup \{i\}), U_i > U_0, U_i > V_k - \rho \forall k \in I\big] \\
&\overset{(e)}{=} \sum_{\emptyset \neq I \subseteq S \setminus \{i\}} (-1)^{|I|+1} \mathbb{P}\big[U_i > U_j \forall j \in S \setminus (I \cup \{i\}), U_i > U_0, U_i > V_j - \rho \forall j \in I\big] \\
&\overset{(f)}{=} \sum_{\emptyset \neq I \subseteq S \setminus \{i\}} (-1)^{|I|+1} \pi\Big(i \mid \textsc{Trad}(S, \boldsymbol{r}_S^{I,\rho})\Big) \\
&\overset{(g)}{=} \sum_{\emptyset \neq I \subseteq S \setminus \{i\}} (-1)^{|I|+1} \frac{e^{v_i - r_i}}{1 + \sum_{j \in S \setminus I} e^{v_j - r_j} + \sum_{j \in I} e^{v_j - \rho}},
\end{aligned}$$

where,

$$r_{S,i}^{I,\rho} = \begin{cases} r_i & i \in S \setminus I \\ \rho & i \in I. \end{cases} \tag{3.3}$$

Here, (a) follows from the definition of $V_q$ for risk-averse customers, (b) uses the fact that it is sufficient to be greater than at least one element of a set to be greater than the minimum element of the set, (c) follows from the inclusion-exclusion principle, (d) follows from the fact that $V_k - \rho > U_k$ for all $k \in I$ due to our assumption that $\rho < \min_{i \in S} r_i$, (e) follows from the fact that the probability is 0 if $i \in I$ since $V_i - \rho > U_i$, (f) follows from the definition of purchase probabilities and (3.3) whereas (g) follows from the definition of MNL purchase probabilities in (3.1). Similarly for the no-purchase option,

$$\pi(0 \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) = \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \frac{1}{1 + \sum_{j \in S \setminus I} e^{v_j - r_j} + \sum_{j \in I} e^{v_j - \rho}}.$$

Finally, for the opaque product, the purchase probability is given by,

$$\begin{aligned}
\pi(q \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) &= 1 - \sum_{i \in S \cup \{0\}} \pi\big(i \mid S_q\big) \\
&= 1 - \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \frac{1 + \sum_{j \in S \setminus I} e^{v_j - r_j}}{1 + \sum_{j \in S \setminus I} e^{v_j - r_j} + \sum_{j \in I} e^{v_j - \rho}} \\
&= \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} - \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \frac{1 + \sum_{j \in S \setminus I} e^{v_j - r_j}}{1 + \sum_{j \in S \setminus I} e^{v_j - r_j} + \sum_{j \in I} e^{v_j - \rho}} \\
&= \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \frac{\sum_{j \in I} e^{v_j - \rho}}{1 + \sum_{j \in S \setminus I} e^{v_j - r_j} + \sum_{j \in I} e^{v_j - \rho}},
\end{aligned}$$

where the second equality follows from the fact that a fixed index set $I$ does not appear in the summation in the purchase probability expansion for product $i$ if and only if $i \in I$, the third equality follows from the fact that

$$\sum_{I \subseteq S} (-1)^{|I|+1} = \sum_{I \subseteq S} (-1)^{|I|+1} = 0.$$

With the probability of purchase of each product defined, the revenue earned from the opaque selling mechanism when customers are risk-averse is,

$$\begin{aligned}
\textsc{Rev}(\textsc{Opq}(S, \boldsymbol{r}_S, \rho)) &= \rho \pi(q \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) + \sum_{i \in S} r_i \pi(i \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) \\
&= \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \frac{\sum_{i \in S \setminus I} r_i e^{v_i - r_i} + \rho \sum_{i \in I} e^{v_i - \rho}}{1 + \sum_{i \in S \setminus I} e^{v_i - r_i} + \sum_{i \in I} e^{v_i - \rho}} \\
&= \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1} \textsc{Rev}\Big(\textsc{Trad}(S, \boldsymbol{r}_S^{I,\rho})\Big).
\end{aligned} \tag{3.4}$$

From the above equation, it can be seen that the revenue earned from the opaque selling mechanism for risk-averse customers is a function of the revenues earned from traditional assortment mechanisms where a subset of products are offered at the opaque price.

## 3.3 Assortment optimization

In this section, we analyze the assortment optimization problem for risk-averse MNL customers. This problem is challenging for several reasons. First, it should be noted that despite having closed form solutions for the purchase probabilities, evaluating the expected revenue for even a single assortment takes $O(n2^n)$ operations. Another challenge is determining the optimal opaque price for a given assortment. In all our experiments, find that the revenue function under opaque selling is unimodal with respect to the opaque price. This claim is challenging to prove formally, but for the sake of our algorithms and experiments, we assume that unimodality holds and the opaque price can be computed using polynomially many calls to the revenue function.

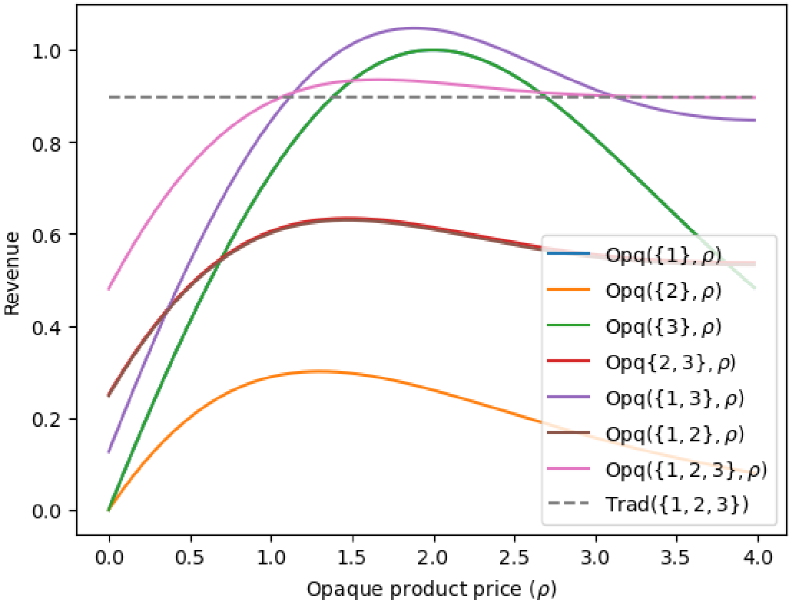

Figure 3.1: Example for risk-averse customers where the optimal assortment under the opaque selling mechanism is not nested-by-revenue. The example has three products, with $\boldsymbol{r} = (4.02, 4.01, 4)$ and $\boldsymbol{v} = (2, 0.1, 2)$. The plots are generated for each assortment by varying the price of the opaque product. The optimal assortment is to offer products $\{1, 3\}$ along with the opaque product, which is not nested-by-revenue since product 2 is excluded from the assortment.

Given these challenges, we focus on analyzing the structure of the optimal assortment, and propose heuristics that perform well while evaluating only a small set of candidate assortments. It is known that assortment optimization is tractable under the MNL choice model for traditional selling mechanism. In particular, the revenue maximizing assortment is **nested-by-revenue**, i.e., it offers all products whose price is above a certain revenue threshold. Consequently, only $n$ candidate assortments need to be evaluated to find the revenue maximizing assortment. However, the nested-by-revenue property does not hold for the revenue maximizing assortment under opaque selling. In Figure 3.1, we give an example with 3 products, where it is optimal to offer $\{1, 3\}$ along with the opaque product, and exclude product 2 from the assortment despite it having a higher revenue than product 3 for risk-averse customers. The idea behind the nested-by-revenue property not holding is that while adding a product with high revenue but low valuation to the assortment did no harm under the traditional mechanism, it can be greatly detrimental for opaque selling, since the valuation of the opaque product is impacted by such a product.

### 3.3.1 Traditional optimal or singleton (TOS) heuristic

Our first heuristic is a simple algorithm that builds on the nested-by-revenue idea. In particular, we consider two types of candidate assortments: (1) the revenue-maximizing traditional assortment and (2) all singleton assortments. This heuristic has the advantage of being very efficient computationally. Both finding the revenue-maximizing singleton assortment under opaque selling, and revenue-maximizing assortment under traditional selling can be performed efficiently. Consequently, the only computationally heavy step for this heuristic is computing the optimal opaque price and revenue for the revenue-maximizing traditional assortment under opaque selling. Further, we show that this heuristic guarantees at least half the revenue of the revenue maximizing assortment under opaque selling.

**Theorem 3.1.** *Let $S^*$ and $\hat{S}$ be the revenue maximizing assortments under opaque and traditional selling, respectively, then*

$$\max\left\{\max_{i\in\mathcal{N}} \textsc{Rev}(\textsc{Opq}^*(\{i\})), \textsc{Rev}\left(\textsc{Opq}^*\left(\hat{S}\right)\right)\right\} \geq \frac{1}{2}\textsc{Rev}(\textsc{Opq}^*(S^*)).$$

*Proof.* Note that the optimal opaque selling revenue can be written as,

$$\textsc{Rev}(\textsc{Opq}^*(S^*)) = \rho^*(S^*)\pi(q \mid \textsc{Opq}^*(S^*)) + \sum_{i\in S^*} r_i\pi(i \mid \textsc{Opq}^*(S^*)).$$

We will proceed by bounding both terms above with different quantities. First, we bound the first term above as follows:

$$\begin{aligned} \rho(S^*)\pi(q \mid \textsc{Opq}^*(S^*)) &\le \max_{i\in S^*} \rho(S^*)\pi(q \mid \textsc{Opq}(\{i\}, \rho^*(S^*))) \\ &\le \max_{i\in \mathcal{N}} \textsc{Rev}(\textsc{Opq}(\{i\}, \rho^*(S^*))) \\ &\le \max_{i\in \mathcal{N}} \textsc{Rev}(\textsc{Opq}^*(\{i\})), \end{aligned} \tag{3.5}$$

where the first inequality follows from the fact that for $S' \subseteq S$, $\pi(q \mid \textsc{Opq}(S, \boldsymbol{r}_S, \rho)) \le \pi(q \mid \textsc{Opq}(S', \boldsymbol{r}_{S'}, \rho))$, and the third inequality follows from the optimality of $\rho(\{i\})$. Next, we bound the second term above as follows:

$$\begin{aligned} \sum_{i\in S^*} r_i\pi(i \mid \textsc{Opq}^*(S^*)) &\le \sum_{i\in S^*} r_i\pi(i \mid \textsc{Trad}(S^*)) = \textsc{Rev}(\textsc{Trad}(S^*)) \\ &\le \textsc{Rev}\left(\textsc{Trad}\left(\hat{S}\right)\right) \le \textsc{Rev}\left(\textsc{Opq}^*\left(\hat{S}\right)\right), \end{aligned} \tag{3.6}$$

where the first inequality follows from substitutability, and the second inequality follows from the definition of $\hat{S}$. Combining (3.5) and (3.6) completes the proof:

$$\begin{aligned} \textsc{Rev}(\textsc{Opq}^*(S^*)) &= \rho^*(S^*)\pi(q \mid \textsc{Opq}^*(S^*)) + \sum_{i\in S^*} r_i\pi(i \mid \textsc{Opq}^*(S^*)) \\ &\le \max_{i\in \mathcal{N}} \textsc{Rev}(\textsc{Opq}^*(\{i\})) + \textsc{Rev}\left(\textsc{Opq}^*\left(\hat{S}\right)\right) \\ &\le 2\max\left\{\max_{i\in \mathcal{N}} \textsc{Rev}(\textsc{Opq}^*(\{i\})), \textsc{Rev}\left(\textsc{Opq}^*\left(\hat{S}\right)\right)\right\}. \end{aligned}$$

□

Note that this proof does not use any specific properties of the MNL choice model other than substitutability, and hence applies to every substitutable choice model (for instance, every random utility (RUM) choice model). While the TOS heuristic is easy to compute and gives very good approximations to the optimal opaque assortment in practice (refer to experimental section for performance details), it is not optimal. One sub-optimal example is the one we considered previously in Figure 3.1. In this example, the optimal opaque assortment, $\{1, 3\}$ is neither the optimal traditional assortment, $\{1, 2, 3\}$, nor the best singleton assortment, $\{3\}$. The reasoning for this is the low valuation for product 2, which negatively impacts the opaque product valuation while not adding to the revenue through traditional sales. This motivates the idea for the next heuristic.

### 3.3.2 Nested by revenue and valuation (NRV) heuristic

The second heuristic we consider builds on the nested-by-revenue heuristic by adding in a valuation threshold on top of the revenue threshold. Consequently, this heuristic increases the cardinality of the set of candidate assortments to $O(n^2)$. While assortment $\{1, 3\}$ in the earlier example is not nested-by-revenue, it is indeed nested by revenue and valuation. In fact, in experiments, we see that this heuristic is optimal in a wide majority of instances.

One aspect to consider for the NRV heuristic is tie-breaks. In the traditional selling mechanism, tie-breaks are not necessary, since if $r_i = r_j$ for two products $i$ and $j$, then $j$ is in the optimal assortment if $i$ is in the optimal assortment. However, this is not always the case for opaque selling, even for identical products.

**Example 3.2** (Adding identical products increases revenue)**.** Consider a setting with two products where $\boldsymbol{r} = (1, 1)$ and $\boldsymbol{v} = (1, 1)$. For this setting, $\text{Rev}(\text{Opq}^*(\{1, 2\})) = 0.67$ and $\text{Rev}(\text{Opq}^*(\{1\})) = 0.5$, i.e., it is optimal to not include the opaque product in any assortment and the optimal assortment includes all identical products. ■

**Example 3.3** (Adding identical products decreases revenue)**.** Consider a setting with two products where $\boldsymbol{r} = (\infty, \infty)$ and $\boldsymbol{v} = (1, 1)$. For this setting, $\text{Rev}(\text{Opq}^*(\{1, 2\})) = 0.34$ and $\text{Rev}(\text{Opq}^*(\{1\})) = 0.57$, i.e., the optimal assortment contains only of the identical products. This observation can be extended to $n$ products as well, since when the prices are set to $\infty$, products can only be sold via the opaque product, and the valuation of the opaque product decreases when more products are added to the assortment. ■

In the two examples above, we show that adding multiple copies of a product to an assortment sometimes helps and sometimes does not with opaque selling. To deal with this, we break ties in revenue arbitrarily by adding small perturbations to the revenues. Our algorithm is as follows.

**Nested by revenue and valuation**

- **Input:** Revenues $\boldsymbol{r} = \{r_i \mid i \in \mathcal{N}\}$ and mean valuations $\boldsymbol{v} = \{v_i \mid i \in \mathcal{N}\}$
- **Step 1:** Add small perturbation to revenues for tie-breaks, $\widetilde{\boldsymbol{r}} = \boldsymbol{r} + \epsilon$.
- **Step 2:** Construct candidate assortments,

$$\mathcal{S} = \left\{S_{v,r} \mid v \in \boldsymbol{v}, r \in \widetilde{\boldsymbol{r}}\right\}$$

  where $S_{v,r} = \{i \in \mathcal{N} \mid v_i \geq v, r_i \geq r\}$.

- **Step 3:** Return revenue maximizing assortment from $\mathcal{S}$,

$$\hat{S} = \arg\max_{S \in \mathcal{S}} \text{REV}(\text{OPQ}^*(S)).$$

Note that despite performing well in practice, the NRV heuristic is also provably sub-optimal.

**Example 3.4** (Bad example for the NRV heuristic)**.** Consider the example with three products, where $\boldsymbol{r} = (25, 5.5, 3)$ and $\boldsymbol{v} = (2, 2, 2)$. For this example, the optimal nested by revenue and valuation assortment is $\{1, 2, 3\}$ with expected revenue 1.03. However, the optimal opaque assortment is $\{2, 3\}$ with expected revenue 1.05. In this example, product 1 is priced exorbitantly and its purchase probability is low. Moreover, adding this product to the assortment reduces the valuation of the opaque product for risk-averse customers, since they consider the minimum valuation across all products, decreasing the overall assortment revenue. ■

Nonetheless, we show that the set of candidate assortments considered by the NRV heuristic, which we denote as $\mathcal{S}$, is a superset of the TOS heuristic, and thus it is at least as good as the TOS heuristic.

**Theorem 3.5.** *If* $\{\hat{i}\}$ *is the best singleton assortment, then* $\{\hat{i}\} \in \mathcal{S}$*, i.e.* $\hat{i}$ *is the only product such that* $v_i \geq v_{\hat{i}}$ *and* $r_i > r_{\hat{i}}$*. This further implies that,*

$$\max_{S \in \mathcal{S}} \textsc{Rev}(\textsc{Opq}^*(S)) \geq \frac{1}{2} \textsc{Rev}(\textsc{Opq}^*(S^*)).$$

*Proof.* Suppose $\{\hat{i}\}$ is the best singleton assortment, i.e.,

$$\hat{i} = \arg\max_{i \in \mathcal{N}} \textsc{Rev}(\textsc{Opq}^*(\{i\})).$$

We prove the first statement by contradiction. Suppose there exists $i'$ such that $v_{i'} \geq v_{\hat{i}}$ and $r_{i'} > r_{\hat{i}}$. Note that for a singleton assortment $\{i\}$, the valuation of the opaque product for is $V_q = \min_j V_j = V_i$. Moreover, since $\rho \leq r_i$, $U_q = V_q - \rho \geq V_i - r_i = U_i$. It follows that the singleton opaque assortment is equivalent to a singleton traditional assortment at a lower price and

$$\textsc{Rev}(\textsc{Opq}^*(\{i\})) = \max_{\rho \leq r_i} \frac{\rho e^{v_i - \rho}}{1 + e^{v_i - \rho}} = \max_{\rho \leq r_i} \textsc{Rev}(\textsc{Trad}(\{i\}, \rho)).$$

Further, $\textsc{Rev}(\textsc{Trad}(\{i\}, \rho))$ increases monotonically with $v_i$. With these two observations, we prove the statement as follows

$$\begin{aligned} \textsc{Rev}(\textsc{Opq}^*(\{i'\})) &= \max_{\rho \leq r_{i'}} \frac{\rho e^{v_{i'} - \rho}}{1 + e^{v_{i'} - \rho}} > \max_{\rho \leq r_{\hat{i}}} \frac{\rho e^{v_{i'} - \rho}}{1 + e^{v_{i'} - \rho}} \geq \max_{\rho \leq r_{\hat{i}}} \frac{\rho e^{v_{\hat{i}} - \rho}}{1 + e^{v_{\hat{i}} - \rho}} \\ &= \textsc{Rev}\big(\textsc{Opq}^*\big(\{\hat{i}\}\big)\big), \end{aligned}$$

where the inequality follows from the fact $r_{i'} > r_{\hat{i}}$ and the second inequality follows from monotonicity and the fact that $v_{i'} \geq v_{\hat{i}}$. This contradicts the claim that $\{\hat{i}\}$ is the optimal singleton assortment. Therefore, $\{\hat{i}\} = S_{v_{\hat{i}}, r_{\hat{i}}}$ and $\{\hat{i}\} \in \mathcal{S}$.

We prove the second statement next. From theorem 3.1 we know that

$$\max\left\{\text{Rev}(\text{Opq}^*(\{\hat{i}\})), \text{Rev}\left(\text{Opq}^*\left(\hat{S}\right)\right)\right\} \geq \frac{1}{2}\text{Rev}(\text{Opq}^*(S^*)),$$

where $\hat{S}$ is the optimal traditional assortment. From the first part of the theorem, we know that $\{\hat{i}\} \in \mathcal{S}$. Further, since the optimal traditional assortment is nested-by-revenue, we know that $\hat{S} = S_{v,r}$ for $v = \min_{i \in \mathcal{N}} v_i$ and $r = \min_{i \in \hat{S}} r_i$. Therefore, $\hat{S} \in \mathcal{S}$. It follows that

$$\begin{aligned}\max_{S \in \mathcal{S}} \text{Rev}(\text{Opq}^*(S)) &\geq \max\left\{\text{Rev}(\text{Opq}^*(\{\hat{i}\})), \text{Rev}\left(\text{Opq}^*\left(\hat{S}\right)\right)\right\} \\ &\geq \frac{1}{2}\text{Rev}(\text{Opq}^*(S^*)).\end{aligned}$$

□

### 3.4 Price optimization

We next consider price optimization under the opaque selling mechanism for risk-averse MNL customers. For traditional selling, it is known that it is optimal to offer every product at the same price (the so called constant-markup policy, which reduces to a uniform pricing policy when there are no costs). Interestingly, we empirically observe that the same solution remains optimal for opaque selling. In fact, it is optimal to not offer the opaque product when prices of traditional products can be optimized. We formally state this conjecture below.

**Conjecture 3.6.** *Suppose $r^*$ is the optimal uniform price for traditional selling, i.e.,*

$$r^* \mathbf{1}_n = \arg\max_{\boldsymbol{r} \in \mathbb{R}^n_{\geq 0}} \text{Rev}(\text{Trad}(\boldsymbol{r})).$$

*Then, for all $\boldsymbol{r} \in \mathbb{R}^n_{\geq 0}$, $\text{Rev}(\text{Opq}^*(\boldsymbol{r})) \leq \text{Rev}(\text{Trad}(r^* \mathbf{1}_n))$.*

We prove this conjecture for two special cases and leave the general case as an open problem.

### 3.4.1 Uniform pricing

We first show that our claim holds when the pricing policies are restricted to uniform pricing, i.e. $r_i = r$ for all products $i \in \mathcal{N}$. Naturally, the price of the opaque product does not necessarily have to be the same as that of the traditional products.

**Theorem 3.7.** *If all products have the same price $r$, then for every opaque price $\rho \leq r$ there exists $\alpha \in [0,1]$ which depends on the primitives $(r, \rho, \boldsymbol{v})$ such that,*

$$\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) - \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n)) = \alpha(\textsc{Rev}(\textsc{Trad}(\rho\mathbf{1}_n)) - \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n))) \tag{3.7}$$

*This further implies that* $\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) \leq \textsc{Rev}(\textsc{Trad}(r^*\mathbf{1}_n))$ *for all $\rho \leq r$.*

*Proof.* We defer the proof of (3.7) to Appendix C.1. Assuming (3.7) to be true, we prove the rest of the theorem by considering two separate cases:

- **Case 1:** $\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) < \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n))$.

  This case is trivial since it directly implies $\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) < \textsc{Rev}(\textsc{Trad}(r^*\mathbf{1}_n))$.

- **Case 2:** $\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) \geq \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n))$.

  For this case, note that,

  $$\begin{aligned}\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) &= \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n)) + \alpha(\textsc{Rev}(\textsc{Trad}(\rho\mathbf{1}_n)) - \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n))) \\ &\leq \textsc{Rev}(\textsc{Trad}(\rho\mathbf{1}_n)) \leq \textsc{Rev}(\textsc{Trad}(r^*\mathbf{1}_n)),\end{aligned}$$

  where the first inequality follows since $\alpha \leq 1$, the second inequality follows from the definition of $r^*$.

$\square$

Note that the theorem only imposes a constraint on the prices, and allows for the products to have arbitrary valuations. The main idea is that when pricing is uniform, adding an opaque product at price $\rho$ moves the expected revenue in the same direction as re-pricing all products to $\rho$ under traditional selling. Moreover, the magnitude of revenue change is smaller for adding an opaque product.

### 3.4.2 Two products

We also show that our claim holds for two products for general prices and valuations. For this result, we show that the revenue function for opaque selling is a unimodal function of the prices and that the unique maximum is $r_1 = r_2 = \rho = r^*$.

**Theorem 3.8.** *For* $n = 2$ *products, the optimal solution for the joint pricing and assortment optimization problem is the same for opaque and traditional selling, i.e.,*

$$\max_{(r_1,r_2)\in\mathbb{R}^2_{\geq 0},\rho\leq\min(r_1,r_2)} \textsc{Rev}(\textsc{Opq}((r_1,r_2),\rho)) = \textsc{Rev}(\textsc{Trad}(r^*\mathbf{1}_2)),$$

*and the maximum is achieved when* $r_1 = r_2 = \rho = r^*$.

We defer the proof of the theorem to Appendix C.2.

## 3.5 Numerical experiments

In this section we numerically evaluate our proposed heuristics in Section 3.3. For the instance bed, we considered instances with different values of $n$ ranging from 2 to 10. We reduced the number of instances considered with $n$ since the computation times increase exponentially. For each product $i$, $r_i$ and $v_i$ are sampled independently from distributions $Lognormal(1, 0.5)$ and $Lognormal(0.5, 0.5)$ respectively. we choose the lognormal distribution for its heavy tail, and choose settings where the price of the product will tend to exceed the valuation to create instances where adding the opaque product is incentivized. In Table 3.1, we summarize the properties of the optimal assortments under opaque selling for our instances, which we denote by $S^*$. We consider the opaque product to be offered only if it has a non-zero purchase probability (i.e., $\rho < \min_{i \in S^*} r_i$). There is a roughly decreasing trend in the percentage of instances in which the opaque product is offered as $n$ increases. Furthermore, we see that the size of the optimal assortment when the opaque product is offered is much smaller than when it is not. This makes sense because offering the opaque product has a higher purchase probability for risk-averse customers when the assortment is smaller.

|  |  |  | average optimal assortment size |  |
|---|---|---|---|---|
| $n$ | #instances | #instances w/ opaque | w/ opaque | w/o opaque |
| 2 | 4000 | 2476 | 1.44 | 1.91 |
| 3 | 2000 | 974 | 2.00 | 2.77 |
| 4 | 2000 | 754 | 2.40 | 3.55 |
| 5 | 1000 | 311 | 2.68 | 4.32 |
| 6 | 1000 | 227 | 3.30 | 4.96 |
| 7 | 500 | 84 | 3.48 | 5.59 |
| 8 | 500 | 84 | 3.57 | 6.14 |
| 9 | 250 | 40 | 3.70 | 6.78 |
| 10 | 100 | 14 | 3.00 | 7.21 |

Table 3.1: Summary for the optimal assortments under opaque selling, for different number of products.

In Tables 3.2 and 3.3, we analyze the performance of our heuristics on the instance bed. The traditional optimal or singleton heuristic is referred to as *TOS*, and the nested by revenue and valuation heuristic is referred to as *NRV*. The analysis is limited to instances where the opaque product is offered, since both heuristics always output the optimal assortment when the opaque product is not offered. Overall, out of the 11350 instances we considered, the opaque product was offered in 4964 instances, of which, TOS and NRV deviated from the optimal assortment on 67 and 4 instances respectively. The largest optimality gap for TOS was 9.3%, while for NRV, it was 0.7%. Both values are much larger than the theoretical guarantees in Theorems 3.1 and 3.5. Moreover, NRV is near-optimal.

In Table 3.2 we show how the sub-optimal instances are distributed across different values of $n$. The vast majority of sub-optimal instances arise when $n = 3$ or $n = 4$. Interestingly, both heuristics were optimal for all instances when $n = 2$. This is expected for NRV, since all possible assortments are nested by revenue and valuation unless $v_1 > v_2$ and $r_1 > r_2$, in which case $\{2\}$ cannot be the optimal assortment (from Theorem 3.5). Suboptimality is only possible for TOS when $n = 2$ if the optimal assortment is to offer both products while the optimal traditional assortment is singleton. This implies that the optimal assortment is always a subset of the optimal traditional assortment in our instances for $n = 2$.

| | | TOS | | NRV | |
|---|---|---|---|---|---|
| $n$ | #instances considered | #subopt | max. optimality gap (%) | #subopt | max. optimality gap (%) |
| 2 | 2476 | 0 | - | 0 | - |
| 3 | 974 | 29 | 9.3 | 1 | 0.7 |
| 4 | 754 | 23 | 2.5 | 3 | 0.6 |
| 5 | 311 | 8 | 5.0 | 0 | - |
| 6 | 227 | 5 | 1.7 | 0 | - |
| 7 | 84 | 0 | - | 0 | - |
| 8 | 84 | 0 | - | 0 | - |
| 9 | 40 | 1 | 0.2 | 0 | - |
| 10 | 14 | 1 | 3.2 | 0 | - |

Table 3.2: Comparing algorithms TOS and NRV for the same instances as Table 3.1 but limiting to instances where the opaque product is offered in the optimal assortment. For each algorithm, "#subopt" is the number of instances where the output by the algorithm is suboptimal; and "max. optimality gap" is maximum ratio of the optimality gap and the optimal revenue.

We also sort the instances by the size of the optimal assortment and present those trends in Table 3.3. A wide majority of the sub-optimal TOS instances occur when $|S^*| = 2$. We also see that the worst case optimality gap reduces as $|S^*|$ grows larger. This indicates that the optimal traditional and opaque assortments are usually identical when they are large.

## 3.6 Conclusions

In this work, we illustrated structural and algorithmic insights for the assortment optimization problem with opaque products for risk-averse MNL customers. We proposed two assortment optimization heuristics. Both heuristics are shown to have a constant factor performance guarantee, that they vastly outperform in experiments. The nested by revenue and valuation heuristic in particular is near-optimal. The gap between the theory and experiments suggests that there is room for improvement in the theoretical guarantees for these heuristics. For instance, the guarantee

| | | TOS | | NRV | |
|---|---|---|---|---|---|
| $\lvert S^*\rvert$ | #instances considered | #subopt | max. optimality gap (%) | #subopt | max. optimality gap (%) |
| 1 | 2482 | 0 | - | 0 | - |
| 2 | 1290 | 52 | 9.3 | 1 | 0.7 |
| 3 | 568 | 12 | 1.6 | 3 | 0.6 |
| 4 | 353 | 3 | 0.2 | 0 | - |
| 5 | 161 | 0 | - | 0 | - |
| 6 | 75 | 0 | - | 0 | - |
| 7 | 21 | 0 | - | 0 | - |
| 8 | 13 | 0 | - | 0 | - |
| 9 | 1 | 0 | - | 0 | - |

Table 3.3: This table is similar to Table 3.2, but the results are represented based on the size of the optimal assortment, instead of the number of products. Again, we only consider instances where the opaque product is in the optimal assortment.

for the traditional optimal or singleton heuristic does not use any properties specific to the MNL choice model. Another interesting open question is to characterize the optimal assortment policy for risk-averse customers. In our experiments, we have seen that the nested by revenue and valuation heuristic is sub-optimal in cases where one of the products has high revenue and consequently a very low purchase probability.

For the price optimization problem, we show that adding an opaque product does not improve revenue. We prove this when pricing is constrained to be uniform and when there are two products. Proving this in more generality is one of our goals for the future. Another question that we want to answer is whether or not this phenomenon is specific to risk-averse MNL customers. It would be good to know if the property still holds for risk-neutral customers, or a setting with costs, or a different valuation paradigm entirely.

# References

[1] G. Gallego and H. Topaloglu, *Revenue management and pricing analytics* (International Series in Operations Research & Management Science). New York, NY, USA: Springer, 2019, vol. 209.

[2] R. L. Phillips, *Pricing and revenue optimization*. Stanford university press, 2021.

[3] N. Golrezaei, H. Nazerzadeh, and P. Rusmevichientong, "Real-time optimization of personalized assortments," *Management Science*, vol. 60, no. 6, pp. 1532–1551, 2014.

[4] F. Bernstein, A. G. Kök, and L. Xie, "Dynamic assortment customization with limited inventories," *Manufacturing & Service Operations Management*, vol. 17, no. 4, pp. 538–553, 2015.

[5] A. N. Elmachtoub and R. D. Lederman, *Revenue management using dynamic customer selection*, US Patent 11,151,604, Oct. 2021.

[6] G. Gallego, G. Iyengar, R. Phillips, and A. Dubey, "Managing flexible products on a network," CORC, Columbia University, New York, NY, Tech. Rep. TR-2004-01, 2004.

[7] Q. Liu and G. Van Ryzin, "On the choice-based linear programming model for network revenue management," *Manufacturing & Service Operations Management*, vol. 10, no. 2, pp. 288–310, 2008.

[8] W. Ma, D. Simchi-Levi, and J. Zhao, "Dynamic pricing (and assortment) under a static calendar," *Management Science*, vol. 67, no. 4, pp. 2292–2313, 2021.

[9] G. Gallego, R. Ratliff, and S. Shebalov, "A general attraction model and sales-based linear program for network revenue management under customer choice," *Operations Research*, vol. 63, no. 1, pp. 212–232, 2015.

[10] J. B. Feldman and H. Topaloglu, "Revenue management under the markov chain choice model," *Operations Research*, vol. 65, no. 5, pp. 1322–1342, 2017.

[11] A. J. Liang, S. Jasin, and J. Uichanco, "Assortment and inventory planning under dynamic substitution with mnl model: Structural results and a near-optimal heuristic," *Available at SSRN 3739047*, 2020.

[12] J. M. Davis, H. Topaloglu, and D. P. Williamson, "Assortment optimization over time," *Operations Research Letters*, vol. 43, no. 6, pp. 608–611, 2015.

[13] X. Chen, W. Ma, D. Simchi-Levi, and L. Xin, "Assortment planning for recommendations at checkout under inventory constraints," MIT, Cambridge, MA, Tech. Rep., 2016.

[14] N. Liu, Y. Ma, and H. Topaloglu, "Assortment optimization under the multinomial logit model with sequential offerings," *INFORMS Journal on Computing*, vol. 32, no. 3, pp. 835–853, 2020.

[15] E. Fata, W. Ma, and D. Simchi-Levi, "Multi-stage and multi-customer assortment optimization with inventory constraints," *Available at SSRN 3443109*, 2019.

[16] Y. Feng, R. Niazadeh, and A. Saberi, "Linear programming based online policies for real-time assortment of reusable resources," *Available at SSRN 3421227*, 2019.

[17] P. Rusmevichientong, M. Sumida, and H. Topaloglu, "Dynamic assortment optimization for reusable products with random usage durations," *Management Science*, vol. 66, no. 7, pp. 2820–2844, 2020.

[18] J. Baek and W. Ma, "Bifurcating constraints to improve approximation ratios for network revenue management with reusable resources," *Available at SSRN 3482457*, 2019.

[19] V. Goyal, G. Iyengar, and R. Udwani, "Asymptotically optimal competitive ratio for online allocation of reusable resources," *arXiv preprint arXiv:2002.02430*, 2020.

[20] S. Jasin and S. Kumar, "A re-solving heuristic with bounded revenue loss for network revenue management with customer choice," *Mathematics of Operations Research*, vol. 37, no. 2, pp. 313–345, 2012.

[21] C. W. Chan and V. F. Farias, "Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems," *Mathematics of Operations Research*, vol. 34, no. 2, pp. 333–350, 2009.

[22] W. Ma and D. Simchi-Levi, "Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios," *Operations Research*, vol. 68, no. 6, pp. 1787–1803, 2020.

[23] R. Cominetti, J. R. Correa, T. Rothvoß, and J. S. Martín, "Optimal selection of customers for a last-minute offer," *Operations research*, vol. 58, no. 4-part-1, pp. 878–888, 2010.

[24] E. Balas, "The prize collecting traveling salesman problem," *Networks*, vol. 19, no. 6, pp. 621–636, 1989.

[25] D. Bienstock, M. X. Goemans, D. Simchi-Levi, and D. Williamson, "A note on the prize collecting traveling salesman problem," *Mathematical programming*, vol. 59, no. 1, pp. 413–420, 1993.

[26] M. X. Goemans and D. P. Williamson, “A general approximation technique for constrained forest problems,” *SIAM Journal on Computing*, vol. 24, no. 2, pp. 296–317, 1995.

[27] A. Archer, M. H. Bateni, M. T. Hajiaghayi, and H. Karloff, “Improved approximation algorithms for prize-collecting steiner tree and tsp,” *SIAM journal on computing*, vol. 40, no. 2, pp. 309–332, 2011.

[28] M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan, “Algorithms for facility location problems with outliers,” in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA ’01, Washington, D.C., USA: SIAM, 2001, 642–651, ISBN: 0898714907.

[29] G. Xu and J. Xu, “An improved approximation algorithm for uncapacitated facility location problem with penalties,” *Journal of combinatorial optimization*, vol. 17, no. 4, pp. 424–436, 2009.

[30] Y. Li, D. Du, N. Xiu, and D. Xu, “Improved approximation algorithms for the facility location problems with linear/submodular penalties,” *Algorithmica*, vol. 73, no. 2, pp. 460–482, 2015.

[31] J. Geunes, R. Levi, H. E. Romeijn, and D. B. Shmoys, “Approximation algorithms for supply chain planning and logistics problems with market choice,” *Mathematical programming*, vol. 130, no. 1, pp. 85–106, 2011.

[32] A. N. Elmachtoub and R. Levi, “From cost sharing mechanisms to online selection problems,” *Mathematics of Operations Research*, vol. 40, no. 3, pp. 542–557, 2015.

[33] A. N. Elmachtoub and R. Levi, “Supply chain management with online customer selection,” *Operations Research*, vol. 64, no. 2, pp. 458–473, 2016.

[34] R. D. Luce, *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.

[35] D. McFadden, “Conditional logit analysis of qualitative choice behavior,” *Frontiers in Econometrics*, pp. 105–142, 1973.

[36] R. L. Plackett, “The analysis of permutations,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 24, no. 2, pp. 193–202, 1975.

[37] K. Talluri and G. Van Ryzin, “Revenue management under a general discrete choice model of consumer behavior,” *Management Science*, vol. 50, no. 1, pp. 15–33, 2004.

[38] N. DeHoratius and A. Raman, “Inventory record inaccuracy: An empirical analysis,” *Management science*, vol. 54, no. 4, pp. 627–641, 2008.

[39] H. H.-C. Chuang and R. Oliva, "Inventory record inaccuracy: Causes and labor effects," *Journal of Operations Management*, vol. 39, pp. 63–78, 2015.

[40] A. E. Lim and J. G. Shanthikumar, "Relative entropy, exponential utility, and robust dynamic pricing," *Operations Research*, vol. 55, no. 2, pp. 198–214, 2007.

[41] A. V. den Boer and B. Zwart, "Simultaneously learning and optimizing using controlled variance pricing," *Management science*, vol. 60, no. 3, pp. 770–783, 2014.

[42] J. Broder and P. Rusmevichientong, "Dynamic pricing under a general parametric choice model," *Operations Research*, vol. 60, no. 4, pp. 965–980, 2012.

[43] O. Besbes and A. Zeevi, "On the (surprising) sufficiency of linear models for dynamic pricing with demand learning," *Management Science*, vol. 61, no. 4, pp. 723–739, 2015.

[44] B. Chen, X. Chao, and H.-S. Ahn, "Coordinating pricing and inventory replenishment with nonparametric demand learning," *Operations Research*, vol. 67, no. 4, pp. 1035–1052, 2019.

[45] J. Li, N. Granados, and S. Netessine, "Are consumers strategic? structural estimation from the air-travel industry," *Management Science*, vol. 60, no. 9, pp. 2114–2137, 2014.

[46] K. Moon, K. Bimpikis, and H. Mendelson, "Randomized markdowns and online monitoring," *Management Science*, vol. 64, no. 3, pp. 1271–1290, 2018.

[47] M. M. Wei and F. Zhang, "Recent research developments of strategic consumer behavior in operations management," *Computers & Operations Research*, vol. 93, pp. 166–176, 2018.

[48] Y. Chen, V. F. Farias, and N. Trichakis, "On the efficacy of static prices for revenue management in the face of strategic customers," *Management Science*, vol. 65, no. 12, pp. 5535–5555, 2019.

[49] Y. Chen and C. Shi, "Joint pricing and inventory management with strategic customers," *Operations Research*, vol. 67, no. 6, pp. 1610–1627, 2019.

[50] X. Chen and D. Simchi-Levi, "Coordinating inventory control and pricing strategies: The continuous review model," *Operations Research Letters*, vol. 34, no. 3, pp. 323–332, 2006.

[51] T. M. Whitin, "Inventory control and price theory," *Management science*, vol. 2, no. 1, pp. 61–68, 1955.

[52] A. Federgruen and A. Heching, "Combined pricing and inventory control under uncertainty," *Operations research*, vol. 47, no. 3, pp. 454–475, 1999.

[53] X. Chen and D. Simchi-Levi, “Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case,” *Operations research*, vol. 52, no. 6, pp. 887–896, 2004.

[54] W. T. Huh and G. Janakiraman, “(s, s) optimality in joint inventory-pricing control: An alternate approach,” *Operations Research*, vol. 56, no. 3, pp. 783–790, 2008.

[55] X. Chao and S. X. Zhou, “Joint inventory-and-pricing strategy for a stochastic continuous-review system,” *IIE Transactions*, vol. 38, no. 5, pp. 401–408, 2006.

[56] H. Chen, O. Q. Wu, and D. D. Yao, “On the benefit of inventory-based dynamic pricing strategies,” *Production and Operations Management*, vol. 19, no. 3, pp. 249–260, 2010.

[57] F. Bernstein, Y. Li, and K. Shang, “A simple heuristic for joint inventory and pricing models with lead time and backorders,” *Management Science*, vol. 62, no. 8, pp. 2358–2373, 2016.

[58] G. Gallego and G. Van Ryzin, “Optimal dynamic pricing of inventories with stochastic demand over finite horizons,” *Management science*, vol. 40, no. 8, pp. 999–1020, 1994.

[59] O. Besbes, A. N. Elmachtoub, and Y. Sun, “Static pricing: Universal guarantees for reusable resources,” *Operations Research*, vol. 70, no. 2, pp. 1143–1152, 2022.

[60] M. Lei, S. Liu, S. Jasin, and A. Vakhutinsky, “Joint inventory and pricing for a one-warehouse multistore problem: Spiraling phenomena, near optimal policies, and the value of dynamic pricing,” *Operations Research*, 2022.

[61] G. Gallego and R. Phillips, “Revenue management of flexible products,” *Manufacturing & Service Operations Management*, vol. 6, no. 4, pp. 321–337, 2004.

[62] A. N. Elmachtoub and M. L. Hamilton, “The power of opaque products in pricing,” *Management Science*, vol. 67, no. 8, pp. 4686–4702, 2021.

[63] S. Fay and J. Xie, “Probabilistic goods: A creative way of selling products and services,” *Marketing Science*, vol. 27, no. 4, pp. 674–690, 2008.

[64] Y. Jiang, “Price discrimination with opaque products,” *Journal of Revenue and Pricing Management*, vol. 6, pp. 118–134, 2007.

[65] K. Jerath, S. Netessine, and S. K. Veeraraghavan, “Revenue management with strategic customers: Last-minute selling and opaque selling,” *Management science*, vol. 56, no. 3, pp. 430–448, 2010.

[66] Z. Zhang, K. Joseph, and R. Subramaniam, “Probabilistic selling in quality-differentiated markets,” *Management Science*, vol. 61, no. 8, pp. 1959–1977, 2015.

[67] A. Strauss, N. Gülpınar, and Y. Zheng, “Dynamic pricing of flexible time slots for attended home delivery,” *European Journal of Operational Research*, vol. 294, no. 3, pp. 1022–1041, 2021.

[68] G. Gallego, G. Iyengar, R. Phillips, and A. Dubey, “Managing flexible products on a network,” *Available at SSRN 3567371*, 2004.

[69] Q. Liu, Y. Xiao, and D. Zhang, “Opaque selling of multiple substitutable products with finite inventories,” *Naval Research Logistics (NRL)*, vol. 69, no. 4, pp. 529–549, 2022.

[70] A. N. Elmachtoub, Y. Wei, and Y. Zhou, “Retailing with opaque products,” *Available at SSRN 2659211*, 2015.

[71] A. N. Elmachtoub, D. Yao, and Y. Zhou, “The value of flexibility from opaque selling,” *Available at SSRN 3483872*, 2019.

[72] S. Alaei, M. Hajiaghayi, and V. Liaghat, “Online prophet-inequality matching with applications to ad allocation,” in *Proceedings of the 13th ACM Conference on Electronic Commerce*, ser. EC ’12, Valencia, Spain: ACM, 2012, 18–35, ISBN: 9781450314152.

[73] A. Asadpour and H. Nazerzadeh, “Maximizing stochastic monotone submodular functions,” *Management Science*, vol. 62, no. 8, pp. 2374–2391, 2016.

# Appendix A: Appendix to Chapter 1

## A.1 Omitted Proofs of Section 1.2

### A.1.1 Proof of Lemma 1.1

*Proof.* Consider any policy $\mathcal{A}$. Let $q_i^j(\mathcal{A})$ be the expected sales of product $i$ to type-$j$ customers under policy $\mathcal{A}$.

$$q_i^j(\mathcal{A}) = \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{I}[X_t(\mathcal{A}) = i, J_t(\mathcal{A}) = j]\right] = \sum_{t=1}^{T} \mathbb{P}[X_t(\mathcal{A}) = i, J_t(\mathcal{A}) = j],$$

which follows from the definition of $J_t(\mathcal{A})$ and $X_t(\mathcal{A})$.

We first show that $\boldsymbol{q}(\mathcal{A})$ is a feasible solution to (SLP). Since sales cannot be negative, $q_i^j(\mathcal{A}) \geq 0$ for all $i, j$. Recall that $Z_i(\mathcal{A})$ and $Y_j(\mathcal{A})$ denote the sales of product $i$ and number of selections of type-$j$ customers in one sample path of the algorithm, respectively. Thus,

$$\mathbb{E}[Z_i(\mathcal{A})] = \sum_{t=1}^{T} \mathbb{P}[X_t(\mathcal{A}) = i] = \sum_{t=1}^{T} \sum_{j \in \mathcal{L}} \mathbb{P}[X_t(\mathcal{A}) = i, J_t(\mathcal{A}) = j] = \sum_{j \in \mathcal{L}} q_i^j(\mathcal{A}), \tag{A.1}$$

$$\begin{aligned} \mathbb{E}[Y_j(\mathcal{A})] &= \sum_{t=1}^{T} \mathbb{P}[J_t(\mathcal{A}) = j] = \sum_{t=1}^{T} \sum_{i \in \mathcal{N} \cup \{0\}} \mathbb{P}[X_t(\mathcal{A}) = i, J_t(\mathcal{A}) = j] \\ &= q_0^j(\mathcal{A}) + \sum_{i \in \mathcal{N}} q_i^j(\mathcal{A}). \end{aligned} \tag{A.2}$$

To show supply constraints and demand constraints, note that $Z_i(\mathcal{A}) \leq c_i$ for all products $i \in \mathcal{N}$

and $Y_j(\mathcal{A}) \leq b_j$ for all customer types $j \in \mathcal{L}$. Therefore,

$$\begin{aligned}\sum_{j \in \mathcal{L}} q_i^j(\mathcal{A}) &\leq c_i, \\ q_0^j(\mathcal{A}) + \sum_{i \in \mathcal{N}} q_i^j(\mathcal{A}) &\leq b_j,\end{aligned}$$

where the first inequality follows from (A.1) and the second inequality follows from (A.2). To show the time horizon constraint, observe that, since the total number of customers is at most $T$, then $\sum_{j \in \mathcal{L}} Y_j(\mathcal{A}) \leq T$. Thus,

$$\mathbb{E}\left[\sum_{j \in \mathcal{L}} Y_j(\mathcal{A})\right] = \sum_{j \in \mathcal{L}}\left(q_0^j(\mathcal{A}) + \sum_{i \in \mathcal{N}} q_i^j(\mathcal{A})\right) \leq T.$$

Finally we show the MNL constraint as follows:

$$\begin{aligned}q_i^j(\mathcal{A}) &= \sum_{t=1}^{T} \mathbb{P}[X_t(\mathcal{A}) = i, J_t(\mathcal{A}) = j] = \sum_{t=1}^{T} \mathbb{P}[X_t(\mathcal{A}) = i \mid J_t(\mathcal{A}) = j]\mathbb{P}[J_t(\mathcal{A}) = j] \\ &= \sum_{t=1}^{T} \pi_i^j(S_t(\mathcal{A}))\mathbb{P}[J_t(\mathcal{A}) = j] = \sum_{t=1}^{T} \mathbb{I}[i \in S_t] u_i^j \pi_0^j(S_t(\mathcal{A}))\mathbb{P}[J_t(\mathcal{A}) = j] \\ &\leq \sum_{t=1}^{T} u_i^j \pi_0^j(S_t(\mathcal{A}))\mathbb{P}[J_t(\mathcal{A}) = j] \\ &= \sum_{t=1}^{T} u_i^j \mathbb{P}[X_t(\mathcal{A}) = 0 \mid J_t(\mathcal{A}) = j]\mathbb{P}[J_t(\mathcal{A}) = j] \\ &= \sum_{t=1}^{T} u_i^j \mathbb{P}[X_t(\mathcal{A}) = 0, J_t(\mathcal{A}) = j] \\ &= u_i q_0^j(\mathcal{A}),\end{aligned}$$

where the third follows from the definition of $\pi_i^j(\cdot)$ and $S_t(\mathcal{A})$, the fourth equality follows from the definition of the MNL choice probabilities in (1.3), and the fifth equality follows from the definition of $\pi_i^j(\cdot)$. Therefore, $\boldsymbol{q}(\mathcal{A})$ is a feasible solution to (SLP).

The objective value at $\boldsymbol{q}(\mathcal{A})$ matches the expected revenue under policy $\mathcal{A}$ (from (1.2)), since,

$$\textsc{Rev}(\mathcal{A}) = \sum_{i=1}^{n} r_i \mathbb{E}[Z_i(\mathcal{A})] = \sum_{i=1}^{n} \sum_{j=1}^{m} r_i q_i^j(\mathcal{A}),$$

where the second equality follows from (A.1). Since $\boldsymbol{q}(\mathcal{A})$ is a feasible solution to (SLP) with objective value $\textsc{Rev}(\mathcal{A})$, then for any policy $\mathcal{A}$, $\textsc{Rev}(\mathcal{A})$ is bounded above by $z^*_{SLP}$ (the optimal value of (SLP)). □

## A.2 Omitted Proofs from Section 1.3

### A.2.1 Proof of Theorem 1.2

*Proof.* Combining Lemmas 1.3 and 1.4 we get,

$$\textsc{Rev}(\text{A1}) \geq \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} r_i \widetilde{q}_i^j \geq \frac{c_{\min} - 1}{4c_{\min} - 2} z^*_{SLP},$$

where the first and second inequalities follow from Lemma 1.4 and Lemma 1.3 respectively. This concludes the proof of the theorem. □

### A.2.2 Proof of Lemma 1.3

*Proof.* In this Lemma, we show that $\widetilde{\boldsymbol{q}}$ is a feasible solution to (SLP) and that the objective value of (SLP) at $\widetilde{\boldsymbol{q}}$ is at least a $1/2 - o(1)$ fraction of the optimal value of the LP upper bound.

We first show that $\widetilde{\boldsymbol{q}}$ (as defined in (1.4)) is a feasible solution to (SLP). To show the MNL constraints, observe that $\widetilde{q}_i^j \leq u_i^j \widetilde{q}_0^j$ since $\pi_i^{\widetilde{J}_t}(\widetilde{S}_t) \leq u_i^j \pi_0^{\widetilde{J}_t}(\widetilde{S}_t)$ for all $i \in \mathcal{N}$ by definition of the MNL purchase probabilities in (1.3). Next, we show demand and supply constraints. According to the update rule for $\widetilde{Y}_j$ and $\widetilde{Z}_i$ used in Algorithm 1, the value that $\widetilde{Y}_j$ and $\widetilde{Z}_i$ will take at the end of the customer selection step is as follows:

$$\widetilde{Y}_j = \sum_{t=1}^{T} \mathbb{I}[\widetilde{J}_t = j], \ \widetilde{Z}_i = \sum_{t=1}^{T} \pi_i^{\widetilde{J}_t}(\widetilde{S}_t). \tag{A.3}$$

We re-write the value of $\widetilde{Y}_j$ and $\widetilde{Z}_i$ in terms of $\widetilde{q}_i^j$ as,

$$\widetilde{q}_0^j + \sum_{i \in \mathcal{N}} \widetilde{q}_i^j = \widetilde{Y}_j, \ \sum_{j \in \mathcal{L}} \widetilde{q}_i^j = \widetilde{Z}_i, \tag{A.4}$$

where both equalities follow from (A.3) and (1.4). Since a customer type $j$ becomes unavailable when $\widetilde{Y}_j = b_j$, it follows that $\widetilde{Y}_j \leq b_j$. Also, since a product $i$ is not offered once $\widetilde{Z}_i > c_i - 1$, it follows that $\widetilde{Z}_i \leq c_i$. This is because the maximum probability of a customer purchasing product $i$ is bounded above by 1, so the first time $\widetilde{Z}_i$ goes above $c_i - 1$, it cannot exceed $c_i$. It follows that $\widetilde{\boldsymbol{q}}$ satisfies supply constraints and demand constraints. To show the time horizon constraint, observe that,

$$\sum_{j \in \mathcal{L}} \left( \sum_{i=1}^{n} \widetilde{q}_i^j + \widetilde{q}_0^j \right) = \sum_{j \in \mathcal{L}} \widetilde{Y}_j = \sum_{t=1}^{T} \sum_{j \in \mathcal{L}} \mathbb{I}[\widetilde{J}_t = j] \leq T,$$

where the first equality follows from (A.4) and the second equality follows from (A.3). Therefore, the customer selection step constructs a feasible solution to (SLP).

Next, we prove the approximation guarantee using dual-fitting. Consider the dual of (SLP),

$$\begin{aligned}
z^*_{SLP} := \min_{\lambda, \eta_j, \theta_i, \beta_i^j} \quad & T\lambda + \sum_{j=1}^{m} b_j \eta_j + \sum_{i=1}^{n} c_i \theta_i && && (D_2) \\
\text{s.t.} \quad & \lambda + \eta_j + \beta_i^j \geq r_i - \theta_i && \forall i \in \mathcal{N}, j \in \mathcal{L}, && \text{(A.5)} \\
& \lambda + \eta_j \geq \sum_{i=1}^{n} u_i^j \beta_i^j && \forall j \in \mathcal{L}, && \text{(A.6)} \\
& \lambda, \eta_j, \theta_i, \beta_i^j \geq 0 && \forall i \in \mathcal{N}, j \in \mathcal{L}. &&
\end{aligned}$$

Here, the dual variables $\lambda$, $\boldsymbol{\eta}$, $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ correspond to the time horizon, demand, supply and MNL constraints of (SLP) respectively.

Fix $\widetilde{Y}_j$ and $\widetilde{Z}_i$ to be the state of the corresponding variables in Algorithm 1 at the end of the execution of the algorithm. Additionally, define $\widetilde{\mathcal{N}}_{T+1} = \{i \in \mathcal{N} \mid \widetilde{Z}_i \leq c_i - 1\}$ and $\widetilde{\mathcal{L}}_{T+1} = \{j \in$

$\mathcal{L} \mid \widetilde{Y}_j \le b_j - 1\}$. Consider the following candidate dual solution:

$$\hat{\lambda} := \max_{j \in \widetilde{\mathcal{L}}_{T+1}, S \subseteq \widetilde{\mathcal{N}}_{T+1}} R_{\pi^j}(S), \ \hat{\theta}_i := \begin{cases} 0 & i \in \widetilde{\mathcal{N}}_{T+1}, \\ r_i & i \in \mathcal{N} \setminus \widetilde{\mathcal{N}}_{T+1}, \end{cases},$$

$$\hat{\eta}_j := \frac{\sum_{t:\widetilde{J}_t = j} (\widetilde{R}_t - \hat{\lambda})}{b_j} \quad \forall j \in \mathcal{L},$$

$$\hat{\beta}_i^j := \max\{r_i - \hat{\theta}_i - \hat{\lambda} - \hat{\eta}_j, 0\} \quad \forall i \in \mathcal{N} \ \forall j \in \mathcal{L}.$$

We show that the candidate dual solution $(\hat{\lambda}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\beta}})$ is a feasible solution of $(D_2)$.

Note that, since $\widetilde{\mathcal{L}}_t \supseteq \widetilde{\mathcal{L}}_{T+1}$ and $\widetilde{N}_t \supseteq \widetilde{N}_{T+1}$ for all $t = 1, \ldots, T$, therefore,

$$\widetilde{R}_t = \max_{j \in \widetilde{\mathcal{L}}_t, S \subseteq \widetilde{\mathcal{N}}_t} R_{\pi^j}(S) \ge \max_{j \in \widetilde{\mathcal{L}}_{T+1}, S \subseteq \widetilde{\mathcal{N}}_{T+1}} R_{\pi^j}(S) = \hat{\lambda}, \tag{A.7}$$

where the first equality follows from the definition of $\widetilde{R}_t$ in Algorithm 1 and the second equality follows from the definition of $\hat{\lambda}$. From the definition of $\hat{\eta}$ and (A.7), it follows that $\hat{\eta}_j \ge 0 \ \forall j \in \mathcal{L}$. Also, by definition, $\hat{\lambda} \ge 0$, $\hat{\theta}_i \ge 0$ and $\hat{\beta}_i^j \ge 0 \ \forall i \in \mathcal{N}, j \in \mathcal{L}$. To show constraint (A.5), observe that by definition, $\hat{\beta}_i^j \ge r_i - \hat{\theta}_i - \hat{\lambda} - \hat{\eta}_j$.

Next, we show constraint (A.6). Let $\hat{S}^{(j)} := \{i \in \mathcal{N} \mid \hat{\beta}_i^j > 0\}$ for all $j \in \mathcal{L}$. Note that $\hat{S}^{(j)} = \{i \in \widetilde{\mathcal{N}}_{T+1} \mid r_i \ge \hat{\lambda} + \hat{\eta}_j\}$ by definition of $(\hat{\lambda}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\beta}})$. We show that for all $j \in \mathcal{L}$, $R_{\pi^j}(\hat{S}^{(j)}) \le \hat{\lambda} + \hat{\eta}_j$. By definition of $\widetilde{R}_t$ and $\widetilde{J}_t$ in Algorithm 1,

$$\widetilde{R}_t = \max_{S \subseteq \widetilde{\mathcal{N}}_t} R_{\pi^{\widetilde{J}_t}}(S) \ge R_{\pi^{\widetilde{J}_t}}(\hat{S}^{(\widetilde{J}_t)}), \tag{A.8}$$

where the inequality follows from the fact that $\hat{S}^{(\widetilde{J}_t)} \subseteq \widetilde{\mathcal{N}}_{T+1} \subseteq \widetilde{\mathcal{N}}_t$. If $j \in \widetilde{\mathcal{L}}_{T+1}$, then,

$$R_{\pi^j}(\hat{S}^{(j)}) \le \hat{\lambda} \le \hat{\lambda} + \hat{\eta}_j, \tag{A.9}$$

where the first inequality follows from the definition of $\hat{\lambda}$ and the fact that $\hat{S}^{(j)} \subseteq \widetilde{\mathcal{N}}_{T+1}$ and the second inequality follows from the non-negativity of $\hat{\eta}_j$. It follows from the definition of $\widetilde{Y}_j$ in (A.3) that $\widetilde{Y}_j = |\{t \mid \widetilde{J}_t = j\}|$. By definition of $\widetilde{\mathcal{L}}_{T+1}$, $\widetilde{Y}_j = b_j$ for $j \in \mathcal{L} \setminus \widetilde{\mathcal{L}}_{T+1}$. Therefore, if $j \in \mathcal{L} \setminus \widetilde{\mathcal{L}}_{T+1}$,

$$\begin{aligned}\hat{\lambda} + \hat{\eta}_j &= \hat{\lambda} + \frac{\sum_{t:\widetilde{J}_t = j}(\widetilde{R}_t - \hat{\lambda})}{b_j} \geq \hat{\lambda} + \frac{\sum_{t:\widetilde{J}_t = j}\left(R_{\pi^j}(\hat{S}^{(j)}) - \hat{\lambda}\right)}{b_j} \\ &= \hat{\lambda} + \frac{b_j\left(R_{\pi^j}(\hat{S}^{(j)}) - \hat{\lambda}\right)}{b_j} = R_{\pi^j}(\hat{S}^{(j)}),\end{aligned} \tag{A.10}$$

where the first equality follows from the definition of $\hat{\eta}_j$, the inequality follows from (A.8) and the second equality follows from the fact that $|\{t \mid \widetilde{J}_t = j\}| = \widetilde{Y}_j = b_j$ for all $j \in \mathcal{L} \setminus \widetilde{\mathcal{L}}_{T+1}$. Therefore, from (A.9) and (A.10),

$$R_{\pi^j}(\hat{S}^{(j)}) \leq \hat{\lambda} + \hat{\eta}_j. \tag{A.11}$$

Now we can show constraint (A.6) as follows:

$$\begin{aligned}\sum_{i=1}^{n} u_i^j \hat{\beta}_i^j - (\hat{\lambda} + \hat{\eta}_j) &= \sum_{i=1}^{n} u_i^j \max\{r_i - \hat{\theta}_i - \hat{\lambda} - \hat{\eta}_j, 0\} - (\hat{\lambda} + \hat{\eta}_j) \\ &= \sum_{i \in \hat{S}^{(j)}} u_i^j (r_i - \hat{\lambda} - \hat{\eta}_j) - (\hat{\lambda} + \hat{\eta}_j) \\ &= \sum_{i \in \hat{S}^{(j)}} u_i^j r_i - \left(1 + \sum_{i \in \hat{S}^{(j)}} u_i^j\right)(\hat{\lambda} + \hat{\eta}_j) \\ &= \left(1 + \sum_{i \in \hat{S}^{(j)}} u_i^j\right)\left(\sum_{i \in \hat{S}^{(j)}} r_i \pi_i^j(\hat{S}^{(j)}) - \hat{\lambda} - \hat{\eta}_j\right) \\ &= \left(1 + \sum_{i \in \hat{S}^{(j)}} u_i^j\right)\left(R_{\pi^j}(\hat{S}^{(j)}) - \hat{\lambda} - \hat{\eta}_j\right) \\ &\leq 0,\end{aligned} \tag{A.12}$$

where the first and second equalities follow from the definitions of $\hat{\beta}_i^j$ and $\hat{S}^{(j)}$ respectively, the fourth equality follows from the definition of the MNL choice model, the fifth equality follows from (1.1) and the inequality follows from (A.11). Therefore, $(\hat{\lambda}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\beta}})$ is a feasible solution of $(D_2)$.

Next, we analyze the objective value of this feasible solution. The contribution of $\hat{\boldsymbol{\theta}}$ to the objective value can be bounded above as follows:

$$\begin{aligned}\sum_{i=1}^{n} c_i \hat{\theta}_i &= \sum_{i \in \mathcal{N} \setminus \widetilde{\mathcal{N}}_{T+1}} c_i r_i \leq \frac{c_{\min}}{c_{\min} - 1} \sum_{i \in \mathcal{N} \setminus \widetilde{\mathcal{N}}_{T+1}} (c_i - 1) r_i \\ &< \frac{c_{\min}}{c_{\min} - 1} \sum_{i \in \mathcal{N} \setminus \widetilde{\mathcal{N}}_{T+1}} r_i \widetilde{Z}_i \leq \frac{c_{\min}}{c_{\min} - 1} \sum_{i=1}^{n} r_i \widetilde{Z}_i,\end{aligned} \tag{A.13}$$

where $c_{\min} = \min_{i \in \mathcal{N}} c_i$. Here, the first inequality follows from the fact that $x/(x-1)$ decreases monotonically for all $x \geq 1$, the second inequality follows from the fact that $\widetilde{Z}_i > c_i - 1$ for all $i \in \mathcal{N} \setminus \widetilde{\mathcal{N}}_{T+1}$, and the third inequality follows from the fact that $\widetilde{Z}_i \geq 0$ for all $i \in \mathcal{N}$. On the other hand, to bound the contribution to the objective value from the remaining variables,

$$\begin{aligned}T\hat{\lambda} + \sum_{j=1}^{m} b_j \hat{\eta}_j &= T\hat{\lambda} + \sum_{j=1}^{m} \sum_{t:\widetilde{J}_t = j} (\widetilde{R}_t - \hat{\lambda}) = T\hat{\lambda} + \sum_{t=1}^{T} \widetilde{R}_t - T\hat{\lambda} = \sum_{t=1}^{T} \widetilde{R}_t \\ &= \sum_{t=1}^{T} \sum_{i=1}^{n} r_i \pi_i^{\widetilde{J}_t}(\widetilde{S}_t) = \sum_{i=1}^{n} r_i \left( \sum_{t=1}^{T} \pi_i^{\widetilde{J}_t}(\widetilde{S}_t) \right) = \sum_{i=1}^{n} r_i \widetilde{Z}_i,\end{aligned} \tag{A.14}$$

where the first equality follows from the definition of $\hat{\eta}_j$, the fourth equality follows from the definition of $\widetilde{R}_t$ and the last equality follows from the definition of $\widetilde{Z}_i$ in (A.3). Adding together (A.13) and (A.14) yields,

$$\begin{aligned}T\hat{\lambda} + \sum_{j=1}^{m} b_j \hat{\eta}_j + \sum_{i=1}^{n} c_i \hat{\theta}_i &\leq \frac{c_{\min}}{c_{\min} - 1} \sum_{i=1}^{n} r_i \widetilde{Z}_i + \sum_{i=1}^{n} r_i \widetilde{Z}_i \\ &= \frac{2c_{\min} - 1}{c_{\min} - 1} \sum_{i=1}^{n} r_i \widetilde{Z}_i = \frac{2c_{\min} - 1}{c_{\min} - 1} \sum_{i=1}^{n} \sum_{j=1}^{k} r_i \widetilde{q}_i^j,\end{aligned}$$

where the last equality follows from (A.4). Since $(\hat{\lambda}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\beta}})$ is a feasible solution of $(D_2)$, the objective value is at least $z^*_{SLP}$. Therefore,

$$z^*_{SLP} \leq T\hat{\lambda} + \sum_{j=1}^{m} b_j \hat{\eta}_j + \sum_{i=1}^{n} c_i \hat{\theta}_i \leq \frac{2c_{\min} - 1}{c_{\min} - 1} \sum_{i=1}^{n} \sum_{j=1}^{k} r_i \widetilde{q}_i^j.$$

This concludes the proof of the lemma. □

### A.2.3 Proof of Lemma 1.4

*Proof.* The intuition behind Lemma 1.4 is as follows. The contribution of product $i$ to the objective value of (SLP) at $\widetilde{\boldsymbol{q}}$ is $r_i \widetilde{Z}_i$. Let $\mathcal{N}_{T+1}$ be the (random) set of products with inventory remaining after the execution of the algorithm. On the one hand, the expected revenue of the algorithm is at least as much as the expected contribution of products in $\mathcal{N} \setminus \mathcal{N}_{T+1}$ to the objective value, since these products stocked out. On the other hand, we also show that the expected revenue of the algorithm is as much as the expected contribution of the products in $\mathcal{N}_{T+1}$ to the objective value. This second result uses the fact that adding products with revenue higher than the expected revenue of the assortment only increases the expected revenue under the MNL choice model. Adding these two results together returns the half approximation.

We first introduce some additional notation for convenience. Let $Z_{i,t}$ denote the number of units of product $i$ purchased until time $t$ and $\mathcal{N}_t$ denote the set of products with positive inventory remaining at the beginning of the $t$-th time period, i.e., $\mathcal{N}_t := \{i \in \mathcal{N} \mid Z_{i,t-1} \leq c_i - 1\}$ for $t = 1, \ldots, T+1$. We use $A^{(j,S)}$ to denote the set of all products with revenue at least $R_{\pi^j}(S)$, i.e., $A^{(j,S)} := \{i \in \mathcal{N} \mid r_i \geq R_{\pi^j}(S)\}$. Also, we use $\sigma(t)$ to denote the round in which $(J_t, R_t)$ were selected in the customer selection step, i.e., $J_t = \widetilde{J}_{\sigma(t)}$ for all $t = 1, \ldots, T$. Therefore, $\widetilde{S}_{\sigma(t)}$ is the assortment that was selected along with customer $J_t$ in round $\sigma(t)$. Recall that $\widetilde{\boldsymbol{q}}$ is the feasible solution to (SLP) constructed in the customer selection step. The contribution of product $i$ to the objective value at $\widetilde{\boldsymbol{q}}$ is $r_i \widetilde{Z}_i$. Also recall that $\mathcal{N}_{T+1}$ is the set of products with inventory remaining after the execution of the algorithm.

Next, we show that $\textsc{Rev}(\text{A1})$ is at least as much as the expected contribution of the products in $\mathcal{N}_{T+1}$ to the objective value of (SLP) at $\widetilde{\boldsymbol{q}}$. We first show that $R_{\pi^{J_t}}(S_t) \geq R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1})$ for all $t = 1, \ldots, T$. Since $R_t = R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)})$, therefore, $A^{(J_t, \widetilde{S}_{\sigma(t)})} = \{i \in \mathcal{N} \mid r_i \geq R_t\}$ by definition of $\sigma(t)$ and $A^{(j,S)}$. Using these observations, $S_t$ can be written as,

$$S_t = \{i \in \mathcal{N} \mid Z_{i,t-1} \leq c_i - 1, r_i \geq R_t\} = \mathcal{N}_t \cap A^{(J_t, \widetilde{S}_{\sigma(t)})}, \tag{A.15}$$

where the first equality follows from the definition of $S_t$ in Algorithm 1 and the second equality follows from the definitions of $\mathcal{N}_t$ and $A^{(J_t, \widetilde{S}_{\sigma(t)})}$. By definition of $S_t$ in Algorithm 1,

$$r_i \geq R_t \quad \forall i \in S_t. \tag{A.16}$$

In Step 1, the assortment $\widetilde{S}_{\sigma(t)}$ was chosen to be optimal in round $\sigma(t)$, and thus the alternative feasible assortment $\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1}$ could not have been better, i.e.,

$$R_t \geq R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1}). \tag{A.17}$$

Combining (A.16) and (A.17),

$$r_i \geq R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1}) \quad \forall i \in S_t. \tag{A.18}$$

Revenue-maximizing assortments are nested-by-revenue under the MNL choice model, with the revenue of each product in the optimal assortment being at least as much as the expected revenue of the optimal assortment. Therefore, for all $i \in \widetilde{S}_{\sigma(t)}$, $r_i \geq R_t$. By definition of $A^{(J_t, \widetilde{S}_{\sigma(t)})}$, this implies that,

$$\widetilde{S}_{\sigma(t)} \subseteq A^{(J_t, \widetilde{S}_{\sigma(t)})}. \tag{A.19}$$

Furthermore,

$$\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1} \subseteq A^{(J_t, \widetilde{S}_{\sigma(t)})} \cap \mathcal{N}_t = S_t, \tag{A.20}$$

where the first relation follows from (A.19) and the fact that $\mathcal{N}_{T+1} \subseteq \mathcal{N}_t$ for all $t = 1, \ldots, T$ and the equality follows from (A.15). It follows from (A.18) and (A.20) that $\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1} \subseteq S_t$ and $r_i \geq R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1})$ for all $i \in S_t \setminus (\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1})$. Due to the nested-by-revenue property of the MNL choice model, adding products in $S_t \setminus (\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1})$ to $\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1}$ keeps the revenue of the final assortment above $R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1})$. Therefore,

$$R_{\pi^{J_t}}(S_t) \geq R_{\pi^{J_t}}(\widetilde{S}_{\sigma(t)} \cap \mathcal{N}_{T+1}). \tag{A.21}$$

Now, $\textsc{Rev}(\text{A1})$ can be bounded below as follows,

$$\begin{aligned} \textsc{Rev}(\text{A1}) &= \sum_{t=1}^{T} \mathbb{E}\left[\sum_{i \in S_t} r_i \pi_i^{J_t}(S_t)\right] \qquad \text{(A.22)} \\ &\geq \sum_{t=1}^{T} \mathbb{E}\left[\sum_{i \in \mathcal{N}_{T+1} \cap \widetilde{S}_{\sigma(t)}} r_i \pi_i^{J_t}(\mathcal{N}_{T+1} \cap \widetilde{S}_{\sigma(t)})\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{T} \sum_{i \in \mathcal{N}_{T+1}} r_i \pi_i^{J_t}(\widetilde{S}_{\sigma(t)})\right] \\ &= \mathbb{E}\left[\sum_{i \in \mathcal{N}_{T+1}} r_i \widetilde{Z}_i\right], \end{aligned}$$

where the first equality follows (1.2), the first inequality follows from (A.21) and (1.1), the second inequality follows from the substitutability property of the MNL choice model and the fact that $\pi_i^{J_t}(\widetilde{S}_{\sigma(t)}) = 0$ for all $i \in \mathcal{N}_{T+1} \setminus \widetilde{S}_{\sigma(t)}$, and the last equality follows from the definition of $\widetilde{Z}_i$ in (A.3).

Next, we show that $\textsc{Rev}(\text{A1})$ is also as much as the expected contribution of the products in

$\mathcal{N} \setminus \mathcal{N}_{T+1}$ to the objective value. For $i \in \mathcal{N} \setminus \mathcal{N}_{T+1}$, $Z_i = c_i \geq \widetilde{Z}_i$. Therefore,

$$\textsc{Rev}(\mathrm{A1}) = \mathbb{E}\left[\sum_{i=1}^{n} r_i Z_i\right] \geq \mathbb{E}\left[\sum_{i \in \mathcal{N} \setminus \mathcal{N}_{T+1}} r_i Z_i\right] \geq \mathbb{E}\left[\sum_{i \in \mathcal{N} \setminus \mathcal{N}_{T+1}} r_i \widetilde{Z}_i\right], \tag{A.23}$$

where the first equality follows from (1.2).

Finally, adding (A.22) and (A.23) yields,

$$\begin{aligned} 2\textsc{Rev}(\mathrm{A1}) &\geq \mathbb{E}\left[\sum_{i \in \mathcal{N}_{T+1}} r_i \widetilde{Z}_i\right] + \mathbb{E}\left[\sum_{i \in \mathcal{N} \setminus \mathcal{N}_{T+1}} r_i \widetilde{Z}_i\right] \\ &= \mathbb{E}\left[\sum_{i \in \mathcal{N}} r_i \widetilde{Z}_i\right] = \sum_{i \in \mathcal{N}} r_i \widetilde{Z}_i. \end{aligned}$$

Therefore,

$$\textsc{Rev}(\mathrm{A1}) \geq \frac{1}{2} \sum_{i \in \mathcal{N}} r_i \widetilde{Z}_i.$$

□

### A.2.4 Tight example for Theorem 1.2

In this section we show that our analysis of the performance of A1 is tight by means of a problem instance on which the worst case guarantee of the policy is tight. The problem instance is as follows:

Suppose there are $n = 6$ products, each with $c_i = T/4$ units of each product $i$ at the start. The revenue earned from the first 4 products is 1, but for tie-break purposes, negligible terms are added in to ensure $r_1 > r_2 > r_3 > r_4$. There are 4 customer types, with $T/4$ customers of each type. Each customer can be thought of as having a preference order over the products. The preference orders for the four customer types are (6,1), (1,2), (5,3) and (1,3,4) respectively. This implies that, for instance, type 1 customers always prefer product 6 if available, product 1 if it is available when product 6 is not available, and opt for no-purchase when none of them are available. Such a model can be approximated by an MNL model by setting attractiveness parameters such that

$u_6^1 >> u_1^1 >> 1 >> \max\{u_2^1, u_3^1, u_4^1, u_5^1\}$. Additionally, for this instance, we set $u_1^2 >> \max\{u_1^1, u_1^4\}$ and $u_3^4 >> u_3^3$. By this assumption we imply that type-2 customers prefer product 1 more than type-1 and type-4 customers, while type-4 customers prefer product 3 more than type-3 customers. In Table A.1 we provide a concrete set of parameters for this problem instance, where $M$ and $\epsilon$ can be assumed to be a large and small constant respectively. Specifically, we can let $\epsilon = T^{-2}/12$ and $M = 3T^2$ to make all our calculations work.

| $i$ | 1 | 2 | 3 | 4 | 5 | 6 |
|---|---|---|---|---|---|---|
| $r_i$ | $1 + 36\epsilon$ | $1 + 24\epsilon$ | $1 + 12\epsilon$ | 1 | $24\epsilon$ | $12\epsilon$ |
| $u_i^1$ | $M$ | $\epsilon$ | $\epsilon$ | $\epsilon$ | $\epsilon$ | $M^2$ |
| $u_i^2$ | $M^4$ | $M$ | $\epsilon$ | $\epsilon$ | $\epsilon$ | $\epsilon$ |
| $u_i^3$ | $\epsilon$ | $\epsilon$ | $M$ | $\epsilon$ | $M^2$ | $\epsilon$ |
| $u_i^4$ | $M^3$ | $\epsilon$ | $M^2$ | $M$ | $\epsilon$ | $\epsilon$ |

Table A.1: Instance parameters

The optimal policy for this problem is to retire products 5 and 6 in the beginning, and then visit all customers of type 1, 2, 3, 4 in that order. Table A.2 describes how the policy would work. Type-$j$ customers purchase all units of product $j$ for $j = \{1, 2, 3, 4\}$ and the optimal revenue is $T$.

| Phase | Length | $J_t^*$ | $S_t^*$ | $X_t^*$ | Revenue |
|---|---|---|---|---|---|
| 1 | $T/4$ | 1 | $\{1, 2, 3, 4\}$ | 1 | $T/4 + O(T^{-1})$ |
| 2 | $T/4$ | 2 | $\{2, 3, 4\}$ | 2 | $T/4 + O(T^{-1})$ |
| 3 | $T/4$ | 3 | $\{3, 4\}$ | 3 | $T/4 + O(T^{-1})$ |
| 4 | $T/4$ | 4 | $\{4\}$ | 4 | $T/4 + O(T^{-1})$ |
| | | | **Total revenue:** | | $T + O(T^{-1})$ |

Table A.2: Optimal policy

Tables A.3 and A.4 describe how A1 performs on this problem instance. In the customer selection step, customers types are chosen in the order 2, 4, 3, 1. In the first set of $T/4$ rounds, all type-2 customers are selected with assortment $\{1\}$ as product 1 has the highest revenue and type-2 customers like product 1 the most. Since none of the remaining customer types have product 2 in their preference order, product 2, it remains unsold and product 3 becomes the next best product. Therefore, in the second set of $T/4$ rounds, type-4 customers are selected as they like product 3 the most. The assortment selected is $\{2, 3\}$ as adding product 2 which has higher revenue than product 3 to the assortment cannot hurt. At this point, none of the remaining customer types refer product 4, thus product 4 remains unsold as well. In the next 2 sets of $T/4$ rounds, type-3 and type-1 customers are selected in that order due to their preference for product 5 and product 6 respectively. Table A.3 provides a concise representation of this step. The sequence to visit the customers when they are sorted by $\widetilde{R}_t$ is exactly the opposite of the sequence they were selected in. Type-1 customers are visited first and offered all products. They purchase their highest preference, i.e., product 6. Then type-3 customers are visited, who too can purchase highest preference, i.e., product 5. At this point product 4 is retired since $r_4 < 1 + \epsilon$. Type-4 customers are visited next and they too can purchase their highest preference, i.e., product 1. Finally, type-2 customers are visited. However, they have no products available, since all units of product 1 were purchased by type-4 customers and the remaining products were retired since their revenue was less than $1 + 3\epsilon$. The expected revenue for A1 on this instance is therefore $T/4$, which is 1/4 times the optimal revenue, making our analysis of the performance of the policy tight.

| Phase | Length | Products Avail. | $\widetilde{J}_t$ | $\widetilde{S}_t$ | $\widetilde{R}_t$ | Revenue |
|---|---|---|---|---|---|---|
| 1 | $T/4$ | {1,2,3,4,5,6} | 2 | $\{1\}$ | $1 + 3\epsilon$ | $T/4 + O(T^{-1})$ |
| 2 | $T/4$ | {2,3,4,5,6} | 4 | $\{2, 3\}$ | $1 + \epsilon$ | $T/4 + O(T^{-1})$ |
| 3 | $T/4$ | {2,4,5,6} | 3 | $\{2, 4, 5\}$ | $2\epsilon$ | $O(T^{-1})$ |
| 4 | $T/4$ | {2,4,6} | 1 | $\{2, 4, 6\}$ | $\epsilon$ | $O(T^{-1})$ |
| | | | | | **Total revenue:** | $T/2 + O(T^{-1})$ |

Table A.3: Customer selection

| Phase | Length | $J_t$ | $S_t$ | $X_t$ | Revenue |
|---|---|---|---|---|---|
| 1 | $T/4$ | 1 | $\{1, 2, 3, 4, 5, 6\}$ | 6 | $O(T^{-1})$ |
| 2 | $T/4$ | 3 | $\{1, 2, 3, 4, 5\}$ | 5 | $O(T^{-1})$ |
| 3 | $T/4$ | 4 | $\{1, 2, 3\}$ | 1 | $T/4 + O(T^{-1})$ |
| 4 | $T/4$ | 2 | $\emptyset$ | 0 | 0 |
| | | | | **Total:** | $T/4 + O(T^{-1})$ |

Table A.4: Final algorithm run

## A.3 Omitted Proofs of Section 1.4

### A.3.1 Proof of Lemma 1.5

*Proof.* Throughout this proof we fix some customer type $j$. First, we show that there is a unique tuple $(\mathcal{O}^j(\boldsymbol{q}), \widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}))$ that satisfies the following properties:

(a) $\mathcal{O}^j(\boldsymbol{q})$ is a collection of distinct nested assortments.

(b) $\widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q})$ satisfies $\widetilde{\tau}^j_S > 0$ for all $S \in \mathcal{O}(\boldsymbol{q})$.

(c) $(\mathcal{O}^j(\boldsymbol{q}), \widetilde{\boldsymbol{\tau}}^j(\boldsymbol{q}))$ satisfy (1.5).

Suppose $(\mathcal{O}^j, \widetilde{\boldsymbol{\tau}}^j)$ is a tuple that satisfies the three above properties. Since $\mathcal{O}^j$ is a collection of nested assortments by property (a), let $\mathcal{O}^j = \{O_1, \dots, O_K\}$ where $O_1 \supset \cdots \supset O_K$. Define $I_l := O_l \setminus O_{l+1}$ for $l = 1, \dots, K-1$ and $I_K = O_K \cup \{0\}$. Note that none of the sets $I_l$ are empty for $l = 1, \dots, K$, since the assortments are all distinct and $I_k$ is non-empty even if $O_K = \emptyset$ since it contains the no-purchase option. We now show that there exist constants $v_l$ for $l = 1, \dots, K$ such that $\forall i \in I_l$, $\frac{q^j_i}{u^j_i} = v_l$. To show this, fix some $l \in [K]$, then for any $i_1, i_2 \in I_l$,

$$\frac{q^j_{i_1}}{u^j_{i_1}} = \frac{1}{u^j_{i_1}} \sum_{k=1}^{l} \widetilde{\tau}_{O_k}(\boldsymbol{q}) \pi^j_{i_1}(O_k) = \frac{1}{u^j_{i_2}} \sum_{k=1}^{l} \widetilde{\tau}_{O_k}(\boldsymbol{q}) \pi^j_{i_2}(O_k) = \frac{q^j_{i_2}}{u^j_{i_2}},$$

, where the first and third equality follow from property (c) and the second equality follows from the definition of the MNL choice model in (1.3). Further, we show that $v_1 < v_2 < \cdots < v_K$. For this, fix some $l \leq K-1$, $i_l \in I_l$, $i_{l+1} \in I_{l+1}$, then

$$\begin{aligned} v_l = \frac{q^j_{i_l}}{u^j_{i_l}} &= \frac{1}{u^j_{i_l}} \sum_{k=1}^{l} \widetilde{\tau}_{O_k}(\boldsymbol{q}) \pi^j_{i_l}(O_k) = \frac{1}{u^j_{i_{l+1}}} \sum_{k=1}^{l} \widetilde{\tau}_{O_k}(\boldsymbol{q}) \pi^j_{i_{l+1}}(O_k) \\ &< \frac{1}{u^j_{i_{l+1}}} \sum_{k=1}^{l+1} \widetilde{\tau}_{O_k}(\boldsymbol{q}) \pi^j_{i_{l+1}}(O_k) = \frac{q^j_{i_{l+1}}}{u^j_{i_{l+1}}} = v_{l+1}, \end{aligned}$$

where the third equality follows from the definition of the MNL choice model in (1.3) and the inequality follows from the fact that $\widetilde{\tau}^j_{O_{l+1}} > 0$ (property (b)) and $\pi^j_{i_{l+1}}(O_{l+1}) > 0$ since $i_{l+1} \in O_{l+1}$. The fact that $v_1 < v_2 < \cdots < v_K$ implies that $I_l$ is uniquely the set which contains all products such that $q^j_i / u^j_i = v_l$, i.e.

$$I_l = \left\{ i \mid \frac{q^j_i}{u^j_i} = v_l \right\}$$

Consequently, the assortment $O_l$ and offering time $\widetilde{\tau}_{O_l}$ are also unique and defined as,

$$O_l = \bigcup_{k=K}^{l} I_k$$
$$\widetilde{\tau}_{O_l} = \frac{v_l - v_{l-1}}{\pi^j_0(O_l)}.$$

Therefore, if the output of Procedure 1 satisfies properties (a)-(c), it is the unique collection of nested assortments and offering times to do so.

Next, we analyze properties of Procedure 1. We first establish some notation. We assume that the last iteration of the while loop in Procedure 1 is iteration $K$. For every iteration $k$, we use the superscript $(k)$ to denote the state of variables in or at the start of the iteration $k$. We use the superscript $(K + 1)$ to denote the state of variables after the termination of the loop.

We first show which that the output of Procedure 1 satisfies properties (a)-(c):

(a) Observe that $S^{(1)} \supset S^{(2)} \supset \ldots$, since in every iteration products are only removed from $S$ and never added back. Since $\mathcal{O}^j(\boldsymbol{q})$ is the collection of all states of $S$, it follows that it is a collection of nested sets, thus proving the statement.

(b) We consider two cases for $S \in \mathcal{O}^j(\boldsymbol{q})$: $S \neq \emptyset$ and $S = \emptyset$. If $S \neq \emptyset$, then this assortment was added to $\mathcal{O}^j(\boldsymbol{q})$ during one of the iterations of the loop. Let that be iteration $k$. From the update rule of $S$, it follows that $\hat{q}^{(k)}_i > 0$ for all $i \in S^{(k)}$, which in turn implies that $\Delta^{(k)} > 0$. Therefore, $\widetilde{\tau}^j_{S^{(k)}}(\boldsymbol{q}) = \Delta^{(k)} > 0$.

Next we consider the second case where $S = \emptyset$. This case occurs if $\hat{q}_0 > 0$ after the termination of the loop. For this case, $\widetilde{\tau}^j_{\emptyset}(\boldsymbol{q}) = \hat{q}_0 > 0$, which proves the statement.

(c) Fix $i \in \mathcal{N} \cup \{0\}$. We consider two cases: $i \in \mathcal{N}$ and $i = 0$. First, we look at the case where $i \in \mathcal{N}$. Let $k_i$ be the iteration of the loop where $i$ is removed from $S$, then,

$$\hat{q}_i^{(k_i+1)} = 0 \Rightarrow \hat{q}_i^{(1)} - \sum_{k=1}^{k_i} \Delta^{(k)} \pi_i^j(S^{(k)}) = 0 \Rightarrow q_i^j = \sum_{k=1}^{k_i} \Delta^{(k)} \pi_i^j(S^{(k)})$$
$$\Rightarrow q_i^j = \sum_{S \in O^j(\boldsymbol{q})} \tilde{\tau}_S^j(\boldsymbol{q}) \pi_i^j(S),$$

where the first equality follows from the fact that $i \notin S^{k_i+1}$, the second equality follows from the update rule of $\hat{q}_i$, the third equality follows from the initialization of $\hat{q}_i$ and the fourth equality follows from the fact that $i \notin S$ if $S \notin \{S^{(1)}, \ldots, S^{(k_i)}\}$. This concludes the current case.

Next we consider the case where $i = 0$. We first show that $\hat{q}_0^{(K+1)} \geq 0$. From the loop condition it follows that $S^{(K)}$ is non-empty. Fix any $i \in S^{(K)}$, then $\hat{q}_0^{(K+1)}$ can be bounded below as,

$$\begin{aligned}
\hat{q}_0^{(K+1)} &= q_0^j - \sum_{k=1}^{K} \Delta^{(k)} \pi_0^j(S^{(k)}) \\
&= q_0^j - \sum_{k=1}^{K} \Delta^{(k)} \pi_i^j(S^{(k)})/u_i^j \geq q_i^j/u_i^j - \sum_{k=1}^{K} \Delta^{(k)} \pi_i^j(S^{(k)})/u_i^j \\
&= 0,
\end{aligned}$$

where the first equality follows from the update rule and initialization for $\hat{q}_0$, the second equality follows from the MNL choice probabilities defined in (1.3), the inequality follows from the MNL constraints of (SLP),the third equality follows from previous case. If $\hat{q}_0^{(K+1)}$ = 0, then the desired result follows from the previous case. On the other hand if $\hat{q}_0^{(K+1)} > 0$, then setting $\tilde{\tau}_\emptyset^j(\boldsymbol{q}) = \hat{q}_0^{(K+1)}$ is sufficient.

Therefore, the output of Procedure 1 uniquely satisfies properties (a)-(c).

Finally, we consider the runtime. Since at least one product is removed from $S$ in every iteration of the loop, the algorithm will terminate in at most $n$ iterations. This concludes the proof of the lemma. $\square$

### A.3.2 Proof of Theorem 1.8

*Proof.* By the theorem assumption, $c_{\min} \geq \frac{2}{\pi\epsilon^2}$, and thus from (1.8) it follows that

$$\gamma(c_{\min}) \geq \gamma\left(\frac{2}{\pi\epsilon^2}\right) = 1 - \epsilon/2, \tag{A.24}$$

where the inequality follows from the fact that $\gamma(\cdot)$ is non-decreasing. By the theorem assumption, $T \geq \frac{2m}{\epsilon}$ and thus

$$\frac{T}{T+m} \geq \frac{\frac{2m}{\epsilon}}{\frac{2m}{\epsilon}+m} = \frac{1}{1+\epsilon/2} \geq 1 - \epsilon/2. \tag{A.25}$$

Since $\hat{\boldsymbol{q}}$ is a feasible solution, Lemma 1.11 can be applied to our policy A2. The expected revenue of A2 can thus be bounded below as,

$$\begin{aligned} \textsc{Rev}(\text{A2}) &\geq \gamma(c_{\min}) \sum_{i=1}^{n} \sum_{j=1}^{m} r_i \hat{q}_i^j = \gamma(c_{\min}) \sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \sum_{S \in O(\boldsymbol{q})} a_{k,S}^j r_i \pi_i^j(S) \\ &\geq \frac{\gamma(c_{\min}) T}{T+m} z_{SLP}^* \geq (1-\epsilon/2)^2 z_{SLP}^* \geq (1-\epsilon) z_{SLP}^*, \end{aligned}$$

where the first inequality follows from Lemma 1.11, the equality follows from the definition of $\hat{q}_i^j$, the second inequality follows from Lemma 1.10(iii), and the third inequality follows from (A.24) and (A.25). $\square$

### A.3.3 Proof of Lemma 1.10

*Proof.* *(i)* Let $\bar{Y}_j$ and $\bar{\tau}_S^j$ denote the value of variables $Y_j$ and $\tau_S^j$ at the end of Step 2a in Algorithm 2. We first show that $\bar{Y}_j \in \mathbb{Z}$ and $\bar{Y}_j \leq b_j$ for all $j \in \mathcal{L}$. From the update steps of $\tau_\emptyset^j$ and $Y_j$ in Step

2a it follows that,

$$\begin{aligned}\bar{Y}_j &= \sum_{S\in O} \bar{\tau}^j_S = \bar{\tau}^j_\emptyset + \sum_{S\in O\setminus\{\emptyset\}} \bar{\tau}^j_S \\ &= \left\lceil \sum_{S\in O\setminus\{\emptyset\}} \tau^j_S \right\rceil - \sum_{S\in O\setminus\{\emptyset\}} \tau^j_S + \sum_{S\in O\setminus\{\emptyset\}} \tau^j_S = \left\lceil \sum_{S\in O\setminus\{\emptyset\}} \tau^j_S \right\rceil,\end{aligned} \tag{A.26}$$

where the first equality follows from the update step of $Y_j$ and the third equality follows from the definition of $\tau^j_\emptyset$. This implies that $\bar{Y}_j \in \mathbb{Z}$. Recall that $\boldsymbol{\tau}^j$ is derived from the output of Procedure 1, $\widetilde{\boldsymbol{\tau}}^j$. It follows that,

$$\sum_{S\in O\setminus\{\emptyset\}} \tau^j_S = \sum_{S\in O^j\setminus\{\emptyset\}} \widetilde{\tau}^j_S \le b_j,$$

where the inequality follows from Lemma 1.5 and (1.5). Combining with (A.26) implies that $\bar{Y}_j \le b_j$. The same must also hold for our final customer selections, $Y_j$, since after Step 2a, the value of $Y_j$ is only decreased by integer amounts. Finally, $\sum_{j=1}^m Y_j \le T$ follows directly from the termination condition of the while loop in Step 2c of Algorithm 2.

*(ii)* Fix $j \in \mathcal{L}$. Let $\bar{Y}_j$ and $\bar{\tau}^j_S$ retain their definitions from the previous part. Suppose $k^*$ and $l^*$ denote the values of $k$ and $l$ at the end of Step 2b of Algorithm 2. We first show the loop invariant that $\sum_{k,S} a^j_{k,S} + \sum_S \hat{\tau}^j_S = \bar{Y}_j$. We then use the invariant to show that $k^* > \bar{Y}_j$ and $l^* > |O|$, which we use to conclude the proof.

**Invariant:** $\sum_{k=1}^{\bar{Y}_j} \sum_{S\in O} a^j_{k,S} + \sum_{S\in O} \hat{\tau}^j_S = \bar{Y}_j$

**Initialization:** At initialization, $a^j_{k,S} = 0$ and $\hat{\tau}^j_S = \bar{\tau}^j_S$ for all $k \in [\bar{Y}_j], S \in O$. Since, $\bar{Y}_j = \sum_{S\in O} \bar{\tau}^j_S$, it follows that the invariant is satisfied at initialization.

**Maintenance:** To see that the invariant is maintained, note that $\sum_{k,S} a^j_{k,S}$ is increased by $\Delta$ while $\sum_S \hat{\tau}^j_S$ is decreased by $\Delta$ in every iteration, keeping the total sum constant. $\square$

From the update rules for $k$ and $l$ it follows that, $\sum_{S\in O} \bar{a}^j_{k,S} = 1 \Leftrightarrow k^* > k$; and $\hat{\tau}^j_{O_l} = 0 \Leftrightarrow$

$l^* > l$. From these two facts it follows that,

$$k^* \le \bar{Y}_j \Leftrightarrow \sum_{k=1}^{\bar{Y}_j} \sum_{S \in O} \bar{a}_{k,S}^j < \bar{Y}_j \Leftrightarrow \sum_{S \in O} \hat{\tau}_S^j > 0 \Leftrightarrow l^* \le |O|, \tag{A.27}$$

where the second implication uses the loop invariant. From (A.27) it follows that if $k^* > \bar{Y}_j$, then $l^* > |O|$ and vice versa. Since at least one of these two conditions must be true for the while loop to terminate, both these conditions are true when the while loop terminates. We can now prove both parts.

(a) Since $k^* > \bar{Y}_j$, $\sum_{S \in O} \bar{a}_{k,S}^j = 1$ for all $k \in [\bar{Y}_j]$.

(b) Since $l^* > |O|$, $\hat{\tau}_S^j = 0$ for all $S \in O$, which implies that $\sum_{k=1}^{\bar{Y}_j} a_{k,S}^j = \bar{\tau}_S^j$. Fix any $S \in O \setminus \{\emptyset\}$, then

$$\sum_{k=1}^{Y_j} a_{k,S}^j \le \sum_{k=1}^{\bar{Y}_j} a_{k,S}^j = \bar{\tau}_S^j = \tau_S^j,$$

where the inequality follows from the fact that $Y_j \le \bar{Y}_j$ and the second equality follows from the fact that $\bar{\tau}_S^j = \tau_S^j$ for all $S \in O \setminus \{\emptyset\}$.

*(iii)* Let $\bar{Y}_j$ and $\bar{\tau}_S^j$ retain their definitions from the previous part. Let $P_i = \{(j,k) \mid j \in \mathcal{L}, k \in [\bar{Y}_j]\}$ and $P_f = \{(j,k) \mid j \in \mathcal{L}, k \in [Y_j]\}$ be the customers selected before and after Step 2c is complete, respectively. If $|P_i| \le T$, $P_f = P_i$ and the statement follows trivially. Henceforth, we assume $|P_i| > T$. For this case, $|P_f| = T$. We can bound the cardinality of $P_i$ above as follows,

$$\begin{aligned} |P_i| &= \sum_{j=1}^{m} \bar{Y}_j = \sum_{j=1}^{m} \left\lceil \sum_{S \in O} \tau_S^j \right\rceil \le m + \sum_{j=1}^{m} \sum_{S \in O} \tau_S^j \\ &= m + \sum_{j=1}^{m} \sum_{S \in O} \sum_{i=0}^{n} \tau_S^j \pi_i^j(S) = m + \sum_{j=1}^{m} \sum_{i=0}^{n} q_i^j \le m + T \end{aligned} \tag{A.28}$$

where the second equality follows from (A.26), the first inequality follows from the fact that $\lceil x \rceil \le x+1$, the fourth equality follows from (1.6) and the second inequality follows from the time horizon constraint of (SLP). Observe that $P_f$ is obtained from $P_i$ by iteratively removing the customer (indexed by $(j,k)$) that minimizes $\sum_{i=1}^{n}\sum_{S\in O} a_{k,S}^{j} r_i \pi_i^j(S)$. This implies that the average value of $\sum_{i=1}^{n}\sum_{S\in O} a_{k,S}^{j} r_i \pi_i^j(S)$ for customers in $P_f$ is higher than it is for customers in $P_i$. Therefore,

$$\begin{aligned} z_{SLP}^{*} &= \sum_{i=1}^{n}\sum_{j=1}^{m} r_i q_i^j = \sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{S\in O} r_i \tau_S^j \pi_i^j(S) = \sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{S\in O}\sum_{k=1}^{\bar{Y}_j} a_{k,S}^{j} r_i \pi_i^j(S) \\ &\le \frac{|P_i|}{|P_f|}\sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{S\in O}\sum_{k=1}^{Y_j} a_{k,S}^{j} r_i \pi_i^j(S) \le \frac{T+m}{T}\sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{S\in O}\sum_{k=1}^{Y_j} a_{k,S}^{j} r_i \pi_i^j(S), \end{aligned}$$

where the first equality follows from the fact that $\boldsymbol{q}$ is an optimal solution of (SLP), the second equality follows from the fact that $\boldsymbol{q}$ satisfies (1.6), the third equality follows from part (ii), the first inequality follows from the fact that the average in $P_f$ is greater than in $P_i$, the second inequality follows from the fact that $|P_f| = T$ and (A.28). □

### A.3.4 Proof of Lemma 1.11

*Proof.* We first construct a random variable $\hat{Z}_i(\mathscr{A})$ for every product $i$, which is a sum of independent Bernoulli random variables, and show that sales of product $i$, i.e., $Z_i(\mathscr{A})$, stochastically dominate $\min\{\hat{Z}_i(\mathscr{A}), c_i\}$. Moreover, expected value of $\hat{Z}_i(\mathscr{A})$ is equal to $\sum_{j=1}^{m}\sum_{k=1}^{Y_j}\sum_{S\in O(\boldsymbol{q})} a_{k,S}^{j} r_i \pi_i^j(S)$. We then conclude the proof by relating expected sales of product $i$ to $\mathbb{E}[\hat{Z}_i(\mathscr{A})]$ through $\mathbb{E}[\min\{c_i, \hat{Z}_i(\mathscr{A})\}]$.

**Constructing $\hat{Z}_i$.** For $i \in \mathcal{N}$, $j \in \mathcal{L}$, and $k = 1, \dots, Y_j$, let $\rho_{i,k}^j$ denote the probability that customer $(j,k)$ purchases product $i$ when offered assortment $D_k^j$. Therefore,

$$\rho_{i,k}^j := \mathbb{E}[\pi_i^j(D_k^j)] = \sum_{S\in O} a_{S,k}^{j} \pi_i^j(S), \tag{A.29}$$

where the second equality follows from the fact that $D_k^j \sim Categorical((a_{k,S}^j)_{S\in O})$. Now we define the Bernoulli random variables that are the building blocks of $\hat{Z}_i$. Let $B_{i,k}^j \sim Bernoulli(\rho_{i,k}^j)$, and define $\hat{Z}_i$ as

$$\hat{Z}_i := \sum_{j=1}^{m} \sum_{k=1}^{Y_j} B^j_{i,k}. \tag{A.30}$$

The expected value of $\hat{Z}_i$ can be written as,

$$\mathbb{E}[\hat{Z}_i] = \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \mathbb{E}[B^j_{i,k}] = \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \rho^j_{i,k} = \sum_{j=1}^{m} \sum_{k=1}^{Y_j} \sum_{S \in O} a^j_{S,k} \pi^j_i(S), \tag{A.31}$$

where the first equality follows from the definition of $\hat{Z}_i$ in (A.30), the second equality follows from the definition of $B^j_{i,k}$ and the third equality follows from (A.29). For $\hat{Z}_i$ constructed as above, we show that $Z_i$ stochastically dominates $\min\{\hat{Z}_i, c_i\}$

**Lemma A.1.** *For all $i \in \mathcal{N}$, $Z_i$ stochastically dominates* $\min\{\hat{Z}_i, c_i\}$.

We provide the proof of Lemma A.1 in Appendix A.3.5.

**Lower bound for** $\min\{\hat{Z}_i, c_i\}$**.** We use the following lemma to bound the loss incurred due to supply (inventory) constraints. Readers should refer to [72] for a proof of the lemma.

**Lemma A.2** (Theorem 2.2 [72])**.** *Let $D_1, D_2, \cdots, D_T$ be independent Bernoulli random variables. Then, for any $c \geq \sum_{t=1}^{T} \mathbb{E}[D_t]$,*

$$\mathbb{E}\left[\min\left\{c,\ \sum_{t=1}^{T} D_t\right\}\right] \geq \gamma(c)\mathbb{E}\left[\sum_{t=1}^{T} D_t\right].$$

First, we verify that $\mathbb{E}[\hat{Z}_i]$ is bounded above by $c_i$. Suppose the inputs to Stage II correspond to optimal solution $\boldsymbol{q}$ of (SLP), then,

$$\begin{aligned}\mathbb{E}[\hat{Z}_i] &= \sum_{j=1}^{m}\sum_{k=1}^{Y_j}\sum_{S\in\mathcal{O}} a^j_{S,k}\pi^j_i(S)\\ &= \sum_{j=1}^{m}\sum_{k=1}^{Y_j}\sum_{S\in\mathcal{O}\setminus\{\emptyset\}} a^j_{S,k}\\ &\le \sum_{j=1}^{m}\sum_{S\in\mathcal{O}\setminus\{\emptyset\}} \tau^j_S\pi^j_i(S) = \sum_{j=1}^{m} q^j_i \le c_i,\end{aligned}$$

where the first equality follows from (A.31), the second equality follows from the fact that $\pi^j_i(\emptyset) = 0$ for all $i \in \mathcal{N}$, the first inequality follows from Lemma 1.10(ii), the third equality follows from the fact that $\boldsymbol{q}$ has consistent nested assortments and (1.6), and the second inequality follows since $\boldsymbol{q}$ is a feasible solution of (SLP) and satisfies the supply constraints. Since $\hat{Z}_i$ is a sum of Bernoulli random variables and $c_i \ge \mathbb{E}[\hat{Z}_i]$, we can use Lemma A.2 to bound $\mathbb{E}[\min\{\hat{Z}_i, c_i\}]$. Therefore, the expected revenue earned by A2 can be bounded below as

$$\begin{aligned}\textsc{Rev}(\text{A2}) &= \sum_{i=1}^{n} r_i\mathbb{E}[Z_i] \ge \sum_{i=1}^{n} r_i\mathbb{E}[\min\{\hat{Z}_i, c_i\}] \ge \sum_{i=1}^{n} r_i\gamma(c_i)\mathbb{E}[\hat{Z}_i]\\ &\ge \gamma(c_{\min})\sum_{i=1}^{n} r_i\mathbb{E}[\hat{Z}_i]\\ &= \gamma(c_{\min})\sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{k=1}^{Y_j}\sum_{S\in\mathcal{O}} r_i a^j_{k,S}\pi^j_i(S),\end{aligned}$$

where the first inequality follows from the stochastic dominance established in Lemma A.1, the second inequality follows from Lemma A.2, the third inequality follows from the fact that $\gamma(\cdot)$ is monotonically increasing, and the equality follows from (A.31). This concludes the proof. □

### A.3.5 Proof of Lemma A.1

*Proof.* We fix a product $i$ and use a coupling argument to prove the stochastic dominance result. Let $(U^j_k)_{k=1}^{Y_j}$ be a sequence of i.i.d. uniform random numbers in $[0, 1]$ for $j \in \mathcal{L}$, and let $(V_t)_{t=1}^{T}$

also be a sequence of i.i.d. uniform random numbers in $[0,1]$. We will use these to couple $Z_i$ and $\hat{Z}_i$.

**Constructing customer sequence.** We use $U_k^j$ to construct $D_k^j$. Recall that

$$D_k^j \sim Categorical((a_{S,k}^j)_{S\in O}).$$

The construction is as follows: if $l$ is the index such that

$$\sum_{l'=1}^{l-1} a_{O_{l'},k}^j < U_k^j \le \sum_{l'=1}^{l} a_{O_{l'},k}^j,$$

then $D_k^j = O_l$. Recall that $J_t$ denotes the customer visited in time $t$. Let $\sigma^j(k)$ be the position in the customer sequence for the $D_k^j$ that we just simulated, i.e., $J_{\sigma^j(k)} = j$.

**Constructing $\hat{Z}_i$.** Recall that $\hat{Z}_i = \sum_{j=1}^m \sum_{t=k}^{Y_j} B_{i,k}^j$. We construct $B_{i,k}^j$ using $U_k^j$ and $V_{\sigma^j(k)}$, by letting

$$B_{i,k}^j = \mathbb{I}[V_{\sigma^j(k)} \le \pi_i^j(D_k^j)], \tag{A.32}$$

where $D_k^j$ is determined using $U_k^j$ as described above. Recall that $B_{i,k}^j$ should have a $Bernoulli(\rho_{i,k}^j)$ distribution according to its definition. This follows from,

$$\begin{aligned}
\mathbb{E}[B_{i,k}^j] &= \mathbb{P}[V_{\sigma^j(k)} \le \pi_i^j(D_k^j)] \\
&= \sum_{S\in O} \mathbb{P}[D_k^j = S]\mathbb{P}[V_{\sigma^j(k)} \le \pi_i^j(D_k^j) \mid D_k^j = S] \\
&= \sum_{S\in O} a_{S,k}^j \pi_i^j(S) = \rho_{i,k}^j,
\end{aligned}$$

where the last equality follows from (A.32) and the fact that $U_k^j$ and $V_{\sigma^j(k)}$ are uniform random variables on $[0,1]$.

**Constructing $Z_i$.** Recall that $Z_i = \sum_{t=1}^{T} \mathbb{I}[X_t = i]$. By definition, $X_t = i$ with probability $\pi_i^{J_t}(S_t)$ for all $i \in \mathcal{N} \cup \{0\}$, where $S_t$ is the random variable denoting the set of available products offered to the customer in time period $t$. We construct $X_t$ using $V_t$ as,

$$X_t = \begin{cases} i \text{ if } V_t \le \pi_i^{J_t}(S_t) \\ i' \text{ w.p. } \pi_{i'}^{J_t}(S_t \setminus \{i\}) \text{ for all } i' \in S_t \setminus \{i\} \cup \{0\} \text{ if } V_t > \pi_i^{J_t}(S_t). \end{cases}$$

We now confirm that $X_t$ takes values of $S_t \cup \{0\}$ according to the distribution $\boldsymbol{\pi}^{J_t}(S_t)$. First, observe that $\mathbb{P}[X_t = i] = \mathbb{P}[V_t \le \pi_i(S_t)] = \pi_i(S_t)$ by construction. Similarly, for $i' \in S_t \setminus \{i\} \cup \{0\}$, it follows that

$$\mathbb{P}[X_t = i'] = \mathbb{P}[V_t > \pi_i^{J_t}(S_t)]\pi_{i'}^{J_t}(S_t \setminus \{i\}) = (1 - \pi_i^{J_t}(S_t))\pi_{i'}^{J_t}(S_t \setminus \{i\}) = \pi_{i'}^{J_t}(S_t),$$

where the last equality follows from the MNL choice probabilities defined in (1.3).

**Proving stochastic dominance.** Note that the retirement time for product $i$, i.e., the last time period the product is offered in if inventory is available, is given by $d_i = |\{(j,k) \mid i \in D_k^j\}|$. Let $\widetilde{S}_t := \{i \mid d_i \ge t\}$, then note that $\widetilde{S}_{\sigma^j(k)} = D_k^j$. Observe that $S_t \subseteq \widetilde{S}_t$ for all $t$. Therefore, by the substitutability property of MNL, $\pi_i^{J_t}(S_t) \ge \pi_i^{J_t}(\widetilde{S}_t)$ if $i \in S_t$. This implies that if $i \in S_t$, then

$$\mathbb{I}[V_t \le \pi_i^{J_t}(\widetilde{S}_t), i \in S_t] \le \mathbb{I}[V_t \le \pi_i^{J_t}(S_t), i \in S_t] = \mathbb{I}[X_t = i]. \tag{A.33}$$

Next, we show that $Z_i \ge \min\{\hat{Z}_i, c_i\}$ under our coupling which directly implies the stochastic dominance result. When $Z_i = c_i$, this is trivially true. When $Z_i < c_i$, it means that $i$ did not stock out ($i \in S_t$) until it was retired (for all $t = 1, \ldots, d_i$).

$$\begin{aligned}\hat{Z}_i &= \sum_{j=1}^{m}\sum_{k=1}^{Y_j} B^j_{i,k} = \sum_{j=1}^{m}\sum_{k=1}^{Y_j} \mathbb{I}[V_{\sigma^j(k)} \le \pi^j_i(D^j_k)] = \sum_{t=1}^{d_i} \mathbb{I}[V_t \le \pi^{J_t}_i(\widetilde{S}_t)] \\ &\le \sum_{t=1}^{d_i} \mathbb{I}[V_t \le \pi^{J_t}_i(S_t)] = \sum_{t=1}^{d_i} \mathbb{I}[X_t = i] = Z_i,\end{aligned}$$

where the first equality follows from (A.30), the second equality follows from (A.32), the third equality follows since $i \in \widetilde{S}_t$ only for $t = 1, \ldots, d_i$ and from the definition of $\widetilde{S}_t$ and the first inequality follows from (A.33) and the assumption that $i \in S_t$ for $t = 1, \ldots, d_i$. ☐

## A.4 Proof of Theorem 1.9

*Proof.* We prove this theorem by construction, i.e., we provide an algorithm to construct an optimal solution to (SLP) with consistent nested assortments when $n = 2$. The key idea of this algorithm is that if there are two customer types such that one is offered the assortment $\{1\}$ and the other is offered $\{2\}$ for a certain amount of time, then there is always a way to offer them nested assortments instead while selling the same number of units and using at most the same time for each customer type.

We first formalize this idea. Suppose there are two customer types $j_1$ and $j_2$, such that assortment $\{1\}$ is offered to $T_1$ type-$j_1$ customers and the assortment $\{2\}$ is offered to $T_2$ type-$j_2$ customers. In other words, $q^{j_1}_1 = u^{j_1}_1 T_1/(1 + u^{j_1}_1)$, $q^{j_2}_2 = u^{j_2}_2 T_2/(1 + u^{j_2}_2)$ and $q^{j_1}_2 = q^{j_2}_1 = 0$ is a feasible solution to the (SLP). We now consider two alternative scenarios:

(A) Assortment $\{1, 2\}$ offered to $t^1_{12}$ type-$j_1$ customers and assortments $\{1, 2\}$ and $\{2\}$ offered to $t^2_{12}$ and $t^2_2$ type-$j_2$ customers, respectively.

(B) Assortment $\{1, 2\}$ offered to $t^2_{12}$ type-$j_2$ customers and assortments $\{1, 2\}$ and $\{1\}$ offered to $t^1_{12}$ and $t^1_1$ type-$j_1$ customers, respectively.

We show that there is a feasible solution in at least one of these scenarios, where total sales of products 1 and 2 is at least $q^{j_1}_1$ and $q^{j_2}_2$, respectively, while the total type-$j_1$ and $j_2$ customers is at most $T_1$ and $T_2$, respectively. We construct LPs corresponding to the two respective scenarios in (A) and (B), and show in Lemma A.3 that at least one of them must be feasible. We defer the proof to Section A.4.1.

$$
\begin{array}{llllllll}
\max & 0 & & & & & & \\
\text{s.t.} & t_1^1 & + & t_{12}^1 & & & \le & T_1, \\
 & & & & & t_{12}^2 & \le & T_2, \\
 & \frac{u_1^{j_1}}{1+u_1^{j_1}} t_1^1 & + & \frac{u_1^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} t_{12}^1 & + & \frac{u_1^{j_2}}{1+u_1^{j_2}+u_2^{j_2}} t_{12}^2 & = & \frac{u_1^{j_1}}{1+u_1^{j_1}} T_1, \\
 & & & \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} t_{12}^1 & + & \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}} t_{12}^2 & = & \frac{u_2^{j_2}}{1+u_2^{j_2}} T_2, \\
 & t_1^1, & & t_{12}^1, & & t_{12}^2 & \ge & 0.
\end{array}
\tag{A}
$$

$$
\begin{array}{llllllll}
\max & 0 & & & & & & \\
\text{s.t.} & & & & & t_{12}^1 & \le & T_1, \\
 & t_2^2 & + & t_{12}^2 & & & \le & T_2, \\
 & & & \frac{u_1^{j_2}}{1+u_1^{j_2}+u_2^{j_2}} t_{12}^2 & + & \frac{u_1^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} t_{12}^1 & = & \frac{u_1^{j_1}}{1+u_1^{j_1}} T_1, \\
 & \frac{u_2^{j_2}}{1+u_2^{j_2}} t_2^2 & + & \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}} t_{12}^2 & + & \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} t_{12}^1 & = & \frac{u_2^{j_2}}{1+u_2^{j_2}} T_2, \\
 & t_2^2, & & t_{12}^2. & & t_{12}^1 & \ge & 0.
\end{array}
\tag{B}
$$

**Lemma A.3.** *For any values of* $(T_1, T_2, u_1^{j_1}, u_2^{j_1}, u_1^{j_2}, u_2^{j_2})$. *at least one of the LPs* (A) *and* (B) *is feasible.*

From Lemma A.3, the algorithm follows simply by applying the same idea at most $m$ times. Each time, there is one less customer type that is offered an assortment of size one. Specifically, the algorithm works by first solving (SLP) and let $t_1^j$, $t_2^j$ and $t_{12}^j$ be the amount of type-$j$ customers shown the assortments $\{1\}, \{2\}, \{1,2\}$, respectively. Then, find two customer types $j_1$ and $j_2$ such that the solution offers $j_1$ the assortment $\{1\}$ and $j_2$ the assortment $\{2\}$. From Lemma A.3, one of (A) and (B) must be feasible. Note that these linear programs results in a new solution where either $j_1$ or $j_2$ no longer has a size one assortment, the total amount of inventory consumed of each product is the same, and the total amount of each customer type chosen does not increase. Update the values of $t_1^j$, $t_2^j$ and $t_{12}^j$ according to the feasible program (either (A) and (B)). Then keep repeating the procedure (at most $m$ times) until the only size one assortment offered is either $\{1\}$ or $\{2\}$. Thus, the final output is feasible, with the same objective value, and has consistent nested assortments. ☐

### A.4.1 Proof of Lemma A.3

*Proof.* Consider the LPs (DA) and (DB), that are the duals of (A) and (B) respectively.

$$
\begin{array}{llllllllll}
\min & T_1S_1 & + & T_2S_1 & + & \frac{u_1^{j_1}}{1+u_1^{j_1}}T_1C_1 & + & \frac{u_2^{j_2}}{1+u_2^{j_2}}T_2C_2 & & \\
\text{s.t.} & S_1 & & & + & \frac{u_1^{j_1}}{1+u_1^{j_1}}C_1 & & & \geq 0, \\
 & S_1 & & & + & \frac{u_1^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}C_1 & + & \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}C_2 & \geq 0 \\
 & & & S_1 & + & \frac{u_1^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}C_1 & + & \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}C_2 & \geq 0 \\
 & S_1, & & S_1 & & & & & \geq 0 \\
 & & & & & C_1, & & C_2 & \text{free.}
\end{array}
\tag{DA}
$$

$$
\begin{array}{llllllllll}
\min & T_1S_1 & + & T_2S_1 & + & \frac{u_1^{j_1}}{1+u_1^{j_1}}T_1C_1 & + & \frac{u_2^{j_2}}{1+u_2^{j_2}}T_2C_2 & & \\
\text{s.t.} & & & S_1 & & & + & \frac{u_2^{j_2}}{1+u_2^{j_2}}C_2 & \geq 0, \\
 & S_1 & & & + & \frac{u_1^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}C_1 & + & \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}C_2 & \geq 0 \\
 & & & S_1 & + & \frac{u_1^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}C_1 & + & \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}C_2 & \geq 0 \\
 & S_1, & & S_1 & & & & & \geq 0 \\
 & & & & & C_1, & & C_2 & \text{free.}
\end{array}
\tag{DB}
$$

We prove that at least one of (DA) and (DB) is feasible and bounded, which is equivalent to the statement we wish to prove. First, note that both LPs are equivalent[1]. Therefore, we focus our

---

[1] The LP (DA) transforms into (DB) by substituting $T_1 = T_2$, $u_2^{j_2} = u_1^{j_1}$, $u_1^{j_2} = u_2^{j_1}$

analysis on (DA) and get the corresponding results for (DB) by symmetry.

Note that (DA) is trivially feasible with objective zero as all zeroes is a feasible solution. Next, we analyze conditions under which (DA) is unbounded towards negative infinity. Observe that it is necessary that $C_2 < 0$ for (DA) to be unbounded, since $S_1 \geq 0$, $S_1 + \frac{u_1^{j_1}}{1+u_1^{j_1}} C_1 \geq 0$, and $T_1(S_1 + \frac{u_1^{j_1}}{1+u_1^{j_1}} C_1)$ is in the objective. Keeping this in mind, we make the following variable substitutions:

$$\begin{aligned}
S_1 + \frac{u_1^{j_1}}{1 + u_1^{j_1}} C_1 &= -\frac{u_2^{j_1}}{1 + u_1^{j_1} + u_2^{j_1}} k_1 C_2 \\
S_1 &= -\frac{u_2^{j_2}}{1 + u_1^{j_2} + u_2^{j_2}} k_2 C_2 \\
C_1 &= x C_2,
\end{aligned}$$

where $k_1, k_2 \geq 0$ and $x$ is free. After these substitutions, the objective of (DA) becomes.

$$C_2 \left( \frac{u_2^{j_2}}{1 + u_2^{j_2}} T_2 - \frac{u_2^{j_1}}{1 + u_1^{j_1} + u_2^{j_1}} T_1 k_1 - \frac{u_2^{j_2}}{1 + u_1^{j_2} + u_2^{j_2}} T_2 k_2 \right).$$

Since $C_2 < 0$, if there are any feasible values for $k_1$ and $k_2$ such that the term on the right becomes positive, then (DA) is unbounded. Therefore, for (DA) to be unbounded, it is necessary that the LP (MA) described below has a positive optimal value.

$$\begin{array}{lll}
\max & \frac{u_2^{j_2}}{1+u_2^{j_2}} T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} T_1 k_1 - \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}} T_2 k_2 & \\
\text{s.t.} & 1 - k_1 & \leq \frac{u_1^{j_1}}{1+u_1^{j_1}} x, \\
& 1 - k_2 & \leq -\frac{u_1^{j_2}}{u_2^{j_2}} x, \\
& -\frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}} k_1 - \frac{u_1^{j_1}}{1+u_1^{j_1}} x & \leq 0 \\
& k_1, k_2 \geq 0, \ x \text{ free.} &
\end{array} \tag{MA}$$

The first and second constraints of (MA) correspond to the second and third constraint of (DA) respectively whereas the third constraint follows from the non-negativity of $S_1$. The following are necessary conditions for any optimal solution of (MA):

$$k_1 = \max\left\{1 - \frac{u_1^{j_1}}{1+u_1^{j_1}}x, -\frac{u_1^{j_1}\left(1+u_1^{j_1}+u_2^{j_1}\right)}{u_2^{j_1}\left(1+u_1^{j_1}\right)}x, 0\right\},\ k_2 = \max\left\{1 + \frac{u_1^{j_2}}{u_2^{j_2}}x, 0\right\}.$$

Alternatively,

$$k_1 = \begin{cases} -\frac{u_1^{j_1}\left(1+u_1^{j_1}+u_2^{j_1}\right)}{u_2^{j_1}\left(1+u_1^{j_1}\right)}x & x \leq -\frac{u_2^{j_1}}{u_1^{j_1}} \\ 1 - \frac{u_1^{j_1}}{1+u_1^{j_1}}x & -\frac{u_2^{j_1}}{u_1^{j_1}} < x \leq \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}} \\ 0 & \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}} < x \end{cases},$$

$$k_2 = \begin{cases} 0 & x \leq -\frac{u_2^{j_2}}{u_1^{j_2}} \\ 1 + \frac{u_1^{j_2}}{u_2^{j_2}}x & -\frac{u_2^{j_2}}{u_1^{j_2}} < x \end{cases}.$$

Note that, since $k_2$ increases with $x$ and $k_1 = 0$ when $x \geq \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$, then at optimality of (MA) we have $x \leq \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$. Similarly, $x \geq -\frac{u_2^{j_2}}{u_1^{j_2}}$ at optimality of (MA). We next look at the objective value of (MA) at the extreme points which lie within this range of values for $x$.

**Case 1:** $-\frac{u_2^{j_1}}{u_1^{j_1}} \leq -\frac{u_2^{j_2}}{u_1^{j_2}}$. Equivalently, this is the case where $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$. The relevant extreme points for this case are

$$(k_1, k_2) = \left(0, \frac{\left(u_1^{j_1}u_2^{j_2} + u_1^{j_2}(1+u_1^{j_1})\right)}{u_1^{j_1}u_2^{j_2}}\right), \left(\frac{\left(u_1^{j_1}u_2^{j_2} + u_1^{j_2}(1+u_1^{j_1})\right)}{\left(1+u_1^{j_1}\right)u_1^{j_2}}, 0\right),$$

which correspond to solutions at $x = \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$ and $x = -\frac{u_2^{j_2}}{u_1^{j_2}}$ respectively when $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$. The objective value of (MA) at $x = \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$ is,

$$\begin{aligned}
&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}T_1k_1 - \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}T_2k_2 \\
=&T_2\left(\frac{u_2^{j_2}}{1+u_2^{j_2}} - \frac{\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)}{u_1^{j_1}\left(1+u_1^{j_2}+u_2^{j_2}\right)}\right) \\
=&T_2\left(\frac{u_1^{j_1}u_2^{j_2}u_1^{j_2} - u_1^{j_2}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_1}\right)}{u_1^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}\right) < 0.
\end{aligned} \tag{A.34}$$

Similarly, the objective value at $x = -\frac{u_2^{j_2}}{u_1^{j_2}}$ is,

$$\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)u_2^{j_1}}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)}T_1.$$

Thus, if $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$, (DA) is unbounded if and only if,

$$\begin{aligned}
&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)u_2^{j_1}}{\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)}T_1 > 0 \\
\Rightarrow&\frac{T_2}{T_1} > \frac{\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)u_2^{j_1}\left(1+u_2^{j_2}\right)}{u_1^{j_2}u_2^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)}.
\end{aligned}$$

By symmetry, when $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$, (DB) is unbounded if and only if,

$$\frac{T_1}{T_2} > \frac{\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)u_1^{j_2}\left(1+u_1^{j_1}\right)}{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}.$$

Combining these two conditions, there exist $T_1$, $T_2$ such that both (DA) and (DB) are unbounded when $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$ if and only if,

$$\begin{aligned}
&\frac{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)} > \frac{u_2^{j_1}\left(1+u_2^{j_2}\right)\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)}{u_1^{j_2}u_2^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)}\\
\Rightarrow & u_1^{j_1}u_2^{j_2}\left(1+u_1^{j_1}+u_2^{j_1}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)\\
&\qquad > \left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)\\
\Rightarrow & u_1^{j_1}u_2^{j_2}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_1}\right)+u_1^{j_1}u_2^{j_2}u_2^{j_1}u_1^{j_2}\\
&\qquad > (u_1^{j_1}u_2^{j_2})^2+u_2^{j_1}u_1^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_2^{j_2}\right)\\
\Rightarrow & (u_1^{j_1}u_2^{j_2}-u_2^{j_1}u_1^{j_2})(1+u_1^{j_1}+u_2^{j_2}) > 0\\
\Rightarrow & u_1^{j_1}u_2^{j_2} > u_2^{j_1}u_1^{j_2}.
\end{aligned}$$

This is not possible since $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$. Therefore, both (DA) and (DB) cannot be unbounded when $u_2^{j_1}u_1^{j_2} \geq u_1^{j_1}u_2^{j_2}$.

**Case 2:** $-\frac{u_2^{j_1}}{u_1^{j_1}} > -\frac{u_2^{j_2}}{u_1^{j_2}}$. Equivalently, this is the case where $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$. The relevant extreme points for this case are

$$\begin{aligned}
(k_1,k_2) = &\left(0, \frac{\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)}{u_1^{j_1}u_2^{j_2}}\right),\\
&\left(\frac{\left(1+u_1^{j_1}+u_2^{j_1}\right)}{\left(1+u_1^{j_1}\right)}, \frac{u_1^{j_1}u_2^{j_2}-u_1^{j_2}u_2^{j_1}}{u_1^{j_1}u_2^{j_2}}\right),\\
&\left(\frac{u_2^{j_2}u_1^{j_1}\left(1+u_1^{j_1}+u_2^{j_1}\right)}{u_2^{j_1}u_1^{j_2}\left(1+u_1^{j_1}\right)}, 0\right),
\end{aligned}$$

which correspond to solutions at $x = \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$, $x = -\frac{u_2^{j_1}}{u_1^{j_1}}$ and $x = -\frac{u_2^{j_2}}{u_1^{j_2}}$ respectively when $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$. We already saw in (A.34) that at $x = \frac{\left(1+u_1^{j_1}\right)}{u_1^{j_1}}$ the objective value is always negative. The objective value at $x = -\frac{u_2^{j_1}}{u_1^{j_1}}$ is,

$$
\begin{aligned}
&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}T_1 k_1 - \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}T_2 k_2 \\
=&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}}T_1 - \frac{u_1^{j_1}u_2^{j_2} - u_1^{j_2}u_2^{j_1}}{u_1^{j_1}\left(1+u_1^{j_2}+u_2^{j_2}\right)}T_2 \\
=&\frac{u_1^{j_2}\left(u_1^{j_1}u_2^{j_2} + u_2^{j_1}(1+u_2^{j_2})\right)}{u_1^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}}T_1. \qquad \text{(A.35)}
\end{aligned}
$$

The objective value at $x = -\frac{u_2^{j_2}}{u_1^{j_2}}$ is,

$$
\begin{aligned}
&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{u_2^{j_1}}{1+u_1^{j_1}+u_2^{j_1}}T_1 k_1 - \frac{u_2^{j_2}}{1+u_1^{j_2}+u_2^{j_2}}T_2 k_2 \qquad \text{(A.36)} \\
=&\frac{u_2^{j_2}}{1+u_2^{j_2}}T_2 - \frac{u_2^{j_2}u_1^{j_1}}{u_1^{j_2}\left(1+u_1^{j_1}\right)}T_1
\end{aligned}
$$

From (A.35) and (A.36), (DA) is unbounded when $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$ if and only if

$$
\begin{aligned}
\frac{T_2}{T_1} &> \min\left\{ \frac{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2} + u_2^{j_1}(1+u_2^{j_2})\right)}, \frac{u_1^{j_1}\left(1+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)} \right\} \\
&= \frac{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2} + u_2^{j_1}(1+u_2^{j_2})\right)},
\end{aligned}
$$

where the last equality holds since,

$$\begin{aligned}
&\frac{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)}\\
=&\frac{u_1^{j_1}\left(1+u_2^{j_2}\right)(u_2^{j_1}u_1^{j_2}+u_2^{j_1}\left(1+u_2^{j_2}\right))}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)}\\
=&\frac{u_1^{j_1}\left(1+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)}-\frac{u_1^{j_1}\left(1+u_2^{j_2}\right)(u_1^{j_1}u_2^{j_2}-u_1^{j_2}u_2^{j_1})}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)}\\
<&\frac{u_1^{j_1}\left(1+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)}.
\end{aligned}$$

By symmetry, (DB) is unbounded when $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$ if and only if,

$$\frac{T_1}{T_2} > \frac{u_1^{j_2}u_2^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)}{u_2^{j_1}\left(1+u_2^{j_2}\right)\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)}.$$

Combining these two conditions, there exist $T_1$, $T_2$ such that both (DA) and (DB) are unbounded when $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$ if and only if,

$$\begin{aligned}
&\frac{u_2^{j_1}\left(1+u_2^{j_2}\right)\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)}{u_1^{j_2}u_2^{j_2}\left(1+u_1^{j_1}\right)\left(1+u_1^{j_1}+u_2^{j_1}\right)} > \frac{u_1^{j_1}u_2^{j_1}\left(1+u_2^{j_2}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)}{u_1^{j_2}\left(1+u_1^{j_1}\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)}\\
\Rightarrow&\left(u_1^{j_1}u_2^{j_2}+u_1^{j_2}(1+u_1^{j_1})\right)\left(u_1^{j_1}u_2^{j_2}+u_2^{j_1}(1+u_2^{j_2})\right)\\
&\qquad\qquad > u_1^{j_1}u_2^{j_2}\left(1+u_1^{j_1}+u_2^{j_1}\right)\left(1+u_1^{j_2}+u_2^{j_2}\right)\\
\Rightarrow&(u_1^{j_1}u_2^{j_2}-u_2^{j_1}u_1^{j_2})(1+u_1^{j_1}+u_2^{j_2}) < 0\\
\Rightarrow&u_1^{j_1}u_2^{j_2} < u_2^{j_1}u_1^{j_2}.
\end{aligned}$$

This is not possible since $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$. Therefore, both (DA) and (DB) cannot be unbounded when $u_2^{j_1}u_1^{j_2} < u_1^{j_1}u_2^{j_2}$.

Combining both cases, we conclude that it is not possible for both (DA) and (DB) to be unbounded. Therefore, at least one of (DA) and (DB) must be feasible bounded. This concludes our proof. □

### A.5 Product Retirement with Independent Demand

In the independent demand model, the preference of the customer does not depend on the available products. Each customer has a desired product, which is sampled from a known distribution. If the desired product is available, the customer purchases that product, else they leave without making a purchase. We let $v_i^j$ denote the probability that product $i$ is the desired product of a type-$j$ customer. Therefore, for any $S \subseteq \mathcal{N}$, any $i \in S$, and any $j \in \mathcal{L}$, we have that

$$\pi_i^j(S) = v_i^j.$$

As a result, the probability of no-purchase when customer type $j$ is offered assortment $S$ becomes

$$\pi_0^j(S) = v_0^j + \sum_{i \in \mathcal{N} \setminus S} v_i^j.$$

Since the desired product is sampled independent of the available products, there is no benefit to be gained from retiring products under this model. Under the independent demand model, it is optimal to never retire any products. Therefore, having one type results in a trivial setting when customers choose according to the independent demand model.

When there are multiple customer types, customer selection decisions become critical. When there are multiple customer types, we show that a dynamic greedy policy can achieve a $1 - 1/e$ approximation to the expected revenue of the optimal dynamic policy. In each time period, the policy selects the customer type with the highest revenue in expectation over the products available at the time. The policy, referred to as A3 is provided in detail in Algorithm 3.

**Theorem A.4.** *Policy* A3 *(described in Algorithm 3) guarantees a* $1 - 1/e$ *approximation to the*

**Algorithm 3** A3: Greedy policy for Independent Choice Model

```
Initialize $Y_j \leftarrow 0 \quad \forall j \in \mathcal{L}, Z_i \leftarrow 0 \quad \forall i \in \mathcal{N}$
for $t \leftarrow 1, \ldots, T$ do
    $S_t \leftarrow \{i \in \mathcal{N} \mid Z_i \leq c_i - 1\}, \mathcal{L}_t \leftarrow \{j \in \mathcal{L} \mid Y_j \leq b_j - 1\}$
    $J_t \leftarrow \arg\max_{j \in \mathcal{L}} \sum_{i \in S_t} r_i v_i^j$
    $X_t \leftarrow$ Product purchased by type-$J_t$ customer when offered $S_t$          // Independent demand model with probability $\boldsymbol{\pi}^{J_t}(S_t)$
    $Z_{X_t} \leftarrow Z_{X_t} + 1, Y_{J_t} \leftarrow Y_{J_t} + 1$
end for
```

*expected revenue of the optimal dynamic policy. Moreover, this analysis is tight.*

The proof of Theorem A.4 uses the result by Asadpour and Nazerzadeh [73] for maximizing stochastic monotone submodular functions over a uniform matroid. The formal proof for the performance guarantee is provided in Appendix A.5.1 while an instance for which the analysis is tight is provided in Appendix A.5.2.

### A.5.1 Proof of Theorem A.4

*Proof.* For convenience, we assume without loss of generality that $b_j = 1$ for all $j \in \mathcal{L}$. As per Algorithm 3, $S_t$ is the set of products with inventory remaining at time $t$ and $J_t$ is the customer chosen at time $t$. Let $\widetilde{X}_j$ be the desired product of the type-$j$ customer. We know that, $\mathbb{P}[\widetilde{X}_j = i] = v_i^j$. Consequently, $X_t$ can be expressed as

$$\widetilde{X}_{J_t} \mathbb{I}[\widetilde{X}_{J_t} \in S_t] \tag{A.37}$$

**Revenue as a set function of selected customers.** Let $\mathcal{J}_t$ be the set of customers selected up to time $t$ in Algorithm 3 and let $Z_{i,t}$ be state of variable $Z_i$ in Algorithm 3 at time $t$. Then the total revenue earned up to time $t$ can be written as,

$$\begin{aligned}\sum_{i=1}^{n} r_i Z_{i,t} &= \sum_{s=1}^{t}\sum_{i=1}^{n} r_i \mathbb{I}[X_s = i] = \sum_{s=1}^{t}\sum_{i=1}^{n} r_i \mathbb{I}[\widetilde{X}_{J_s} = i]\mathbb{I}[i \in S_s] \\ &= \sum_{s=1}^{t}\sum_{i=1}^{n} r_i \mathbb{I}[\widetilde{X}_{J_s} = i]\mathbb{I}[Z_{i,s-1} \le c_i] \\ &= \sum_{i=1}^{n} r_i \min\left\{\sum_{s=1}^{t} \mathbb{I}[\widetilde{X}_{J_s} = i], c_i\right\}, \qquad \text{(A.38)}\end{aligned}$$

where the second equality follows from (A.37), the third equality follows from the definition of $S_t$ in Algorithm 3. Observe that the right-hand side of (A.38) is order-independent. It only depends on the selected customer and not on the order in which they were selected. We abuse notation to let $R(\mathcal{J}_t)$ denote the right-hand side of (A.38). More generally, let

$$R(\mathcal{J}) := \sum_{i=1}^{n} r_i \min\left\{\sum_{j \in \mathcal{J}} \mathbb{I}[\widetilde{X}_j = i], c_i\right\}$$

be the set function denoting the expected revenue earned from selecting the set of customers $J$.

**Revenue function is monotone submodular.** We show that $R(\mathcal{J})$ is monotone submodular. Note that if $\mathcal{J} \subseteq \mathcal{K} \subseteq \mathcal{L}$,

$$\sum_{j \in \mathcal{J}} \mathbb{I}[\widetilde{X}_j = i] \le \sum_{j \in \mathcal{K}} \mathbb{I}[\widetilde{X}_j = i] \qquad \forall i \in \mathcal{N}. \qquad \text{(A.39)}$$

We first show monotonicity. Note that, if $\mathcal{J} \subseteq \mathcal{K} \subseteq \mathcal{L}$, then

$$\begin{aligned} R(\mathcal{J}) &= \sum_{i=1}^{n} r_i \min\left\{\sum_{j \in \mathcal{J}} \mathbb{I}[\widetilde{X}_j = i], c_i\right\} \\ &\le \sum_{i=1}^{n} r_i \min\left\{\sum_{j \in \mathcal{K}} \mathbb{I}[\widetilde{X}_j = i], c_i\right\},\end{aligned}$$

where the inequality follows from (A.39). Next we show submodularity. Observe that, for any $k \in \mathcal{L} \setminus \mathcal{K}$,

$$\begin{aligned} R(\mathcal{J} \cup \{k\}) - R(\mathcal{J}) &= \sum_{i=1}^{n} r_i \mathbb{I}[\widetilde{X}_k = i] \mathbb{I}\left[\sum_{j \in \mathcal{J}} \widetilde{X}_j \leq c_i - 1\right] \\ &\geq \sum_{i=1}^{n} r_i \mathbb{I}[\widetilde{X}_k = i] \mathbb{I}\left[\sum_{j \in \mathcal{K}} \widetilde{X}_j \leq c_i - 1\right] \\ &= R(\mathcal{K} \cup \{k\}) - R(\mathcal{K}), \end{aligned}$$

where the inequality follows from (A.39). Therefore, $R(\mathcal{J})$ is monotone submodular.

**Finishing up.** Asadpour and Nazerzadeh [73] showed that for maximizing the expected value of a stochastic monotone submodular function over a uniform matroid, a dynamic greedy policy that selects the element maximizing the expected marginal value in each time period guarantees an approximation ratio of $1 - 1/e$ with respect to the optimal adaptive policy. Suppose $x_{\mathcal{J}}$ is the realization of the desired products of all customers in $\mathcal{J}$. Then the marginal expected revenue in time period $t$ for choosing customer type $k$ is expressed as,

$$\begin{aligned} &\mathbb{E}[R(\mathcal{J}_{t-1} \cup \{k\}) - R(\mathcal{J}_{t-1}) \mid x_{\mathcal{J}_{t-1}}] \\ &= \sum_{i=1}^{n} r_i \mathbb{E}\left[\mathbb{I}[\widetilde{X}_k = i] \mathbb{I}\left[\sum_{j \in \mathcal{J}_{t-1}} \widetilde{X}_j \leq c_i - 1\right] \middle| x_{\mathcal{J}_{t-1}}\right] \\ &= \sum_{i=1}^{n} r_i \mathbb{P}[\widetilde{X}_k = i] \mathbb{I}\left[Z_{i,t-1} \leq c_i - 1\right] \\ &= \sum_{i=1}^{n} r_i \mathbb{P}[\widetilde{X}_k = i] \mathbb{I}[i \in S_t] \\ &= \sum_{i \in S_t} r_i v_i^k. \end{aligned}$$

Here, the second equality follows from the definition of $Z_{i,t-1}$ and the fact that $\widetilde{X}_k$ and $\{\widetilde{X}_j \mid j \in \mathcal{J}_{t-1}\}$ are independent, the third equality follows from the definition of $S_t$ in Algorithm 3 and the fourth equality follows from the definition of $\widetilde{X}_j$. Therefore, $J_t$ is the customer that maximizes the expected marginal contribution to the revenue. This implies that Algorithm 3 guarantees a $1 - 1/e$ approximation to the expected revenue of the optimal adaptive policy. $\square$

### A.5.2 Tight example for Theorem A.4

In this section we provide an instance for which the $1 - 1/e$ guarantee in Theorem A.4 is tight.

**Instance parameters.** For this instance, the number of customer types is the same as the number of products, i.e., $m = n$. The set of potential customers has $T$ customers of each type, i.e., $b_j = T$ for $j = 1, \ldots, n$ where recall $T$ is the time horizon. We now define the choice model parameters for each customer type. For type-1 customers, $v_1^1 = \frac{(n-1)^{n-1}}{n^{n-1}}$, $v_i^1 = \frac{(n-1)^{n-i}}{n^{n-i+1}}$ for $i = 2, \ldots, n$ and $v_0^1 = 0$. Note that the parameters are valid, i.e., they sum up to 1.

$$\sum_{i=1}^{n} v_i^1 = \frac{(n-1)^{n-1}}{n^{n-1}} + \sum_{i=2}^{n} \frac{(n-1)^{n-i}}{n^{n-i+1}} \tag{A.40}$$

$$= \frac{(n-1)^{n-1}}{n^{n-1}} \left(1 + \frac{1}{n-1} \sum_{i=2}^{n} \left(\frac{n}{n-1}\right)^{i-2}\right) \tag{A.41}$$

$$= \frac{(n-1)^{n-1}}{n^{n-1}} \left(1 + \left(\frac{1}{n-1}\right)\left(\frac{\left(\frac{n}{n-1}\right)^{n-1} - 1}{\frac{1}{n-1}}\right)\right) \tag{A.42}$$

$$= 1. \tag{A.43}$$

Also, note that

$$v_2^1 < v_3^1 < \ldots < v_n^1. \tag{A.44}$$

For $j = 2, \ldots, n$, the choice model parameters are $v_j^j = n v_j^1 = (1 - 1/n)^{n-j}$, $v_0^j = 1 - v_j^j$ and $v_i^j = 0$ for the remaining products. The initial inventory for every product $i$ is given by $c_i = T v_i^1$ for all $i \in \mathcal{N}$. The price of the first product is $r_1 = 1$ while that of the remaining products is $r_i = 1 + 1/T$

for $i = 2, \dots, n$. Note that for $i = 2, \dots, n$, $\sum_{k=1}^{i} v_k^1 = nv_i^1 = v_i^i$.

$$\begin{aligned}
\sum_{k=1}^{i} v_k &= \frac{(n-1)^{n-1}}{n^{n-1}} + \sum_{k=2}^{i} \frac{(n-1)^{n-k}}{n^{n-k+1}} \\
&= \frac{(n-1)^{n-1}}{n^{n-1}} \left( 1 + \frac{1}{n-1} \sum_{k=2}^{i} \left( \frac{n}{n-1} \right)^{k-2} \right) \\
&= \frac{(n-1)^{n-1}}{n^{n-1}} \left( 1 + \left( \frac{1}{n-1} \right) \left( \frac{\left(\frac{n}{n-1}\right)^{i-1} - 1}{\frac{1}{n-1}} \right) \right) \\
&= \left( \frac{n-1}{n} \right)^{n-i} = nv_i^1 = v_i^i,
\end{aligned} \tag{A.45}$$

where the last equality follows from the definition of $v_i^i$. The above implies that for $i = 2, \dots, n$,

$$\begin{aligned}
R_{\pi^1}([i]) &= \sum_{k=1}^{i} r_k v_k^1 = v_1^1 + \left(1 + \frac{1}{T}\right) \sum_{k=2}^{i} v_k^1 < \left(1 + \frac{1}{T}\right) \sum_{k=1}^{i} v_k^1 \\
&= \left(1 + \frac{1}{T}\right) nv_i^1 = r_i v_i^i = R_{\pi^i}([i]),
\end{aligned} \tag{A.46}$$

where the first and last equality follow from (1.1), the second equality follows from the definition of $r_i$, the third equality follows from (A.45) and the fourth equality follows from the definition of $r_i$ and $v_i^i$ for $i > 1$.

**Revenue from selecting type-1 customers.** Consider the customer selection policy Type-1 that always selects type-1 customers. Assume that $T$ is large, then expected revenue for this policy is,

$$\begin{aligned}
\textsc{Rev}(\textsc{Type-1}) = \sum_{i=1}^{n} r_i \mathbb{E}[Z_i] \geq \sum_{i=1}^{n} \mathbb{E}[Z_i] = \sum_{i=1}^{n} \mathbb{E}[\min\{c_i, \hat{Z}_i\}] &= \sum_{i=1}^{n} \mathbb{E}[\min\{Tv_i^1, \hat{Z}_i\}] \\
&\approx \sum_{i=1}^{n} Tv_i^1 = T,
\end{aligned} \tag{A.47}$$

where $\hat{Z}_i$ denotes the number of customers who desire product $i$. The first equality follows from (1.2), the inequality follows from the fact that $r_i \geq 1$ for all $i \in \mathcal{N}$, the second equality follows from the definition of $\hat{Z}_i$, the third equality follows from the definition of $c_i$, for the approximation we use the fact that since $\mathbb{E}[\hat{Z}_i] = Tv_i^1$, for large $T$, $\mathbb{E}[\min\{Tv_i^1, \hat{Z}_i\}] \approx Tv_i^1$ [2], and the fourth equality follows from (A.40).

**Revenue from A3.** Since policy A3 selects customers greedily, from (A.44) and (A.45) it follows that the policy will first select type-$n$ customers. Since these customers only purchase product $n$, these customers will continue to be selected until product $n$ stocks out. Following this, type-$(n-1)$ customers will be selected and so on, until finally only product 1 remains, at which point type-1 customers will be selected. The expected number of type-$i$ customers that must be visited to stock out product $i$ for $i = 2, \ldots, n$ is,

$$\frac{c_i}{\pi_i^i([i])} = \frac{Tv_i^1}{v_i^i} = \frac{T}{n}. \tag{A.48}$$

For large enough $T$, we can assume that products $2, \ldots, n$ stock out before the time horizon and $T/n$ customers of each type are visited. The expected revenue earned under these assumptions is bounded above as follows,

$$\begin{aligned} \textsc{Rev}(\text{A3}) &\approx \frac{T}{n} r_1 v_1^1 + T \sum_{i=2}^{n} r_i v_i^1 = \frac{T}{n} v_1^1 + (T+1) \sum_{i=2}^{n} v_i^1 \\ &\leq \frac{T}{n} v_1^1 + T(1 - v_1^1) + 1 = T\left(1 - \left(1 - \frac{1}{n}\right)^n\right) + 1, \end{aligned} \tag{A.49}$$

where the approximation follows from the fact that for large $T$ products $2, \ldots, n$ are sold out and $T/n$ type-1 customers are visited, the first equality follows from the definition of $r_i$, the inequality follows from (A.40) and the fact that $v_1^1 \geq 0$, and the second equality follows from the definition of $v_1^1$.

From (A.47) and (A.49) it follows that for large $T$ and $n$, A3 achieves at most $1 - 1/e$ fraction of the optimal revenue for this instance.

---

[2] This can be shown rigorously using Lemma A.2.

# Appendix B: Appendix to Chapter 2

## B.1 Proofs

### B.1.1 Proof of Lemma 2.5

*Proof.* An MHR valuation distribution implies that for any feasible prices $p, p'$ such that $p < p'$, it holds that $\frac{f(p)}{1-F(p)} \leq \frac{f(p')}{1-F(p')}$. Recall that $1 - F(p) = \lambda(p)/\Lambda$. It follows that,

$$p(\lambda) = F^{-1}\left(1 - \frac{\lambda}{\Lambda}\right) \Rightarrow p'(\lambda) = \frac{-1}{\Lambda f(p(\lambda))} \Rightarrow \frac{f(p)}{1 - F(p)} = \frac{-1}{\lambda(p)p'(\lambda(p))}.$$

Since $\lambda(p)$ is monotonically decreasing in $p$, the MHR condition can equivalently be expressed as: for any rates $\lambda, \lambda'$ such that $\lambda > \lambda'$, it holds that

$$-\lambda p'(\lambda) \geq -\lambda' p'(\lambda').$$

Under the lost sales model, the inventory level is always non-negative, thus Lemma 2.1(iv) implies that for any positive integers $i, j \leq S^*$ such that $i > j$, $\lambda_i \geq \lambda_j$. Therefore, from Lemma 2.1(i) it follows that

$$-\lambda_i p'(\lambda_i) \geq -\lambda_j p'(\lambda_j) \Rightarrow \frac{hi + \gamma^*}{\lambda_i^*} \geq \frac{hj + \gamma^*}{\lambda_j^*}.$$

The above relationship establishes an upper bound on the rate at which the optimal arrival rate increases with the inventory level, under the optimal pricing policy for any MHR valuation distribution. Since $\pi_i(S^*, \boldsymbol{\lambda}^*) \propto 1/\lambda_i^*$ from (2.3), it follows that

$$\frac{hi + \gamma^*}{\lambda_i^*} \geq \frac{hj + \gamma^*}{\lambda_j^*} \Rightarrow \frac{hi + \gamma^*}{hj + \gamma^*} \geq \frac{\pi_j(S^*, \boldsymbol{\lambda}^*)}{\pi_i(S^*, \boldsymbol{\lambda}^*)} \Rightarrow \frac{i}{j} \geq \frac{\pi_j(S^*, \boldsymbol{\lambda}^*)}{\pi_i(S^*, \boldsymbol{\lambda}^*)},$$

where the last inequality follows from the fact that $\gamma^* > 0$. Using this fact, the minimum holding cost rate over all MHR valuation distributions for a fixed order-up-to level $S$ and optimal arrival rates can be characterized as follows:

$$\begin{aligned} \min_{\pi_i} \quad & \textstyle\sum_{i=1}^{S} hi\pi_i \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^{S} \pi_i = 1 \\ & i\pi_i \geq j\pi_j \quad 1 \leq j < i \leq S \\ & \pi_i \geq 0 \qquad i = 1, \ldots, S. \end{aligned} \tag{B.1}$$

Intuitively, the optimal solution for the above linear program is the point for which the decrease in the steady-state probabilities is maximized, i.e., $i\pi_i = j\pi_j$ for all $1 \leq j < i \leq S$. We formally prove this using induction over $S$.

**Base Case:** For $S = 1$, the hypothesis is trivially true.

**Induction Step:** Suppose the hypothesis is true for order-up-to level $S-1$. Let $\boldsymbol{\pi}^*$ be an optimal solution for order-up-to level $S$. Observe that it follows by the induction hypothesis that $i\pi_i^* = j\pi_j^*$ for $1 \leq j < i \leq S - 1$. Moreover, if $z_S$ and $z_{S-1}$ are the optimal values of (B.1) for $S$ and $S - 1$ respectively, then,

$$z_S = (1 - \pi_S^*)z_{S-1} + \pi_S^* hS.$$

Since, $z_{S-1} \leq h(S - 1) < hS$, $\pi_S^*$ must be as small as possible for $z_S$ to be the optimal value. Therefore, $\pi_S^* = i\pi_i^*/S$ for all $i = 1, \ldots, S - 1$. Thus, the hypothesis holds for $S$.

It follows that the optimal solution and optimal value to (B.1) are

$$\pi_i^* = \frac{\frac{1}{i}}{\sum_{j=1}^{S} \frac{1}{j}} \text{ and } \sum_{i=1}^{S} hi\pi_i = \frac{hS}{\sum_{i=1}^{S} \frac{1}{i}} \tag{B.2}$$

respectively. Therefore, the holding cost rate for any MHR distribution with optimal order-up-to level $S^*$ is be bounded below as follows.

$$\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) = \sum_{i=1}^{S^*} hi\pi_i(S^*, \boldsymbol{\lambda}^*) \geq \frac{hS^*}{\sum_{i=1}^{S^*} \frac{1}{i}} \geq \frac{hS^*}{1 + \int_1^{S^*} \frac{1}{x} dx} = \frac{hS^*}{1 + \ln(S^*)}.$$

□

### B.1.2 Proof of Lemma 2.6

*Proof.* To see the upper bound on the ordering cost ratio observe that

$$\frac{\textsc{Order}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)} = \frac{S^*}{S_{static}} = \frac{S^*}{\left\lceil \frac{S^*}{\sqrt{1+\ln(S^*)}} \right\rceil} \leq \frac{S^*}{\frac{S^*}{\sqrt{1+\ln(S^*)}}} = \sqrt{1 + \ln(S^*)},$$

which follows from substituting (2.11) in (2.8). For the holding cost ratio, note that,

$$\begin{aligned}\frac{\textsc{Hold}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} &\leq \frac{(1 + \ln(S^*))(S_{static} + 1)}{2S^*} = \frac{(1 + \ln(S^*))\left(\left\lceil \frac{S^*}{\sqrt{1+\ln(S^*)}} \right\rceil + 1\right)}{2S^*} \\ &\leq \frac{(1 + \ln(S^*))\left(\frac{S^*}{\sqrt{1+\ln(S^*)}} + 2\right)}{2S^*} = \frac{\sqrt{1 + \ln(S^*)}}{2} + \frac{1 + \ln(S^*)}{S^*}\end{aligned}$$

which follows from substituting (2.11) in (2.10) and using the fact that $\lceil x \rceil \leq x + 1$. For $S^* \geq 3$.

$$\frac{\sqrt{1 + \ln(S^*)}}{2} + \frac{1 + \ln(S^*)}{S^*} \leq \sqrt{1 + \ln(S^*)} \Rightarrow \frac{\textsc{Hold}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} \leq \sqrt{1 + \ln(S^*)}.$$

For $S^* = 1, 2$, the above bound is inadmissible. When $S^* = 1$, $S_{static} = 1$ and the optimal pricing policy is static, thus the holding cost ratio is 1. When $S^* = 2$, it follows from (B.2) that the lower bound on the optimal holding cost rate is $4h/3$. On the other hand, when $S^* = 2$, $S_{static} = 2$, thus the holding cost rate for our static policy is $h(S_{static} + 1)/2 = 3h/2$. Therefore, the holding cost ratio for $S^* = 2$ is bounded above by $9/8$ which is less than $\sqrt{1 + \ln(2)} \approx 1.3$. Therefore, ordering and holding cost ratios are both bounded above by $\sqrt{1 + \ln(S^*)}$ for $S_{static}$ as defined in (2.11) for all $S^* \geq 1$. □

### B.1.3 Proof of Theorem 2.4

*Proof.* We first formulate the optimal rates under linear demand. Suppose the demand model is given by $\lambda = a(1 - bp)$ and the optimal profit rate is $\gamma^*$. Then from Lemma 2.1(i) it follows that,

$$p'(\lambda_i^*) = -\frac{1}{ab} = -\frac{hi + \gamma^*}{(\lambda_i^*)^2} \Rightarrow \lambda_i^* = \sqrt{abhi + ab\gamma^*}. \tag{B.3}$$

Having this exact relation of how the rates vary with inventory level for the optimal pricing policy allows us to lower bound the holding cost rate for the optimal dynamic policy by

$$\begin{aligned}\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) &= \frac{\sum_{i=1}^{S^*} \frac{hi}{\lambda_i^*}}{\sum_{i=1}^{S^*} \frac{1}{\lambda_i^*}} = \frac{\sum_{i=1}^{S^*} \frac{hi}{\sqrt{abhi+ab\gamma^*}}}{\sum_{i=1}^{S^*} \frac{1}{\sqrt{abhi+ab\gamma^*}}} \geq \frac{\sum_{i=1}^{S^*} \frac{hi}{\sqrt{abhi}}}{\sum_{i=1}^{S^*} \frac{1}{\sqrt{abhi}}} = \frac{h\sum_{i=1}^{S^*} \sqrt{i}}{\sum_{i=1}^{S^*} \frac{1}{\sqrt{i}}} \\ &\geq \frac{h\int_0^{S^*} \sqrt{x}dx}{1 + \int_1^{S^*} \frac{1}{\sqrt{x}}dx} = \frac{\frac{2h}{3}(S^*)^{3/2}}{1 + 2(\sqrt{S^*} - 1)} = \frac{2h(S^*)^{3/2}}{6\sqrt{S^*} - 3},\end{aligned} \tag{B.4}$$

where the first equality follows from (2.4), the second equality follows from (B.3), and the first inequality uses the fact that the expression is monotonically decreasing in $\gamma^*$ and $\gamma^* \geq 0$, and the second inequality follows from interpreting the summations in the numerator and denominator as right and left Riemann sums respectively. We can now prove both statements of the theorem.

**(i)** From (2.9) and (B.4) it follows that,

$$\frac{\textsc{Hold}(S^*, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} \leq \frac{3(2\sqrt{S^*} - 1)(S^* + 1)}{4(S^*)^{3/2}} = \frac{3S^*}{2S^*} + \frac{3(2\sqrt{S^*} - 1 - S^*)}{4(S^*)^{3/2}} = \frac{3}{2} - \frac{3(\sqrt{S^*} - 1)^2}{4(S^*)^{3/2}} \leq 1.5. \tag{B.5}$$

**(ii)** Our approach to this proof is similar to the proof of the total cost ratio bound for Theorem 2.3. Let $S_{static}$ be

$$S_{static} = \left\lceil \frac{S^*}{\sqrt{1.5}} \right\rceil. \tag{B.6}$$

The ordering cost ratio is bounded above by $\sqrt{1.5}$ since,

$$\frac{\textsc{Order}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Order}(S^*, \boldsymbol{\lambda}^*)} = \frac{S^*}{S_{static}} = \frac{S^*}{\left\lceil \frac{S^*}{\sqrt{1.5}} \right\rceil} \leq \frac{S^*}{\frac{S^*}{\sqrt{1.5}}} = \sqrt{1.5},$$

which follows from substituting (B.6) in (2.8). For the holding cost ratio note that,

$$\begin{aligned}\frac{\textsc{Hold}(S_{static}, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(S^*, \boldsymbol{\lambda}^*)} &\leq \frac{3(2\sqrt{S^*}-1)(S_{static}+1)}{4(S^*)^{3/2}} = \frac{3(2\sqrt{S^*}-1)\left(\left\lceil \frac{S^*}{\sqrt{1.5}} \right\rceil + 1\right)}{4(S^*)^{3/2}} \\ &\leq \frac{3(2\sqrt{S^*}-1)\left(\frac{S^*}{\sqrt{1.5}} + 2\right)}{4(S^*)^{3/2}} \leq \sqrt{1.5} + \frac{3}{S^*} - \sqrt{\frac{3}{8S^*}}.\end{aligned}$$

The above bound is bounded above by $\sqrt{1.5}$ for $S^* \geq 24$. When $S^* \leq 23$, we can numerically check that the upper bound of $\sqrt{1.5}$ holds directly through computations by using the following intermediate bound from (B.4), $\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) \geq \frac{h\sum_{i=1}^{S^*}\sqrt{i}}{\sum_{i=1}^{S^*}\frac{1}{\sqrt{i}}}$. Therefore, all cost ratios are bounded above by $\sqrt{1.5}$ when $S_{static}$ is constructed according to (B.6). □

### B.1.4 Proof of Lemma 2.11

*Proof.* This lemma proves several properties for instance $k$ as defined in (2.18) required to show tightness of Theorem 2.4.

**(i)** To see that optimal order-up-to level and profit for instance $k$ are $k$ and 0, respectively, we check the conditions in Lemma 2.1. Note that $g_{i,0}(\lambda_i^*)$ for $i = 1, \ldots, k$ is

$$\begin{aligned}g_{i,0}(\lambda_i^*) &= p(\lambda_i^*) - \frac{h_k i}{\lambda_i^*} = \frac{1}{b}\left(1 - \frac{a}{\lambda_i^*}\right) - \frac{h_k i}{\lambda_i^*} = \frac{1}{b}\left(1 - \frac{\sqrt{abh_k i}}{a}\right) - \frac{h_k i}{\sqrt{abh_k i}} = \frac{1}{b} - 2\sqrt{\frac{\frac{a}{4bk}i}{ab}} \\ &= \frac{1}{b}\left(1 - \sqrt{\frac{i}{k}}\right)\end{aligned}$$

where the first equality follows from (2.5), the third equality follows from (B.3) using $\gamma = 0$, and the fourth equality follows from the definition of $h_k$. One can check that $g_{k,0}(\lambda_k^*) = 0$ and

$K_k = \sum_{i=1}^{k} g_{i,0}(\lambda_i^*)$, thus proving our claim.

**(ii)** The holding cost rate for the optimal dynamic policy is,

$$\textsc{Hold}(k, \boldsymbol{\lambda}^*) = \frac{\sum_{i=1}^{k} \frac{hi}{\lambda_i^*}}{\sum_{i=1}^{k} \frac{1}{\lambda_i^*}} = \frac{h_k \sum_{i=1}^{l} \sqrt{i}}{\sum_{i=1}^{k} \frac{1}{\sqrt{i}}} \leq \frac{\int_1^{k+1} \sqrt{x} dx}{\int_1^{k+1} \frac{1}{\sqrt{x}} dx} = \frac{h_k}{3} \cdot \frac{(k+1)^{3/2} - 1}{(k+1)^{1/2} - 1},$$

where the second equality uses the fact that $\lambda_i^* = \sqrt{abh_k i}$ from (B.3) and the fact that the profit is 0 from part (i). As the holding cost ratio for static policy $(k, \lambda_{static} \cdot \mathbf{1})$ is $h_k(k+1)/2$, the holding cost ratio for policy $(k, \lambda_{static} \cdot \mathbf{1})$ can be bounded below as,

$$\frac{\textsc{Hold}(k, \lambda_{static} \cdot \mathbf{1})}{\textsc{Hold}(k, \boldsymbol{\lambda}^*)} \geq \frac{3}{2} \cdot \frac{(k+1)^{1/2} - 1}{(k+1)^{1/2} - \frac{1}{k+1}} \geq \frac{3}{2}\left(1 - \frac{1}{\sqrt{k+1}}\right).$$

**(iii)** For the total cost ratio, we repeat the analysis from Example 2.9. We know that for any order-up-to level $S = 1, \ldots, k$,

$$\frac{\textsc{Costs}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(k, \boldsymbol{\lambda}^*)} \geq \min\left\{\frac{k}{S}, \frac{S+1}{2} \cdot \frac{3((k+1)^{1/2} - 1)}{(k+1)^{3/2} - 1}\right\}.$$

Let $\widetilde{S}(k)$ be the minimizer of right-hand side. Then,

$$\widetilde{S}(k) := -\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2k((k+1)^{3/2} - 1)}{3((k+1)^{1/2} - 1)}}.$$

Consequently, the total cost ratio for instance $k$ and policy $(S, \lambda_{static} \cdot \mathbf{1})$ for any order-up-to level $S = 1, \ldots, k$ is bounded below as follows.

$$\frac{\textsc{Costs}(S, \lambda_{static} \cdot \mathbf{1})}{\textsc{Costs}(k, \boldsymbol{\lambda}^*)} \geq \frac{\widetilde{S}(k) + 1}{2} \cdot \frac{(k+1)^{1/2} - 1}{(k+1)^{3/2} - 1} \geq \sqrt{\frac{3k((k+1)^{1/2} - 1)}{2((k+1)^{3/2} - 1)}} \geq \sqrt{\frac{3k}{2(k+1)}\left(1 - \frac{1}{\sqrt{k+1}}\right)},$$

where the last inequality follows from part (ii) □

### B.1.5 Proof of Theorem 2.13

*Proof.* Let $T_+(S, \boldsymbol{\lambda})$, $T_\circ(\boldsymbol{\lambda})$, $T_-(s, \boldsymbol{\lambda})$ be the expected amount of time that the system spends in positive, 0 and negative inventory states between orders. Then,

$$T_+(S, \boldsymbol{\lambda}) = \sum_{i=1}^{S} \frac{1}{\lambda_i}, \qquad T_\circ(\boldsymbol{\lambda}) = \frac{1}{\lambda_0}, \qquad T_-(s, \boldsymbol{\lambda}) = \sum_{i=s+1}^{-1} \frac{1}{\lambda_i}.$$

We also divide our revenues and cost rates across these three regimes as follows:

$$\begin{aligned}
&\textsc{Rev}_+(S,\boldsymbol{\lambda}) = \frac{\sum_{i=1}^{S} p(\lambda_i)}{T_+(S,\boldsymbol{\lambda})}, && \textsc{Order}_+(S,\boldsymbol{\lambda}) = \frac{K}{T_+(S,\boldsymbol{\lambda})}, && \textsc{Hold}_+(S,\boldsymbol{\lambda}) = \sum_{i=1}^{S} ih\pi_{i,+}(S,\boldsymbol{\lambda}),\\
&\textsc{Rev}_\circ(\boldsymbol{\lambda}) = \frac{p(\lambda_0)}{T_\circ(\boldsymbol{\lambda})}, && \textsc{Order}_\circ(\boldsymbol{\lambda}) = \frac{K}{T_\circ(\boldsymbol{\lambda})}, && \textsc{Hold}_\circ(\boldsymbol{\lambda}) = 0,\\
&\textsc{Rev}_-(s,\boldsymbol{\lambda}) = \frac{\sum_{i=s+1}^{-1} p(\lambda_i)}{T_-(s,\boldsymbol{\lambda})}, && \textsc{Order}_-(s,\boldsymbol{\lambda}) = \frac{K}{T_-(s,\boldsymbol{\lambda})}, && \textsc{Hold}_-(s,\boldsymbol{\lambda}) = \sum_{i=s+1}^{-1} -i\pi\pi_{i,-}(s,\boldsymbol{\lambda}),
\end{aligned} \tag{B.7}$$

where $\pi_{i,y}(\cdot, \boldsymbol{\lambda}) = \frac{1/\lambda_i}{T_y(\cdot,\boldsymbol{\lambda})}$ for $y \in \{+, -\}$. For any metric, the overall value can be re-constructed as follows:

$$\textsc{Metric}(s,S,\boldsymbol{\lambda}) = \textsc{Metric}_+(S,\boldsymbol{\lambda})\frac{T_+(S,\boldsymbol{\lambda})}{T(s,S,\boldsymbol{\lambda})} + \textsc{Metric}_\circ(\boldsymbol{\lambda})\frac{T_\circ(\boldsymbol{\lambda})}{T(s,S,\boldsymbol{\lambda})} + \textsc{Metric}_-(s,\boldsymbol{\lambda})\frac{T_-(s,\boldsymbol{\lambda})}{T(s,S,\boldsymbol{\lambda})}. \tag{B.8}$$

Observe the similarities between the metrics in (B.7) and those defined for the case without backlogging in (2.4). For the positive inventory regime, the metrics are equivalent to those for an instance without backlogging with parameters $(h, K)$ for holding and ordering costs and an order-up-to level $S$. For the negative inventory regime, the metrics are equivalent to those for an instance without backlogging with parameters $(\pi, K)$ for holding and ordering costs and an order-up-to level $-s-1$. Moreover, the system in the positive inventory regime has no dependence on $s$ or $\lambda_i$ for $i = s+1, \ldots, 0$ and vice versa for the negative inventory regime and $S$, $\lambda_i$ for i in $0, \ldots S$.

Combining this and the definition of $\boldsymbol{\lambda}_{3p}$, it follows from Section 2.3.3 that

$$\frac{\textsc{Rev}_+(S,\lambda_+\cdot\mathbf{1})}{\textsc{Rev}_+(S^*,\boldsymbol{\lambda}^*)}\ge 1,\quad \frac{\textsc{Order}_+(S,\lambda_+\cdot\mathbf{1})}{\textsc{Order}_+(S^*,\boldsymbol{\lambda}^*)}=\frac{S^*}{S},\qquad \frac{\textsc{Hold}_+(S,\lambda_+\cdot\mathbf{1})}{\textsc{Hold}_+(S^*,\boldsymbol{\lambda}^*)}\le\frac{\ln(eS^*)(S+1)}{2S^*},$$

$$\frac{\textsc{Rev}_\circ(\lambda_\circ\cdot\mathbf{1})}{\textsc{Rev}_\circ(,\boldsymbol{\lambda}^*)}=1,\qquad \frac{\textsc{Order}_\circ(\lambda_\circ\cdot\mathbf{1})}{\textsc{Order}_\circ(,\boldsymbol{\lambda}^*)}=1,\qquad \frac{\textsc{Hold}_\circ(\lambda_\circ\cdot\mathbf{1})}{\textsc{Hold}_\circ(,\boldsymbol{\lambda}^*)}=1,$$

$$\frac{\textsc{Rev}_-(s,\lambda_-\cdot\mathbf{1})}{\textsc{Rev}_-(s^*,\boldsymbol{\lambda}^*)}\ge 1,\quad \frac{\textsc{Order}_-(s,\lambda_-\cdot\mathbf{1})}{\textsc{Order}_-(s^*,\boldsymbol{\lambda}^*)}=\frac{-s^*-1}{-s-1},\quad \frac{\textsc{Hold}_-(s,\lambda_-\cdot\mathbf{1})}{\textsc{Hold}_-(s^*,\boldsymbol{\lambda}^*)}\le\frac{\ln(e(-s^*-1))(-s)}{2(-s^*-1)}. \tag{B.9}$$

In Theorem 2.13 we analyze the policy at $S=S^*$ and $s=s^*$. Note that, when $S=S^*$ and $s=s^*$, the time spent in every regime is the same for both static and dynamic policies. Using this fact, we prove the bounds stated in the theorem.

The revenue ratio can be found by seeing that

$$\begin{aligned}
\textsc{Rev}(s^*,S^*,\boldsymbol{\lambda}_{3p}) &= \textsc{Rev}_+(S^*,\lambda_+\cdot\mathbf{1})\frac{T_+(S^*,\lambda_+\cdot\mathbf{1})}{T(s^*,S^*,\boldsymbol{\lambda}_{3p})}+\textsc{Rev}_\circ(\lambda_\circ\cdot\mathbf{1}))\frac{T_\circ(\lambda_\circ\cdot\mathbf{1})}{T(s^*,S^*,\boldsymbol{\lambda}_{3p})}\\
&\quad+\textsc{Rev}_-(s^*,\lambda_-\cdot\mathbf{1})\frac{T_-(s^*,\lambda_-\cdot\mathbf{1})}{T(s^*,S^*,\boldsymbol{\lambda}_{3p})}\\
&\ge \textsc{Rev}_+(S^*,\boldsymbol{\lambda}^*)\frac{T_+(S^*,\boldsymbol{\lambda}^*)}{T(s^*,S^*,\boldsymbol{\lambda}^*)}+\textsc{Rev}_\circ(\boldsymbol{\lambda}^*)\frac{T_\circ(\boldsymbol{\lambda}^*)}{T(s^*,S^*,\boldsymbol{\lambda}^*)}+\textsc{Rev}_-(s^*,\boldsymbol{\lambda}^*)\frac{T_-(s^*,\boldsymbol{\lambda}^*)}{T(s^*,S^*,\boldsymbol{\lambda}^*)}\\
&= \textsc{Rev}(s^*,S^*,\boldsymbol{\lambda}^*),
\end{aligned}$$

where the inequality follows from (B.9) and the fact that $(s^*,S^*,\boldsymbol{\lambda}_{3p})$ spends the same time in every inventory regime as the optimal policy by construction and the equalities follow from (B.8).

Similarly, the ordering cost ratio can be found by seeing that

$$\begin{aligned}\textsc{Order}(s^*, S^*, \boldsymbol{\lambda}_{3p}) &= \textsc{Order}_+(S^*, \lambda_+ \cdot \mathbf{1})\frac{T_+(S^*, \lambda_+ \cdot \mathbf{1})}{T(s^*, S^*, \boldsymbol{\lambda}_{3p})} + \textsc{Order}_\circ(\lambda_\circ \cdot \mathbf{1}))\frac{T_\circ(\lambda_\circ \cdot \mathbf{1})}{T(s^*, S^*, \boldsymbol{\lambda}_{3p})} \\ &\quad + \textsc{Order}_-(s^*, \lambda_- \cdot \mathbf{1})\frac{T_-(s^*, \lambda_- \cdot \mathbf{1})}{T(s^*, S^*, \boldsymbol{\lambda}_{3p})} \\ &= \textsc{Order}_+(S^*, \boldsymbol{\lambda}^*)\frac{T_+(S^*, \boldsymbol{\lambda}^*)}{T(s^*, S^*, \boldsymbol{\lambda}^*)} + \textsc{Order}_\circ(\boldsymbol{\lambda}^*)\frac{T_\circ(\boldsymbol{\lambda}^*)}{T(s^*, S^*, \boldsymbol{\lambda}^*)} \\ &\quad + \textsc{Order}_-(s^*, \boldsymbol{\lambda}^*)\frac{T_-(s^*, \boldsymbol{\lambda}^*)}{T(s^*, S^*, \boldsymbol{\lambda}^*)} \\ &= \textsc{Order}(s^*, S^*, \boldsymbol{\lambda}^*).\end{aligned}$$

Finally, the holding cost ratio can be found by seeing that

$$\begin{aligned}\textsc{Hold}(S^*, \lambda_{static} \cdot \mathbf{1}) &\leq \max\left\{\frac{\ln(eS^*)}{2}\left(1 + \frac{1}{S^*}\right), \frac{\ln(e(-s^* - 1))}{2}\left(1 + \frac{1}{-s^* - 1}\right)\right\}\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) \\ &\leq \frac{\ln(e(S^* - s^*))}{2}\left(1 + \frac{1}{S^* - s^*}\right)\textsc{Hold}(S^*, \boldsymbol{\lambda}^*) \leq (1 + \ln(S^* - s^*))\textsc{Hold}(S^*, \boldsymbol{\lambda}^*),\end{aligned}$$

where the last inequality uses the fact that $S^* - s^* \geq 1$. □

# Appendix C: Appendix to Chapter 3

## C.1 Omitted sections of the proof of Theorem 3.7

We restate (3.7), which we will prove in this section.

**Lemma C.1.** *If all products have the same price $r$, then for every opaque price $\rho < r$ and assortment $S \subseteq \mathcal{N}$, there exists $\alpha \in [0,1]$ which depends on the primitives $(S, r, \rho, \boldsymbol{v})$ such that,*

$$\textsc{Rev}(\textsc{Opq}(r\mathbf{1}_n, \rho)) - \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n)) = \alpha(\textsc{Rev}(\textsc{Trad}(\rho\mathbf{1}_n)) - \textsc{Rev}(\textsc{Trad}(r\mathbf{1}_n))).$$

We introduce some notation simplifications before we proceed. We shorten $\textsc{Rev}(\cdot)$, $\textsc{Trad}(\cdot)$, $\textsc{Opq}(\cdot)$ to $\mathrm{R}(\cdot)$, $\mathrm{T}(\cdot)$, $\mathrm{O}(\cdot)$ to optimize space. We also drop the **1** when stating prices since it is understood that pricing of traditional products is uniform. Accordingly, $\mathrm{O}(r, \rho)$ denotes the opaque selling mechanism where the traditional and opaque products are priced at price $r$ and $\rho$ respectively. Similarly, $\mathrm{T}(r)$ denotes the traditional selling mechanism where traditional products are priced at price $r$. We also overload notation, to let $\mathrm{T}(r, (\rho, I))$ denote the traditional assortment where all products in $I \subseteq S$ are priced at $\rho$ while products in $S \setminus I$ are still priced at $r$.

Additionally, we introduce some new notation to make the proof cleaner. For any $k \in S$ and $I \subseteq S$, we use the following shorthand:

$$\delta_k := \frac{e^{v_k}}{\sum_{i \in S} e^{v_i}} \quad \text{and} \quad \delta(I) = \sum_{i \in I} \delta_i.$$

We also let

$$u(r) := 1 + e^{-r} \sum_{i \in S} e^{v_i}$$

and

$$u(r,(\rho,I)) = 1 + e^{-r}\sum_{i\in S\setminus I} e^{v_i} + e^{-\rho}\sum_{i\in I} e^{v_i} = (1-\delta(I))u(r) + \delta(I)u(\rho) \tag{C.1}$$

denote the total attractiveness of traditional assortments where all products are priced uniformly at $r$ and products in $I$ are priced at $\rho$ respectively.

### C.1.1 Helper lemmas

In this section we state and prove lemmas that will be useful in the proof of Lemma C.1.

**Lemma C.2.** *The expected revenue earned from assortment* $\mathrm{T}(r,(\rho,I))$ *is a convex combination of the revenue earned from assortments* $\mathrm{T}(r)$ *and* $\mathrm{T}(\rho)$*. In particular,*

$$\mathrm{R}(\mathrm{T}(r,(\rho,I))) = \frac{(1-\delta(I))\mathrm{R}(\mathrm{T}(r))u(r) + \delta(I)\mathrm{R}(\mathrm{T}(\rho))u(\rho)}{(1-\delta(I))u(r) + \delta(I)u(\rho)}.$$

*Proof.*

$$\begin{aligned}
\mathrm{R}(\mathrm{T}(r,(\rho,I))) &= \frac{re^{-r}\sum_{i\in S\setminus I} e^{v_i} + \rho e^{-\rho}\sum_{i\in I} e^{v_i}}{1 + e^{-r}\sum_{i\in S\setminus I} e^{v_i} + e^{-\rho}\sum_{i\in I} e^{v_i}} \\
&= \frac{\frac{\sum_{i\in S\setminus I} e^{v_i}}{\sum_{i\in S} e^{v_i}} \cdot \frac{re^{-r}\sum_{i\in S} e^{v_i}}{1+e^{-r}\sum_{i\in S} e^{v_i}} u(r) + \frac{\sum_{i\in I} e^{v_i}}{\sum_{i\in S} e^{v_i}} \cdot \frac{\rho e^{-\rho}\sum_{i\in S} e^{v_i}}{1+e^{-\rho}\sum_{i\in S} e^{v_i}} u(\rho)}{(1-\delta(I))u(r) + \delta(I)u(\rho)} \\
&= \frac{(1-\delta(I))\mathrm{R}(\mathrm{T}(r))u(r) + \delta(I)\mathrm{R}(\mathrm{T}(\rho))u(\rho)}{(1-\delta(I))u(r) + \delta(I)u(\rho)},
\end{aligned}$$

where the second equality uses the definition of $u(r)$ in the numerator and (C.1) in the denominator, the third equality uses the definitions of $\delta(\cdot)$ and $\mathrm{T}(\cdot)$. □

**Lemma C.3.** *For sequence of scalars* $(z_i)_{i=0}^n$ *and scalar* $c$*, the following identity holds,*

$$\prod_{i=0}^{n}(z_i + c) - \prod_{i=0}^{n} z_i = c\sum_{j=0}^{n}\left(\prod_{i=0}^{j-1}(z_i + c)\prod_{i=j+1}^{n} z_i\right).$$

*Proof.* We prove the identity by induction on $n$.

**Base case:** Consider $n = 0$. The LHS for this case is $z_0 + c - z_0 = c$, while the RHS is $c \cdot 1 = 1$.

**Induction step:** Assume the identity is true for $n - 1$.

$$\begin{aligned}
\prod_{i=0}^{n}(z_i + c) - \prod_{i=0}^{n} z_i &= z_n \prod_{i=0}^{n-1}(z_i + c) + c\prod_{i=0}^{n-1}(z_i + c) - \prod_{i=0}^{n} z_i \\
&= z_n\left(\prod_{i=0}^{n-1}(z_i + c) - \prod_{i=0}^{n-1} z_i\right) + c\prod_{i=0}^{n-1}(z_i + c) \\
&= z_n c \sum_{j=0}^{n-1}\left(\prod_{i=0}^{j-1}(z_i + c) \prod_{i=j+1}^{n-1} z_i\right) + c\prod_{i=0}^{n-1}(z_i + c) \\
&= c\left(\sum_{j=0}^{n-1}\left(\prod_{i=0}^{j-1}(z_i + c) \prod_{i=j+1}^{n} z_i\right) + \prod_{i=0}^{n-1}(z_i + c)\right) \\
&= c\sum_{j=0}^{n}\left(\prod_{i=0}^{j-1}(z_i + c) \prod_{i=j+1}^{n} z_i\right),
\end{aligned}$$

where the third equality follows from the induction hypothesis. □

### C.1.2 Proof of Lemma C.1

Note that, from (3.4) it follows that,

$$\begin{aligned}
\mathrm{R}(\mathrm{O}(r, \rho)) - \mathrm{R}(\mathrm{T}(r)) &= \sum_{\emptyset \neq I \subseteq S} (-1)^{|I|+1}\mathrm{R}(\mathrm{T}(r, (\rho, I))) - \mathrm{R}(\mathrm{T}(r)) \qquad \text{(C.2)} \\
&= \sum_{I \subseteq S} (-1)^{|I|+1}\mathrm{R}(\mathrm{T}(r, (\rho, I))).
\end{aligned}$$

To prove the lemma, we will prove the following more general claim: for any sets $S' \subset S$ and $S'' := \{k_1, k_2, \dots, k_l\}$ such that $S'' \subseteq S \setminus S'$ and $|S''| = l > 0$,

$$\sum_{I \subseteq S''} (-1)^{|I|+1}\mathrm{R}(\mathrm{T}(r, (\rho, S' \cup I))) = \alpha(S', S'')(\mathrm{R}(\mathrm{T}(\rho)) - \mathrm{R}(\mathrm{T}(r))),$$

where,

$$\alpha(S', S'') := u(r)u(\rho)(u(\rho) - u(r))^{l-1}\left(\prod_{i=1}^{l} \delta_{k_i}\right) \sum_{\alpha \in \mathcal{P}_l} \frac{1}{\prod_{i=0}^{l} u(r, (\rho, S' \cup S''_{\alpha([i])}))} \in [0, 1]. \quad \text{(C.3)}$$

Here, $\mathcal{P}_l$ is the set of all permutations of $[l]$. For a specific permutation $\alpha \in \mathcal{P}_l$ and integer $i = 0, \ldots, l$, we define $S''_{\alpha([i])} := \left\{k_{\alpha_j} \mid j \in [i]\right\}$. For example, if $S'' = \{1, 2, 3\}$ and $\alpha = \{3, 1, 2\}$, then $S''_{\alpha([1])} = \{3\}$ and $S''_{\alpha([2])} = \{3, 1\}$. Note that setting $S' = \emptyset$ and $S'' = S$ in the induction hypothesis proves the lemma.

We perform induction on the cardinality of $S''$, i.e., $l$ to prove the above claim.

**Base case:** We first prove the hypothesis is true when $l = 1$. Fix $S'$ and $k \in S \setminus S'$. Then,

$$\begin{aligned}
&\mathrm{R}(\mathrm{T}(r, (\rho, S' \cup \{k\}))) - \mathrm{R}(\mathrm{T}(r, (\rho, S')))\\
&\overset{(a)}{=} \frac{(1 - \delta(S') - \delta_k)\mathrm{R}(\mathrm{T}(r))u(r) + (\delta(S') + \delta_k)\mathrm{R}(\mathrm{T}(\rho))u(\rho)}{(1 - \delta(S') - \delta_k)u(r) + (\delta(S') + \delta_k)u(\rho)} - \frac{(1 - \delta(S'))\mathrm{R}(\mathrm{T}(r))u(r) + \delta(S')\mathrm{R}(\mathrm{T}(\rho))u(\rho)}{(1 - \delta(S'))u(r) + \delta(S')u(\rho)}\\
&\overset{(b)}{=} \frac{\delta_k(1 - \delta(S') + \delta(S'))u(r)u(\rho)(\mathrm{R}(\mathrm{T}(\rho)) - \mathrm{R}(\mathrm{T}(r)))}{u(r, (\rho, S'))u(r, (\rho, S' \cup \{k\}))} = \frac{\delta_k u(r)u(\rho)(\mathrm{R}(\mathrm{T}(\rho)) - \mathrm{R}(\mathrm{T}(r)))}{u(r, (\rho, S'))u(r, (\rho, S' \cup \{k\}))}\\
&\overset{(c)}{=} \alpha(S', S' \setminus \{k\})(\mathrm{R}(\mathrm{T}(\rho)) - \mathrm{R}(\mathrm{T}(r))),
\end{aligned}$$

where (a) follows from Lemma C.2, (b) uses (C.1) and (c) follows from (C.3). To see that $\alpha(S', S' \setminus \{k\}) \in [0, 1]$ note that,

$$\begin{aligned}
u(r, (\rho, S'))u(r, (\rho, S' \cup \{k\})) &= ((1 - \delta(S'))u(r) + \delta(S')u(\rho))((1 - \delta(S') - \delta_k)u(r) + (\delta(S') + \delta_k)u(\rho))\\
&\geq u(r)u(\rho)((1 - \delta(S'))(\delta(S') + \delta_k) + \delta(S')(1 - \delta(S') - \delta_k))\\
&= u(r)u(\rho)(\delta_k + 2\delta(S')(1 - \delta(S') - \delta_k))\\
&\geq u(r)u(\rho)\delta_k,
\end{aligned}$$

where the first equality follows from (C.1), the second inequality follows from the fact that $\delta(S') +$

$\delta_k = \delta(S' \cup \{k\}) \leq 1$.

**Induction step:** Assume that the induction hypothesis is true for all $S'$ and $S''$ such that $|S''| = l-1$. We show that the induction hypothesis also holds for all $S'$ and $S''$ such that $|S''| = l$. Fix $S'$ and $S''$ such that $|S''| = l$.

$$\begin{aligned}
&\sum_{I \subseteq S''} (-1)^{|I|+1} \mathrm{R}(\mathrm{T}(r, (\rho, S' \cup I))) \\
=& \sum_{I \subseteq S'' \setminus \{k_l\}} (-1)^{|I|+1} \mathrm{R}(\mathrm{T}(r, (\rho, S' \cup I))) - (-1)^{|I|+1} \mathrm{R}(\mathrm{T}(r, (\rho, S' \cup \{k_l\} \cup I))) \\
=& (\alpha(S', S'' \setminus \{k_l\}) - \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}))(\mathrm{R}(\mathrm{T}(\rho)) - \mathrm{R}(\mathrm{T}(r))),
\end{aligned}$$

where the last equality follows from the induction hypothesis. What remains to be shown is that $\alpha(S', S'' \setminus \{k_l\}) - \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}) = \alpha(S', S'')$. To see this note that

$$\begin{aligned}
&\alpha(S', S'' \setminus \{k_l\}) - \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}) \qquad \text{(C.4)} \\
=& u(r)u(\rho)(u(\rho) - u(r))^{l-2} \left( \prod_{i=1}^{l-1} \delta_{k_i} \right) \sum_{\alpha \in \mathcal{P}_{l-1}} \left( \frac{1}{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])}))} - \frac{1}{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])}))} \right) \\
=& u(r)u(\rho)(u(\rho) - u(r))^{l-2} \left( \prod_{i=1}^{l-1} \delta_{k_i} \right) \sum_{\alpha \in \mathcal{P}_{l-1}} \frac{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])})) - \prod_{i=0}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])}))}{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])})) u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])}))},
\end{aligned}$$

where the first equality follows from (C.3). From (C.1) we know that,

$$u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])})) = u(r, (\rho, S' \cup S''_{\alpha([i])})) + \delta_{k_l}(u(\rho) - u(r)).$$

Therefore, using Lemma C.3 with $z_i = u(r, (\rho, S' \cup S''_{\alpha([i])}))$ for $i = 1, \ldots, l-1$ and $c = \delta_{k_l}(u(\rho) -$

$u(r)$) in (C.4) we get,

$$\begin{aligned}
&\alpha(S', S'' \setminus \{k_l\}) - \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}) \\
=&u(r)u(\rho)(u(\rho) - u(r))^{l-1}\left(\prod_{i=1}^{l} \delta_{k_i}\right) \sum_{\alpha \in \mathcal{P}_{l-1}} \frac{\sum_{j=0}^{l-1} \prod_{i=0}^{j-1} u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])})) \prod_{i=j+1}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])}))}{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])}))u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])}))} \\
=&u(r)u(\rho)(u(\rho) - u(r))^{l-1}\left(\prod_{i=1}^{l} \delta_{k_i}\right) \sum_{\alpha \in \mathcal{P}_{l-1}} \sum_{j=0}^{l-1} \frac{1}{\prod_{i=0}^{j} u(r, (\rho, S' \cup S''_{\alpha([i])})) \prod_{i=j}^{l-1} u(r, (\rho, S' \cup \{k_l\} \cup S''_{\alpha([i])}))} \\
=&u(r)u(\rho)(u(\rho) - u(r))^{l-1}\left(\prod_{i=1}^{l} \delta_{k_i}\right) \sum_{\alpha \in \mathcal{P}_{l}} \frac{1}{\prod_{i=0}^{l-1} u(r, (\rho, S' \cup S''_{\alpha([i])}))} \\
=&\alpha(S', S'').
\end{aligned}$$

To see that $\alpha(S', S'') \in [0, 1]$, note that

$$\alpha(S', S'' \setminus \{k_l\}), \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}) \leq 1 \text{ and } \alpha(S', S'' \setminus \{k_l\}) \geq \alpha(S' \cup \{k_l\}, S'' \setminus \{k_l\}),$$

where the first inequality follows from the induction hypothesis and the second inequality follows from (C.3) and the fact that $u(r, (\rho, S' \cup \{k_l\} \cup I)) \geq u(r, (\rho, S' \cup I))$ for all $I \subseteq S'' \setminus \{k_l\}$. This completes the proof of the induction hypothesis. From (C.2) it follows that setting $S' = \emptyset$ and $S'' = S$ in the induction hypothesis proves the lemma. ☐

## C.2 Proof of Theorem 3.8

Let $R(\cdot, \cdot)$ denote the expected revenue as a function of the prices under traditional selling. For prices $(x_1, x_2)$,

$$R(x_1, x_2) := \frac{x_1 e^{v_1 - x_1} + x_2 e^{v_2 - x_2}}{1 + e^{v_1 - x_1} + e^{v_2 - x_2}}.$$

Further, let $R_q(\cdot, \cdot, \cdot)$ denote the expected revenue as a function of the prices (traditional and opaque) under opaque selling. If the traditional products are offered at prices $(r_1, r_2)$ and the opaque product is offered at price $\rho$, then revenue for risk-averse customers under opaque selling is given by,

$$R_q(r_1, r_2, \rho) := R(r_1, \rho) + R(\rho, r_2) - R(\rho, \rho),$$

which follows from (3.4) as long as $\rho \leq \min\{r_1, r_2\}$.

The first-order conditions of $R_q(r_1, r_2, \rho)$ are given by

$$\begin{aligned}
\frac{\partial R_q(r_1, r_2, \rho)}{\partial r_1} =& \frac{\partial R(r_1, \rho)}{\partial r_1} = \frac{e^{v_1 - r_1}}{1 + e^{v_1 - r_1} + e^{v_2 - \rho}} (R(r_1, \rho) + 1 - r_1) = 0 \\
\frac{\partial R_q(r_1, r_2, \rho)}{\partial r_2} =& \frac{\partial R(\rho, r_2)}{\partial r_2} = \frac{e^{v_2 - r_2}}{1 + e^{v_1 - \rho} + e^{v_2 - r_2}} (R(\rho, r_2) + 1 - r_2) = 0 \\
\frac{\partial R_q(r_1, r_2, \rho)}{\partial \rho} =& \frac{e^{v_2 - \rho}(R(r_1, \rho) + 1 - \rho)}{1 + e^{v_1 - r_1} + e^{v_2 - \rho}} + \frac{e^{v_1 - \rho}(R(\rho, r_2) + 1 - \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - r_2}} - \frac{(e^{v_1 - \rho} + e^{v_2 - \rho})(R(\rho, \rho) + 1 - \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} = 0
\end{aligned}$$

$$\begin{aligned}
\frac{\partial R_q(r_1, r_2, \rho)}{\partial \rho} =& e^{v_1 - \rho} \left( \frac{R(\rho, r_2) + 1 - \rho}{1 + e^{v_1 - \rho} + e^{v_2 - r_2}} - \frac{R(\rho, \rho) + 1 - \rho}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} \right) + e^{v_2 - \rho} \left( \frac{R(r_1, \rho) + 1 - \rho}{1 + e^{v_1 - r_1} + e^{v_2 - \rho}} - \frac{R(\rho, \rho) + 1 - \rho}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} \right) \\
\geq& e^{v_1 - \rho} \left( \frac{R(\rho, r_2) + 1 - \rho - (R(\rho, \rho) + 1 - \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} \right) + e^{v_2 - \rho} \left( \frac{R(r_1, \rho) + 1 - \rho - (R(\rho, \rho) + 1 - \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} \right) \\
=& e^{v_1 - \rho} \cdot \frac{R(\rho, r_2) - R(\rho, \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}} + e^{v_2 - \rho} \cdot \frac{R(r_1, \rho) - R(\rho, \rho)}{1 + e^{v_1 - \rho} + e^{v_2 - \rho}}
\end{aligned}$$

Let $r_1^*, r_2^*, \rho^*$ denote the optimal prices under the opaque setting and $r^*$ be the optimal price under the traditional setting. We know that

$$R(r_1^*, \rho^*) = \max_{r_1} R(r_1, \rho^*) \geq R(\rho^*, \rho^*),\ R(\rho^*, r_2^*) = \max_{r_2} R(\rho^*, r_2) \geq R(\rho^*, \rho^*).$$

Therefore, we have

$$0 = \frac{\partial R_q(r_1, r_2, \rho)}{\partial \rho}\Big|_{(r_1^*, r_2^*, \rho^*)} \geq e^{v_1 - \rho^*} \cdot \frac{R(\rho^*, r_2^*) - R(\rho^*, \rho^*)}{1 + e^{v_1 - \rho^*} + e^{v_2 - \rho^*}} + e^{v2 - \rho^*} \cdot \frac{R(r_1^*, \rho^*) - R(\rho^*, \rho^*)}{1 + e^{v_1 - \rho^*} + e^{v_2 - \rho^*}} \geq 0.$$

It is known that the traditional revenue function is unimodal in each of the prices, which combined with the above implies that $r_1^* = \rho^* = r_2^*$. Moreover, since $\frac{\partial R(r_1, r_2)}{\partial r_1}\big|_{(r_1^*, r_2^*)} = \frac{\partial R(r_1, r_2)}{\partial r_2}\big|_{(r_1^*, r_2^*)} = 0$, it holds that $r_1^* = r_2^* = \rho^* = r^*$. □