Approximation Schemes for Dynamic Pricing with Opaque Products

Jacob Feldman, Yukai Huang

Dept. of Supply Chain, Operations & Technology, Olin Business School, Washington University in St Louis, jbfeldman@wustl.edu, yukaihuang@wustl.edu

Xingxing Chen

Dept. of Analytics & Operations, Robins School of Business, University of Richmond

chen@richmond.edu

Abstract

Problem definition: In this paper, we consider a multi-period, multi-product dynamic pricing problem in which each product is endowed with an exogenous starting inventory level, and there is the added complexity of an opaque selling option. That is, alongside traditional (transparent) products, the retailer or platform also has the option to create and price an opaque product, which corresponds to a dummy product comprised potentially of any subset of the displayed transparent products. In the event that a customer selects the opaque product, the platform has the freedom to choose any of the opaque product's constituents to satisfy this demand. All-in-all, we are left with a classical dynamic pricing problem with a twist, since the addition of the opaque selling option gives the platform an extra lever of flexibility to balance supply and demand. **Methodology/results:** We begin by studying a variant of the problem in which the retailer only offers and prices an opaque product. In this initial setting, we exploit the special structural of the optimal solution to a fluid approximation of our dynamic pricing problem to yield a policy whose performance relative to optimal grows from $1 - \frac{1}{e}$ to 1 as the initial inventory levels tend to infinity. Next, we consider a setting with both transparent and opaque products, and provide a constant factor approximation scheme for this more nuanced version of the problem. Our approach builds on top of the inventory-tracking basis function approximation originally conceived by Ma et al. (2020) for network revenue management problems. We also include two distinct sets of the computational experiments, the first of which demonstrates the efficacy of our approximation scheme for the opaque-only setting, whereas the second uses our approach for the general setting to study the revenue gains afforded to platforms that exploit the use of an opaque selling option. **Managerial implications:** We provide the first approximation schemes for multi-period, multi-product dynamic pricing problems with an opaque selling option. Moreover, we exploit these novel algorithms to show that effectively pricing said opaque option can lead to 5-7% revenue gains over an approach that offers only transparent products.

Key words: approximation schemes, opaque products, dynamic pricing, online matching

1. Introduction

It is without a doubt that online matching has become a foundational topic in the field of revenue management and beyond. In its most general and abstract form, this ever-popular problem setting captures the task of matching a finite collection of heterogeneous supply units (e.g. products, appointment times, computing resources, etc ...) to stochastic demand in a multi-period dynamic environment. Its steady ascent to its current status as a central topic of study can be, in part, attributed to its initial relevance, and eventual prominence, within a host of industries ranging from retail to hospitality to ridesharing, where various iterations of online matching problems have been at the forefront of internal operations. This picture of popularity is completed by pointing to the intriguing breadth and depth of algorithmic questions that have arisen in this context, and that continue to spawn new and exciting research topics.

Indeed, variants of the traditional online bi-partite matching problem studied by the seminal works of Karp et al. (1990), Mehta et al. (2007), and Aggarwal et al. (2011), to name a few, have been re-shaped and extended in a multitude of ways that add novel practical and theoretical nuances. One such example includes the classical network revenue management problem (Talluri and Van Ryzin 1998, Adelman 2007, Topaloglu 2009), which considers a variant of online matching in which matched units of demand consume a combination of atomic resources (e.g. the purchase of an airline ticket potentially consumes seats on multiple flight legs). Initial formulations of the network revenue management problem were expanded by adding elements of substitution behavior via customer choice models (Talluri and Van Ryzin 2004, Liu and Van Ryzin 2008, Gallego et al. 2015b). Applications in this context range from multi-period dynamic pricing akin to the one studied in the work at-hand (Ma et al. 2021) to appointment scheduling in hospital clinics (Gallego et al. 2015a). The recent works of Rusmevichientong et al. (2020), Gong et al. (2022) and Feng et al. (2019) add a completely new dimension to choice-based online matching problems by considering the prospect of reusable resources, i.e. once a supply unit is consumed, it returns for future use after a (possibly random) duration of time. Finally, Aouad and Saban (2022) innovate by studying a two-sided matching problem wherein there is an element of choice on both the demand and supply sides.

Our work is intended to represent another branch in the growing stream of research that considers practically motivated adaptations of the classical choice-based online matching problem. Namely, we study a multi-period, multi-product dynamic pricing problem with an opaque selling option. More specifically, in addition to the traditional assortment/pricing decision levers, we assume that the retailer or platform has the option to offer a supplementary opaque option, which is a dummy product comprised of some subset of the available alternatives. A customer who selects the opaque option is at the mercy of the platform with regards to the item she receives, i.e. she is guaranteed to receive an item that is a constituent of the opaque option, but the platform selects the item she ultimately receives. As such, there is inherent consumer-side uncertainty in selecting the opaque option, which needs to be accounted for within the presumed choice model that governs purchasing behavior. Platforms generally counter this uncertainty by offering the opaque option at a discount, as is depicted in the example at Hotwire in Figure 1. All-in-all, the opaque selling option gives platforms such as Hotwire an additional lever to effectively balance supply and demand. Specifically, they gain the option to offer a discounted opaque product in exchange for flexibility in the subsequent allocation decision if said opaque product is purchased. In the remainder of this section, we formalize the exact nature of this trade-off through a rigorous formulation of our dynamic pricing problem of interest (Section 1.1), summarize past work that considers managerial and algorithmic implications of adopting opaque selling strategies (Section 1.2), and detail our main contributions, which are mostly algorithmic in nature (Section 1.3).

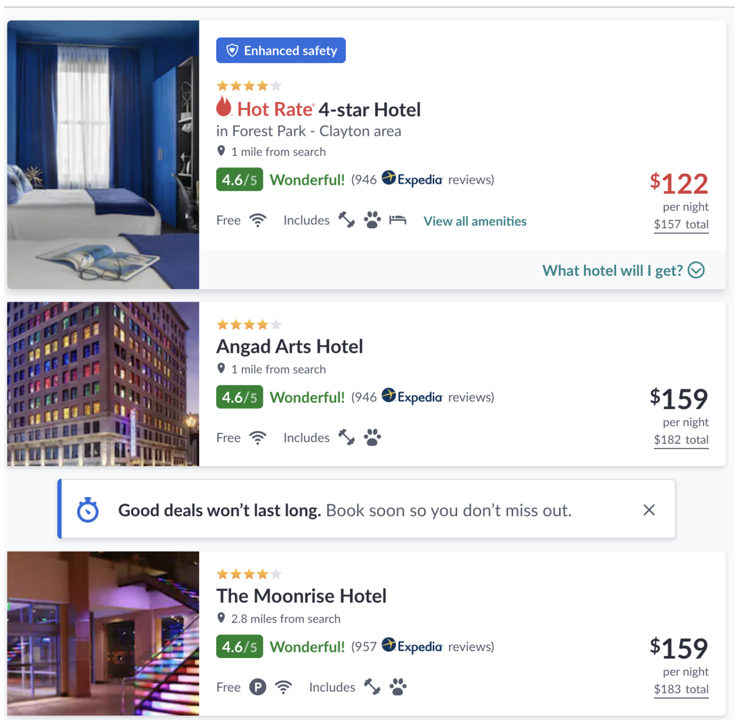

Figure 1 Four-star filtered hotel search on Hotwire. Note that there is a discounted opaque option followed by two transparent options.

1.1. Problem Formulation

We consider the dynamic pricing of n products indexed by the set $\mathcal{N} = \{1, \ldots, n\}$, each endowed with exogenously given starting inventory levels $U_1 \in \mathbb{N}^n$, over a finite selling horizon that has been partitioned into T time periods. In each time period, the retailer selects an assortment $S \subseteq \mathcal{N}$ of transparent products, as well as a subset of products $S_{\mathcal{O}} \subseteq S$ constituting the opaque product, to make available for purchase. The transparent products represent "typical" products, whose identity is fully revealed to the customer upon being made available for purchase. The opaque product, on the the other hand, is a dummy product represented by the subset of products $S_{\mathcal{O}}$. A customer who selects the opaque product agrees to allow the retailer to provide them with any product $i \in S_{\mathcal{O}}$.

The restriction that $S_{\mathcal{O}} \subseteq S$ is imposed to prevent consumer-side confusion/frustration that might arise if a particular product appears only within the opaque product, and hence cannot be bought outright as a transparent product. As is eventually formalized in the sequel, we enforce that only in-stock products can be offered in each period, either as transparent options or as a constituent of the opaque product.

In addition to the two-fold assortment decisions described above, the retailer is also tasked with pricing each product $i \in S$, as well the opaque product $S_{\mathcal{O}}$ (if $S_{\mathcal{O}} \neq \emptyset$), from a discrete set of L potential price levels. We use p_l to denote the l-th price level and the binary vector $(x_{i\ell})_{i \in \mathcal{N} \cup \{\mathcal{O}\}, \ell \in [L]}$ to represent the pricing decisions, where $x_{i\ell} = 1$ if product i is priced at level ℓ, and $x_{i\ell} = 0$ otherwise. For fixed assortment decisions S and $S_{\mathcal{O}}$, we let

$$\mathcal{X}(S, S_{\mathcal{O}}) = \Bigg\{ (x_{i\ell})_{i \in \mathcal{N} \cup \{\mathcal{O}\}, \ell \in [L]} : \sum_{\ell \in [L]} x_{i\ell} = \mathbb{1}_{i \in S} \ \forall i \in \mathcal{N}, \ \sum_{\ell \in [L]} x_{\mathcal{O}\ell} = \mathbb{1}_{S_{\mathcal{O}} \neq \emptyset}, \\ p_i(x) \geq p_{\mathcal{O}}(x) \ \forall i \in S_{\mathcal{O}} \Bigg\}$$

denote all feasible pricing decisions, where $p_j(x) = \sum_{\ell \in [L]} p_\ell x_{j\ell}$ denotes the price charged for any product $j \in S \cup \{\mathcal{O}\}$. The first set of constraints enforces that all products offered as a transparent product must be priced at one of the L levels, while the second singleton constraint encodes a similar notion for the opaque product. The final set of constraints restricts the price of the opaque product to be below the posted prices of all transparent products that make-up the opaque product. This latter constraint reflects what is typically observed in practice (See Figure 1), namely, the opaque product is priced below all of its constituent transparent products.

The general demand model. During each period, a single customer arrives and is assumed to make an MNL-based choice among the offered transparent products and their opaque counterpart. A noteworthy modeling challenge that arises with the addition of opaque selling is how to aptly quantify the choice probability of the opaque option when customers have different risk profiles. For example, risk-averse customers are likely to be pessimistic, and thus assume that they will be allocated their least preferred product among those that comprise the opaque product. On the other hand, less risk-averse or more flexible customers are likely to take on a more balanced perception of the opaque product, and hence value its contents using an aggregate measure (e.g., average) of their perceived preferences of the products that make-up the opaque product. Our initially proposed MNL-based demand model, detailed subsequently, makes no explicit assumptions on how customers value the opaque product, however, for tractability, future sections will require more rigid assumptions regarding this modeling choice. When it becomes time to levy such assumptions, we make explicit note of them, while concurrently providing practical motivation to back their use.

In our general demand model, under assortment S of transparent products and opaque product $S_{\mathcal{O}}$, along with pricing decisions $x \in \mathcal{X}(S, S_{\mathcal{O}})$, we use $\pi_i(x, S_{\mathcal{O}})$ to denote the probability that product $i \in \mathcal{N} \cup \{\mathcal{O}\}$ is selected for purchase [1]. We assume that $\pi_i(x, S_{\mathcal{O}})$ is governed by an MNL model, where $w_{i\ell} = w_i e^{\beta p_\ell}$ denotes the so-called preference weight of transparent product i when priced at level ℓ. Since the price coefficient is homogeneous across products, we can assume without loss of generality that the products are indexed in increasing order of weight for every price level $\ell \in [L]$, i.e. $w_{1\ell} \le w_{2\ell} \le \ldots \le w_{n\ell}$. Moreover, we use $w_{\mathcal{O}}(x, S_{\mathcal{O}})$ to capture the MNL-based weight of the opaque product, which, for now, we model as an arbitrary function of the posted prices for the transparent products and the contents of the opaque product itself. We also assume the weight of the opaque product is a decreasing function of its price. As such, following the well-established MNL framework (Luce 1959, McFadden 1974), we have that for transparent product $i \in \mathcal{N}$, the choice probability is

$$\pi_i(x, S_{\mathcal{O}}) = \frac{w_i(x)}{1 + \sum_{j \in \mathcal{N}} w_j(x) + w_{\mathcal{O}}(x, S_{\mathcal{O}})},$$

and for the opaque product, we have

$$\pi_{\mathcal{O}}(x, S_{\mathcal{O}}) = \frac{w_{\mathcal{O}}(x, S_{\mathcal{O}})}{1 + \sum_{j \in \mathcal{N}} w_j(x) + w_{\mathcal{O}}(x, S_{\mathcal{O}})}.$$

where $w_i(x) = \sum_{\ell \in [L]} w_{i\ell} x_{i\ell}$ denotes the weights of product i under pricing decision x.

Dynamic pricing with opaque products. In what follows, we formulate our dynamic pricing problem of interest as a dynamic program, whose value functions $V_t(U_t)$ represent the maximum expected reward that can be garnered by the retailer across periods $t, \ldots, T$, given that the inventories of all product at the start of period t is denoted by the vector $U_t = (u_{1t}, \ldots, u_{nt})$. For the purpose of formally defining the value functions, we let $\mathcal{N}(U_t) = \{i \in [n] : u_{it} > 0\}$ denote the set of products that have yet to stock-out, as specified by the inventory vector U_t. With this notation in-hand, we present the Bellman equations of our dynamic program below

$$\begin{aligned}
V_t(U_t) = \max_{\substack{S \subseteq \mathcal{N}(U_t),\\ S_{\mathcal{O}} \subseteq S}} \max_{x \in \mathcal{X}(S, S_{\mathcal{O}})} \Bigg\{ & \underbrace{\sum_{i \in S} \pi_i(x, S_{\mathcal{O}}) \cdot \Big(p_i(x) + V_{t+1}(U_t - e_i) \Big)}_{\text{expected revenue from transparent products}} + \\
& \underbrace{\pi_{\mathcal{O}}(x, S_{\mathcal{O}}) \cdot \Big(p_{\mathcal{O}}(x) + \max_{k \in S_{\mathcal{O}}} V_{t+1}(U_t - e_k) \Big)}_{\text{expected revenue from opaque product}} + \underbrace{\Big(1 - \sum_{i \in S} \pi_i(x, S_{\mathcal{O}}) - \pi_{\mathcal{O}}(x, S_{\mathcal{O}}) \Big) \cdot V_{t+1}(U_t)}_{\text{expected revenue if customer does not make a purchase}} \Bigg\} \\
= \max_{\substack{S \subseteq \mathcal{N}(U_t),\\ S_{\mathcal{O}} \subseteq S}} \max_{x \in \mathcal{X}(S, S_{\mathcal{O}})} \Bigg\{ & \sum_{i \in S} \pi_i(x, S_{\mathcal{O}}) \cdot \Big(p_i(x) - \Delta V^i_{t+1}(U_t) \Big) + \\
& \pi_{\mathcal{O}}(x, S_{\mathcal{O}}) \cdot \Big(p_{\mathcal{O}}(x) - \min_{k \in S_{\mathcal{O}}} \Delta V^k_{t+1}(U_t) \Big) \Bigg\} + V_{t+1}(U_t) \qquad (1)
\end{aligned}$$

[1] The dependence of this choice probability on S is not explicitly noted since the pricing decision x encodes the assortment of transparent products offered

where e_i is the unit vector with a single one in the i-th component and zeros in every other component, and $\Delta V_t^i(U_t) = V_t(U_t) - V_t(U_t - e_i)$ is the marginal value of single unit of product i at time t. We also define base cases of $V_{T+1}(\cdot) = 0$. Note that the recursion given in (1) reflects the fact that, if product $i \in S \cup \{\mathcal{O}\}$ is purchased, the retailer accrues a revenue of $p_i(x)$, and one unit of product i is consumed. Additionally, if the arriving customer selects the opaque product, then it is optimal to assign her a single unit of product $\arg\min_{k \in S_{\mathcal{O}}} \Delta V_{t+1}^k(U_t)$, i.e., the least valuable unit among those available in $S_{\mathcal{O}}$. Finally, it is important to note that, due to the high dimensional state space, it is not computationally tractable to compute the optimal policy via the dynamic program given in (1).

Modeling remark. Up until this point, we have yet to handle the important corner case in which $|S_{\mathcal{O}}| = 1$, i.e. the opaque option consists of only a single product. Since offering an opaque product that consists of a single product is equivalent to offering just the transparent version of this product, it is imperative that either our demand model reflects this nuance (as written, it currently does not), or that we disallow such a decision. We adopt the latter, which leads to the following assumption that is applied to both retroactively and to all future sections in which he retailer has the option to offer both transparent and opaque products.

Assumption 1. *When the retailer has the option to offer both transparent and opaque products, we enforce that the opaque option cannot be comprised of only a single product, i.e. we implicitly enforce the constraint that $|S_{\mathcal{O}}| \neq 1$ for all decisions concerning $S_{\mathcal{O}}$.*

1.2. Literature Review

In what follows, we summarize two camps of existing works that relate to opaque selling strategies.

Opaque selling - managerial insights. To the best of our knowledge, Fay and Xie (2008) were one of the first to consider the potential benefits of offering opaque products. They refer to such selling strategies as "probabilistic selling", since the products within the opaque option are to be allocated with fixed exogenous probabilities, which are known to the consumers. In a simplistic two product setting, this work shows that probabilistic selling can be an effective tool to deal with market uncertainty. Huang and Yu (2014) relax the assumption that consumers fully know the retailer's probabilistic allocation policy, and show that even in their so-called bounded rationality setting, opaque selling can still dominate traditional selling strategies. Jerath et al. (2010) and Ren and Huang (2022) both show that opaque selling often dominates traditional last minute discount policies when it comes to unsold inventory. Finally, Post and Spann (2012) detail the substantial revenue growth that ensued after Germanwings, one of Germany's top low-cost airline, adopted opaque selling for its unsold tickets.

Opaque selling - algorithmic insights and results. Anderson and Xie (2012) appear to be one of the first works to consider dynamic pricing with an opaque product, however their main focus is on developing and estimating a featurized nested logit model to capture customer choice is settings where a opaque option is available. Elmachtoub et al. (2015) and Elmachtoub and Hamilton (2021) were the first to explicitly model how customers might value an opaque product from a utility perspective. The former paper presents optimal re-stocking and allocation policies in a dynamic multi-period setting, while the latter paper considers a static pricing setting and shows that optimal opaque pricing is guaranteed to garner at least 71.9% of the revenue that could be afforded to a policy that optimally prices all products individually. Both of the two just-mentioned works are restricted to environments with two products. To the best of our knowledge, Liu et al. (2022) is the lone other paper that considers a choice-based multi-product, multi-period revenue management problem with opaque selling, however, they do not explicitly model how the contents of the opaque product influence choice, nor do they give general approximation guarantees.

1.3. Contributions

Below, we provide a brief summary of our two main algorithmic results.

The opaque-only setting (Section 2). We initially consider a setting where the platform can only offer an opaque option, and hence the decision in each period is two fold. Namely, the platform must first select the set of products that comprise the opaque option, and then must select the give-away product in the event of a purchase. In this so-called opaque-only setting, we consider a general demand model that draws inspiration from that of Elmachtoub and Hamilton (2021), who consider a mix of risk-averse and risk-neutral customers. The former group acts under the assumption that the platform will allocate their least preferred item if they select the opaque option, while the latter group assumes they will be allocated a product uniformly at random from those within the opaque option. Under this demand model, we propose a randomized policy that is guided by the optimal solution to a fluid approximation of the original problem. We show that the performance of this policy relative to the optimal expected revenue converges from $(1 - \frac{1}{e})$ to 1 as the minimum inventory level tends to infinity. We subsequently provide a simple dynamic-programming-based procedure to derandomize our policy at no cost to its performance.

The general setting (Section 3). In this section, we reintroduce the platform's ability to offer transparent products alongside the opaque option, which is required to be comprised of a subset of the transparent products. Moreover, the opaque option, if offered, must contain at least two products. In this more complex setting, we adopt a variant of the classical independent demand model in which, informally speaking, each customer arrives with the intention of purchasing a particular (transparent) product, but could be swayed to purchase the opaque option if it contains her preferred product. We develop a constant factor approximation scheme, which is based on a carefully crafted value function approximation. More specifically, we propose a policy that exploits an approximation of the optimal value functions $V_t(U_t)$ by decomposing the contribution of each product in terms of its value when sold as a transparent product and its value when used within the opaque option.

Computational experiments (Sections 4 and 5). In Section 4, we present an extensive set of computational experiments, which demonstrate the efficacy of the approximation scheme presented in Section 2 for the opaque-only setting. Generally speaking, we show that our approach performs far better than its worst-case theoretical guarantee, and for the majority of test instances, it is within 5% of optimal. In Section 5, our intention is to use the landscape of the general setting considered in Section 3 to assess the revenue gain afforded to a platform that introduces an opaque selling option. We build a suite of realistic test cases using a publicly available data set from Expedia, and benchmark our approach against a pricing policy that ignores the opaque option. Ultimately, we find that exploiting the flexibility of the opaque option through our approximation scheme leads to revenue gains of between 5-7%.

2. The Opaque-Only Setting

In this section, we consider a somewhat simplified setting in which the retailer offers (and prices) only an opaque product in each period. From a technical standpoint, this initial setting still possesses the defining operational trade-off that is the backbone of this work; namely, even in the absence of transparent products, the challenge of optimally utilizing the opaque selling option as a lever to effectively control/balance inventory levels remains a non-trivial task. Moreover, there are indeed multiple platforms that offer an opaque-only selling option. For example, as in seen in Figure 2, Priceline's Pricebreaker option offers three similar hotels (in terms of location, free internet, star rating, etc ...) within an opaque product at a discount price. Consequently, the opaque-only setting not only serves as an instrumental technical warm-up before tackling the more general setting with both transparent and opaque products, but it also has its roots in important practical applications.

The updated dynamic program. In an effort to formalize the opaque-only setting, we present an updated version of the dynamic program given in (1), which reflects the notion that the retailer can now only offer and price an opaque product. For this purpose, we use $\pi(\ell, S_{\mathcal{O}})$ to denote the choice probability for an opaque product consisting of products $S_{\mathcal{O}}$ priced at level ℓ, whose exact structure is formalized shortly. With this notation in-hand, the value functions given in (1) simplify to

$$V_t(U_t) = \max_{\substack{S_{\mathcal{O}} \subseteq \mathcal{N}(U_t), \\ \ell \in [L]}} \left\{ \pi\left(\ell, S_{\mathcal{O}}\right) \cdot \left(p_\ell - \min_{k \in S_{\mathcal{O}}} \Delta V_{t+1}^k\left(U_t\right) \right) \right\} + V_{t+1}(U_t), \tag{2}$$

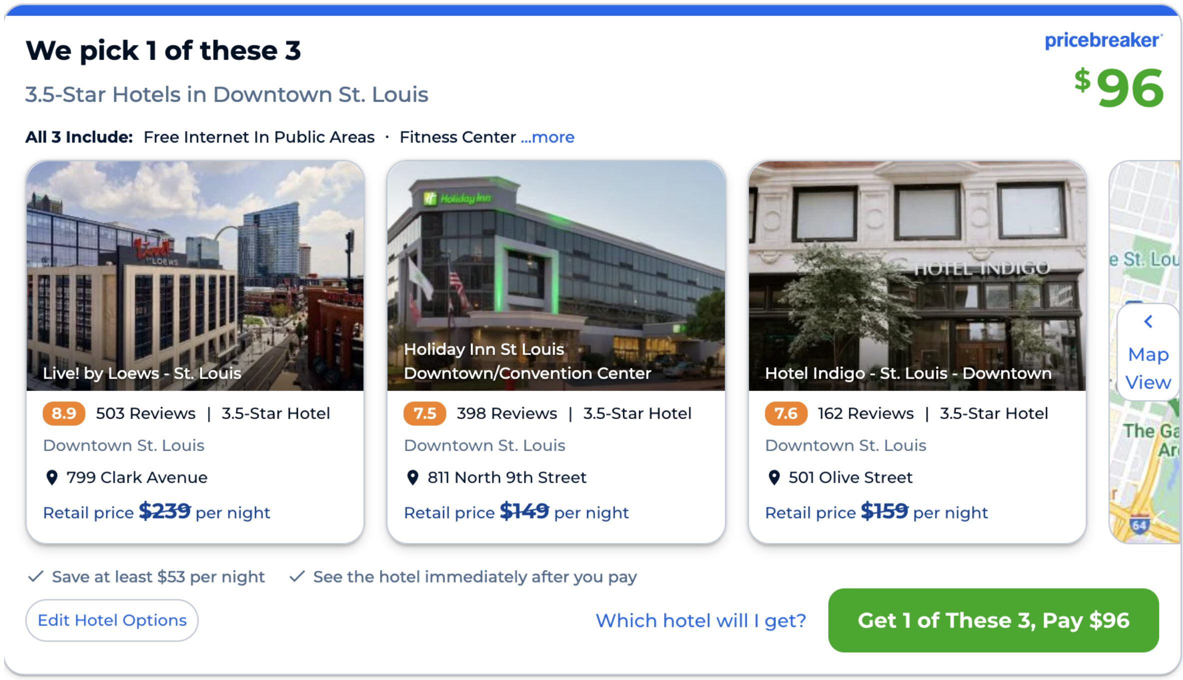

Figure 2 **Opaque-only selling option at Priceline**

with base cases $V_{T+1}(\cdot) = 0$. The recursion in (2) reflects the notion that in the opaque-only setting, the retailer's decision in each period concerns selecting the contents of the opaque product and its price. It is important to note that computing the optimal policy via backward induction remains intractable due to the curse of dimensionality, i.e. the size of the state space grows exponentially in n.

The demand model. Following almost exactly the demand framework of the seminal work of Elmachtoub and Hamilton (2021), we assume that the customer population consists of a mix of risk-averse (RA) and risk-neutral (RN) customers. The risk-averse customers associate a preference weight of $w_\ell^{\text{RA}}(S_{\mathcal{O}}) = \min_{i \in S_{\mathcal{O}}} w_{i\ell}$ with opaque product $S_{\mathcal{O}}$ priced at level ℓ, while the risk-neutral customers associate a weight of $w_\ell^{\text{RN}}(S_{\mathcal{O}}) = \frac{\sum_{i \in S_{\mathcal{O}}} w_{i\ell}}{|S_{\mathcal{O}}|}$. Informally speaking, risk-averse customers assume they will be allocated their least preferred product, while risk-neutral customers are more optimistic, and essentially operate under the assumption that they will be allocated one of the $|S_{\mathcal{O}}|$ products uniformly at random. In sum, letting α denote the fraction of the population that is risk-averse, the choice probability for opaque product $S_{\mathcal{O}}$ priced at level ℓ is

$$\pi(\ell, S_{\mathcal{O}}) = \alpha \cdot \frac{w_\ell^{\text{RA}}(S_{\mathcal{O}})}{1 + w_\ell^{\text{RA}}(S_{\mathcal{O}})} + (1-\alpha) \cdot \frac{w_\ell^{\text{RN}}(S_{\mathcal{O}})}{1 + w_\ell^{\text{RN}}(S_{\mathcal{O}})}.$$

It is important to note that, for risk-neutral customers, it is quite possible that the weight of the opaque product will decrease if a low weight product is added to the opaque option. Consequently, our demand model breaks from the traditional random-utility-maximization-based choice models in that removing a product from the offered assortment can actually decrease the appeal of the no-purchase option. The opposite is true under any random-utility-maximization-based choice model, [2] where the removal of an offered product always raises the chances the no-purchase option will be selected.

Main theorem. Having fully formalized the details of the opaque-only setting, we state our main theorem, and devote the remainder of the section to the details of its proof. For the remainder of the paper, we use $\mathrm{OPT} = V_1(U_1)$ to denote the optimal revenue for both the opaque-only and general setting.

THEOREM 1. *There exists a deterministic policy that garners an expected revenue of at least* $\mathrm{APX}(u_{\min}) \cdot \mathrm{OPT}$*, where* $u_{\min} = \min_{i \in \mathcal{N}} u_{1i}$ *is the smallest starting inventory level across all products and*

$$\mathrm{APX}(u) = 1 - \frac{u^u}{u!} \cdot e^{-u}.$$

Before beginning the exposition of our proof of the above theorem, we first note that $\mathrm{APX}(u)$ increasing from $1 - \frac{1}{e}$ to 1 as u increases from 1 to ∞. Consequently, our algorithm performs near-optimally as long as the starting inventory levels of all products are reasonably large.

2.1. Technical overview

In what follows, we provide a summary of the three main steps required to prove Theorem 1.

Step 1: The fluid approximation (Section 2.2). In this first step, we present a deterministic fluid approximation of our problem, whose optimal objective value is well-known to provide an upper bound on OPT. We formulate this fluid approximation as a linear program with decision variable $h(\ell, S_{\mathcal{O}}, i)$ for each price level $\ell \in [L]$, opaque product $S_{\mathcal{O}} \subseteq \mathcal{N}$, and product $i \in S_{\mathcal{O}}$, that represents the (potentially fractional) number of periods in which we offer opaque product $S_{\mathcal{O}}$ at price p_ℓ, and if it is selected, we give away product i. As such, it is straightforward to see that this linear program has $O(2^n L)$ decision variables, and hence cannot likely be readily solved if n is large. Our main result of this first step shows that for each product $i \in \mathcal{N}$ and each price level $\ell \in [L]$, there exists a lone opaque product $S_{\mathcal{O}} \in \{S \subseteq \mathcal{N} : i \in S\}$ such that $h(\ell, S_{\mathcal{O}}, i) > 0$ at optimality. As a result, the originally exponentially-sized fluid linear program can be immediately recast as an equivalent linear program with nL decision variables. Moreover, we show that the optimal basis for this reduced fluid linear program consists of a decision variable associated with a single price level for all products, except for one product, which might be represented at two price levels.

[2] The majority of popular choice models, such as MNL, nested logit, Markov chain and mixed-MNL, fall under this framework

Step 2: The randomized policy (Section 2.3). In this step, we analyze the randomized policy that results from directly following the recommendation of the fluid approximation. That is, if the optimal decision variables are given by $\{h^*(\ell, S_{\mathcal{O}}, i) : \ell \in [L], S_{\mathcal{O}} \subseteq \mathcal{N}, i \in S_{\mathcal{O}}\}$, then we simply offer product $S_{\mathcal{O}}$ at price p_ℓ, giving away product i if selected, for $h^*(\ell, S_{\mathcal{O}}, i)$ periods. Our analysis proceeds by comparing the expected revenue earned by our randomized policy to the expected revenue garnered by the fluid approximation across all assortments wherein the same product is prescribed as the focal give-away product, i.e. the product allocated to the customer if a purchase is made. We show that for any particular focal give-away product, these two expected revenues are within a factor of $\mathrm{APX}(u_{\min})$, which directly implies the performance bound stated in Theorem 1 related to the overall earnings of our policy.

Step 3: Derandomization (Section 2.4). In this final step, we provide a simple means to fully derandomize our policy, while strictly improving its revenue performance. In the process, we argue that our derandomized policy is indeed feasible, i.e. we are never at risk of offering a product that is stocked out within the opaque option.

2.2. The fluid approximation

In this section, we present and analyze a relaxed version of problem (2) that is best viewed as a fluid approximation in which demand takes on its expected value. We note that it is revenue management folklore (Liu and Van Ryzin 2008, Gallego et al. 2015b) that such fluid approximations yield an upper bound on OPT. Following this just-mentioned stream of literature, we formulate this deterministic fluid approximation as a linear program with an exponential number of decision variables. We then show that we can identify a subset of only $O(nL)$ decision variables that are the sole candidates to be part of an optimal basis, which gives way to a reduced linear programming formulation of our fluid approximation. Finally, we conclude the section by showing additional structure on this optimal basis that is critical in the analysis of the policy we put forth in Section 2.3.

The fluid linear program. The decision variables $\{h(\ell, S_{\mathcal{O}}, i) : \ell \in [L], S_{\mathcal{O}} \subseteq \mathcal{N}, i \in S_{\mathcal{O}}\}$ denote the number of periods in which opaque option $S_{\mathcal{O}}$ is offered at price level ℓ, and if selected, product i will be allocated. Our linear program of interest is given below:

$$\begin{aligned}
\mathrm{OPT}_{\text{fluid}} = \max \quad & \sum_{S_{\mathcal{O}} \subseteq \mathcal{N}} \sum_{i \in S_{\mathcal{O}}} \sum_{\ell \in [L]} p_\ell \pi(\ell, S_{\mathcal{O}}) h(\ell, S_{\mathcal{O}}, i) \\
\text{s.t.} \quad & (1) \quad \sum_{S_{\mathcal{O}} : i \in S_{\mathcal{O}}} \sum_{\ell \in [L]} \pi(\ell, S_{\mathcal{O}}) h(\ell, S_{\mathcal{O}}, i) \leq u_{i1} \ \forall i \in \mathcal{N} \\
& (2) \quad \sum_{S_{\mathcal{O}} \subseteq \mathcal{N}} \sum_{i \in S_{\mathcal{O}}} \sum_{\ell \in [L]} h(\ell, S_{\mathcal{O}}, i) \leq T \\
& (3) \quad h(\ell, S_{\mathcal{O}}, i) \geq 0.
\end{aligned} \tag{OPA-LP}$$

The constraints in (1) ensure that the expected demand for each product does not exceed its initial capacity, while the constraint in (2) encodes the notion that we only have T periods in the selling horizon.

Structural result - the best opaque option. Recalling that products are indexed in increasing order of weight, for each product $i \in \mathcal{N}$ and price level $\ell \in [L]$ pair, let

$$S_{\mathcal{O}}(\ell, i) = \arg\max_{k \in [i,n]} \pi(\ell, \{i\} \cup [k, n])$$

be the opaque product with the largest choice probability among all those that includes i and some subset of the highest weight products with index larger than i at price level ℓ [3]. The following claim reveals that, for any $i \in \mathcal{N}$ and $\ell \in [L]$, the opaque product that maximizes the choice probability for this pair is in fact $S_{\mathcal{O}}(\ell, i)$.

Claim 1. *For any product $i \in \mathcal{N}$ and price level $\ell \in [L]$, we have*

$$S_{\mathcal{O}}(\ell, i) = \arg\max_{S \subseteq \mathcal{N}: i \in S} \pi(\ell, S). \tag{3}$$

As an important side-note, observe that we can easily recover $\{S_{\mathcal{O}}(\ell, i)\}_{\ell \in [L], i \in \mathcal{N}}$ by enumerating over the $O(n)$ options for $S_{\mathcal{O}}(\ell, i)$ to solve (3) for each product $i \in \mathcal{N}$ and price level $\ell \in [L]$ pair, of which there are at most $O(nL)$.

Structural result - the optimal basis. The next lemma reveals that there exists an optimal solution to OPA-LP in which $h(\ell, S_{\mathcal{O}}, i) > 0$ only if $S_{\mathcal{O}} = S_{\mathcal{O}}(\ell, i)$. In other words, for focal give-away product i and price level ℓ, the optimal objective value of OPA-LP remains unchanged if we remove all decision variables related to this focal product-price level pair in which $S_{\mathcal{O}} \neq S_{\mathcal{O}}(\ell, i)$.

Lemma 1. *There exists an optimal solution $\{h^*(\ell, S_{\mathcal{O}}, i) : \ell \in [L], S_{\mathcal{O}} \subseteq \mathcal{N}, i \in S_{\mathcal{O}}\}$ to OPA-LP such that $h^*(\ell, S_{\mathcal{O}}, i) = 0$ if $S_{\mathcal{O}} \neq S_{\mathcal{O}}(\ell, i)$.*

The reduced linear program. Exploiting Lemma 1, the originally exponentially-sized OPA-LP can be reformulated as the following reduced linear program that includes only the $O(nL)$ decision variables $\{h(\ell, S_{\mathcal{O}}(\ell, i)\}_{\ell \in [L], i \in \mathcal{N}}$.

$$\begin{aligned}
\mathrm{OPT}_{\text{fluid}} = \max \ & \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} p_\ell \pi(\ell, S_{\mathcal{O}}(\ell, i)) h(\ell, S_{\mathcal{O}}(\ell, i), i) \\
\text{s.t.} \quad & (1) \quad \sum_{\ell \in [L]} \pi(\ell, S_{\mathcal{O}}(\ell, i)) h(\ell, S_{\mathcal{O}}(\ell, i), i) \le u_{i1} \ \forall i \in \mathcal{N} \\
& (2) \quad \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} h(\ell, S_{\mathcal{O}}(\ell, i), i) \le T \\
& (3) \quad h(\ell, S_{\mathcal{O}}(\ell, i), i) \ge 0.
\end{aligned} \tag{Reduced-OPA-LP}$$

[3] For positive integers i, j such that $i < j$, we use $[i, j] = \{i, i+1, \ldots, j\}$.

We conclude this section by showing that in the optimal solution to Reduced-OPA-LP, there is at most a single product $i \in \mathcal{N}$ for which there exists two distinct price levels $\ell, \ell' \in [L]$ such that $h(\ell, S_{\mathcal{O}}(\ell, i), i), h(\ell', S_{\mathcal{O}}(\ell', i), i) > 0$, while all other products have at most a single price level in the optimal basis.

Claim 2. *Let $\{h^*(\ell, S_{\mathcal{O}}(\ell, i), i)\}_{\ell \in [L], i \in \mathcal{N}}$ denote the optimal solution to Reduced-OPA-LP, and let $\mathcal{L}_i = \{\ell \in [L] : h^*(\ell, S_{\mathcal{O}}(\ell, i), i) > 0\}$. There exists at most a single product $i \in \mathcal{N}$ such that $|\mathcal{L}_i| > 1$, and if such a set $\mathcal{L}_i$ exists, we have that $|\mathcal{L}_i| = 2$.*

2.3. The randomized policy

In this section, we present our randomized opaque pricing policy, which is derived from the optimal solution to Reduced-OPA-LP. After describing our policy, we prove that it achieves the performance guarantee stated in Theorem 1.

The Reduced-OPA-LP-based policy. Re-hashing the definitions of Claim 2, we let $\{h^*(\ell, S_{\mathcal{O}}(\ell, i), i)\}_{\ell \in [L], i \in \mathcal{N}}$ denote the optimal solution to Reduced-OPA-LP, and $\mathcal{L}_i = \{\ell \in [L] : h^*(\ell, S_{\mathcal{O}}(\ell, i), i) > 0\}$ denote the price levels used when product i is the focal give-away product. Our proposed policy is quite simple: we sequentially proceed over the sets $\{\mathcal{L}_i\}_{i \in \mathcal{N}}$ in increasing order of i, and for each $\ell \in \mathcal{L}_i$, we offer opaque product $S_{\mathcal{O}}(\ell, i)$ at price level ℓ, allocating product i if selected, for $h^*(\ell, S_{\mathcal{O}}(\ell, i), i)$ periods. To formally define this procedure, let $(\ell_1, i_1), (\ell_2, i_2), \ldots, (\ell_Q, i_Q)$ denote the price level-product pairs such that for $q \in [Q]$, we have that $\ell_q \in \mathcal{L}_{i_q}$. These tuples are indexed in increasing order of product index, where we tie-break by choosing the higher of two price levels for the lone product that might appear twice, as was established in Claim 2. Algorithm 1 formalizes how we sequentially proceed over this tuple sequence, "filling" each time period with the h^*-value at-hand until it is exhausted.

The analysis. The randomized policy presented in Algorithm 1 progresses sequentially over the focal give away products in increasing order of index. As such, we can analyze the performance guarantee of our approach relative to OPT on a product-by-product basis. More specifically, for each product $i \in \mathcal{N}$ and price level $\ell \in \mathcal{L}_i$, let $D_{i\ell}$ denote the random number of sales of opaque product $S_{\mathcal{O}}(\ell, i)$. Furthermore, let $R_i = \sum_{\ell \in \mathcal{L}_i} p_\ell \cdot \mathbb{E}[D_{i\ell}]$ denote the expected revenue earned by our randomized policy when product i is the focal give-away product. We let $\mathrm{ALG} = \sum_{i \in \mathcal{N}} R_i$ represent the total expected revenue earned by our policy. Similarly, define $\mathrm{OPT}_i = \sum_{\ell \in \mathcal{L}_i} p_\ell \pi(\ell, S_{\mathcal{O}}(\ell, i)) h^*(\ell, S_{\mathcal{O}}(\ell, i), i)$ to be product i's contribution as the focal give-away product to $\mathrm{OPT}_{\text{fluid}}$. The following lemma, whose proof can be found in Appendix A, relates R_i to OPT_i.

Lemma 2. *For any $i \in \mathcal{N}$ such that $\mathcal{L}_i \neq \emptyset$, we have that*

$$\frac{R_i}{\mathrm{OPT}_i} \geq \mathrm{APX}(u_{i1}).$$

Algorithm 1 Randomized Pricing Policy

Initialization: $T_0 = 0$

for $q \in [Q]$ **do**

$$T_q = \sum_{q' \in [q]} h^*(\ell_{q'}, S_{\mathcal{O}}(\ell_{q'}, i_{q'}), i_{q'})$$

end for

for $t \in [T]$ **do**

Offer the opaque product $S_{\mathcal{O}}(\ell_q, i_q)$ at price ℓ_q, giving away product i_q w.p.

$$P_{qt} = \max\{([t-1,t] \cap [T_{q-1}, T_q]) \cup \{0\}\} - \min\{([t-1,t] \cap [T_{q-1}, T_q]) \cup \{0\}\}$$

Note: Above, we use $[x, y]$ to denote the continuous interval from x to y.

end for

Via Lemma 2 and the fact that $\text{OPT}_{\text{fluid}} \geq \text{OPT}$, we get that

$$\frac{\text{ALG}}{\text{OPT}} \geq \min_{i \in \mathcal{N}: \mathcal{L}_i \neq \emptyset} \frac{R_i}{\text{OPT}_i} = \text{APX}(u_{\min}),$$

which is precisely the guarantee stated in Theorem 1.

2.4. Derandomization

In this section, we describe our procedure for derandomizing Algorithm 1, which results in a fully deterministic state-dependent policy that can be computed efficiently, and is guaranteed to garner a larger expected revenue than its randomized counterpart. We start by introducing a means to express the expected revenue earned by Algorithm 1 through a dynamic program. We then leverage this dynamic program to reveal a simple way to derandomize our approach, while strictly improving its revenue performance. Ultimately, our approach is a modified version of the classical derandomization method of conditional expectations

Computing **ALG** *recursively.* To begin, we establish the following intermediate claim, which acts as a stepping stone to the dynamic program we present to compute the expected revenue earned by Algorithm 1. The proof of this claim is presented in Appendix A.

Claim 3. *Under Algorithm 1, if product i is the focal give-away product in period t, then*

(i) For each $q \in [Q]$ such that $i_q < i$, and any $\tau > t$, we have $P_{q\tau} = 0$, i.e. any product indexed lower than i will never be included within the opaque product offered in periods $t+1, \ldots, T$.

(ii) For each $q \in [Q]$ such that $i_q > i$, we have $u_{i_q 1}$ remaining inventory units of product i_q at the start of period $t+1$, i.e. product i_q has not been consumed over periods $1, \ldots, t$.

Moving forward, for ease of notation, we use $\pi_q = \pi(\ell_q, S_{\mathcal{O}}(\ell_{i_q}, i_q)$ as a shorthand for the choice probability of opaque product $S_{\mathcal{O}}(\ell_{i_q}, i_q)$ priced at level ℓ_{i_q}. With Claim 3 in-mind, we define $\mathcal{R}_t(i,u)$ to be the expected revenue earned by Algorithm 1 over periods $t,\dots,T$, when product i was the last product that was the focal give-away product if a purchase was made, and at the start if period t, product i has u units of remaining inventory. This function can be expressed recursively as

$$\begin{aligned}\mathcal{R}_t(i,u) = \mathbb{1}_{u>0} \cdot \left(\sum_{\substack{q\in[Q]:\\ i_q=i}} P_{qt} \cdot \left(\pi_q \cdot \left(p_{\ell_q} + \mathcal{R}_{t+1}(i,u-1) - \mathcal{R}_{t+1}(i,u)\right) + \mathcal{R}_{t+1}(i,u)\right)\right) \\ + \sum_{\substack{q\in[Q]:\\ i_q>i}} P_{qt} \cdot \left(\pi_q \cdot \left(p_{\ell_q} + \mathcal{R}_{t+1}(i_q,u_{i_q1}-1) - \mathcal{R}_{t+1}(i_q,u_{i_q1})\right) + \mathcal{R}_{t+1}(i_q,u_{i_q1})\right), \qquad (4)\end{aligned}$$

with base cases of $\mathcal{R}_{T+1}(\cdot,\cdot)=0$. In the above recursion, the first term corresponds to the case in which the randomization results in an opaque product where i is the focal give-away product, and the second term corresponds to the case in which we offer an opaque product with focal give-away product $i_q > i$. This latter case only needs to consider $q\in[Q]$ such that $i_q > i$, since by property (i) of Claim 3, we have that $P_{qt}=0$ for any other $q\in[Q]$ with index no larger than i over the remaining periods $t+1,\dots,T$. Additionally, by property (ii) of Claim 3, we know that any product $i_q > i$ must not have been consumed over the first $t-1$ periods, and hence its remaining inventory in period t is u_{i_q1}. All-in-all, we see that $\text{ALG} = \mathcal{R}_1(\cdot, u_{i_11})$, where we set $i=\cdot$ to reflect the fact that this parameter is undefined in the first period.

Derandomization. Our derandomization procedure is based on the following alternative dynamic program, in which the random offer/pricing decision in each period captured by the probabilities $\{P_{qt}\}_{q\in[Q],t\in[T]}$, is replaced by a deterministic decision that is guaranteed to represent an improvement revenue-wise over the randomized policy. Specifically, define

$$\begin{aligned}\hat{\mathcal{R}}_t(i,u) = \max\Bigg\{ \mathbb{1}_{u>0} \cdot \left(\max_{\substack{q\in[Q]:\\ P_{qt}>0, i_q=i}} \left\{ \pi_q \cdot \left(p_{\ell_q} + \hat{\mathcal{R}}_{t+1}(i,u-1) - \hat{\mathcal{R}}_{t+1}(i,u)\right) + \hat{\mathcal{R}}_{t+1}(i,u)\right\}\right), \\ \max_{\substack{q\in[Q]:\\ P_{qt}>0, i_q>i}} \left\{ \pi_q \cdot \left(p_{\ell_q} + \hat{\mathcal{R}}_{t+1}(i_q,u_{i_q1}-1) - \hat{\mathcal{R}}_{t+1}(i_q,u_{i_q1})\right) + \hat{\mathcal{R}}_{t+1}(i_q,u_{i_q1})\right\}\Bigg\}, \qquad (5)\end{aligned}$$

with base cases of $\hat{\mathcal{R}}_{T+1}(\cdot,\cdot)=0$. We go from $\mathcal{R}_t(i,u)$ to $\hat{\mathcal{R}}_t(i,u)$ by simply choosing the opaque option that leads to the largest future expected revenue among all those with positive offer probabilities. As such, it is clear that $\hat{\mathcal{R}}_t(i,u) \geq \mathcal{R}_t(i,u)$, and hence the policy derived from following this dynamic program from the initial state of $(\cdot, u_{i_1})$ in period 1 is guaranteed to improve upon the revenue performance of Algorithm 1. In the event that $u=0$ and and the outer maximization prescribes continuing to offer product i as the give-away product, then we earn no revenue in the current period because of the indicator $\mathbb{1}_u > 0$. In this case, to avoid a scenario in which we offer a product within the opaque option that is stocked out, we can simply offer nothing, which clearly results in the same revenue accumulation.

The running time of derandomization. The overall running time of our derandomization procedure is the precisely the running time needed to compute the value functions in (5). In total, there $O(Tn \cdot \max_{i\in[n]} u_{i1})$ such value functions, each of which can be computed in running time of $O(Q) = O(nL)$ time by exhaustive enumeration over all $q \in [Q]$ to solve the inner maximization problems.

3. The General Setting

In this section, we revert back to the original setting in which there are both transparent and opaque products. When we move to this general setting, the stochastic element of the demand becomes far more difficult to tame, since now demand can be realized through both the transparent and the opaque offering. In short, in the absence of transparent products, we used the highly attractive products to continuously boost the appeal of the opaque option while never being depleted, until it became their turn to become the focal give-away product. However, when transparent products are reintroduced as a selling option, such a policy might be rendered ineffective due to the fact that any product offered within the opaque option must also be offered as a transparent option. As such, the use of highly attractive products as "decoys" within the opaque option is rendered less effective due to the sales of such products through the tranparent channel. Nonetheless, under a variant of the classical independent demand model (Talluri and Van Ryzin 1999, Adelman 2007, Topaloglu 2009) described next, we are able to overcome these technical hurdles to derive a policy that garners a constant fraction of the optimal expected revenue.

The opaque-specific (OS) independent demand model The distinguishing feature of the standard independent demand model is that the demand for each product is assortment-independent, i.e. each customer arrives with the intent of purchasing a particular product, and will do so if her product of interest is made available for purchase. We consider a variant of this model in which customers only substitute between their product of interest, and an opaque option that contains this focal product. Formally, in period t, the arriving customer is interested in product $i \in \mathcal{N}$ with probability λ_{it}. We refer to these customers as type-i customers for the remainder of this section. A type-i customer associates an MNL-based weight of $w_{i\ell}$ with the transparent version of product i priced at price level ℓ, and an MNL-based weight of $w^{\mathcal{O}}_{i\ell}(S_{\mathcal{O}}) = \mathbb{1}_{i\in S_{\mathcal{O}}} \cdot f_i(|S_{\mathcal{O}}|, \ell)$ with the opaque option, where $f_i(|S_{\mathcal{O}}|, \ell)$ is an arbitrary function that is non-increasing in price and the cardinality of $S_{\mathcal{O}}$. As such, an arriving type-i customer will only consider the opaque product if it includes her product of interest. Moreover, the attractiveness of this opaque option is influenced by its price, and the number of additional products included alongside product i. AS a result, under assortment S of transparent products and opaque product $S_{\mathcal{O}}$, along with pricing decisions $x \in \mathcal{X}(S, S_{\mathcal{O}})$, we denote the choice probability of transparent i conditioned on the arrival of a type-i customer as

$$\pi_i(x, S_{\mathcal{O}}) = \frac{w_i(x)}{1 + w_j(x) + w_i^{\mathcal{O}}(x, S_{\mathcal{O}})},$$

and the choice probability of the opaque option as

$$\pi_i^{\mathcal{O}}(x, S_{\mathcal{O}}) = \frac{w_i^{\mathcal{O}}(x, S_{\mathcal{O}})}{1 + w_i(x) + w_i^{\mathcal{O}}(x, S_{\mathcal{O}})}.$$

In the above expressions, we use $w_i(x) = \sum_{\ell \in [L]} w_{i\ell} x_{i\ell}$ and $w_i^{\mathcal{O}}(x, S_{\mathcal{O}}) = \sum_{\ell \in [L]} w_{i\ell}^{\mathcal{O}}(S_{\mathcal{O}}) x_{\mathcal{O}\ell}$ to denote the weights of product i and the opaque product under pricing decision x. Finally, we make the following natural assumption, which states that a type-i customers associate a higher weight with transparent product i when its price matches that of the opaque option.

Assumption 2. *For any pricing vectors $x, x' \in \mathcal{X}(S, S_{\mathcal{O}})$ such that $p_i(x) = p_{\mathcal{O}}(x')$, we hav that $w_i(x) \geq w_i^{\mathcal{O}}(x', S_{\mathcal{O}})$.*

Remark. The OS independent demand model is indeed a simplified version of the general demand model conceived in Section 1, since (i) each customer substitutes only between their product of interest and the opaque option, and (ii) the weight of the opaque option for each customer type is influenced only by the cardinality of $S_{\mathcal{O}}$, and not by its specific contents. That said, as alluded to above, the OS independent demand model represents one natural translation of the classical independent demand model into an opaque selling environment, and perhaps more importantly, we are still left with a highly non-trivial dynamic pricing problem, even under this seemingly simple demand model.

The updated dynamic program. Under the above-described opaque-specific independent demand model, the updated recursion for the general problem given in (1) can be written as

$$\begin{aligned} V_t(U_t) = \max_{\substack{S \subseteq \mathcal{N}(U_t), \\ S_{\mathcal{O}} \subseteq S}} \max_{x \in \mathcal{X}(S, S_{\mathcal{O}})} \Bigg\{ \sum_{i \in S} \lambda_{it} \cdot \Bigg(& \pi_i(x, S_{\mathcal{O}}) \cdot \Big(p_i(x) - \Delta V_{t+1}^i(U_t) \Big) + \\ & \pi_i^{\mathcal{O}}(x, S_{\mathcal{O}}) \cdot \Big(p_{\mathcal{O}}(x) - \min_{k \in S_{\mathcal{O}}} \Delta V_{t+1}^k(U_t) \Big) \Bigg) \Bigg\} + V_{t+1}(U_t), \end{aligned} \tag{6}$$

to reflect the arrival probabilities of the n different customer types.

Main theorem. Our main result of this section is a constant factor approximation scheme that builds on the approximate dynamic programming ideas presented in Ma et al. (2020) for the network revenue management problem. The remainder of this section is devoted to presenting our approach and proving that it matches the guarantee stated in the theorem below.

Theorem 2. *There exists a deterministic policy that garners an expected revenue of at least $\frac{1}{8}$OPT.*

3.1. Technical overview

In what follows, we summarize the three steps that go into proving Theorem 2

Step 1: The value function approximation (Section 3.2). In this first step, we propose alternative value functions $H_t(U_t)$ that serve as approximations to $V_t(U_t)$, as seen in (6). These approximate value functions possess the following two key properties. First, they can be efficiently computed for any time period t and inventory vector U_t. Second, their estimate of the marginal value of a unit of each product can be upper bounded by a term that is independent of current inventory levels, which will prove critical throughout the latter two steps.

Step 2: The roll-out policy (Section 3.3). In this second step, we describe and analyze the roll-out policy that results from plugging in the approximate value functions into the recursion in (6). Via an inductive argument, we show that the expected revenue garnered by this policy is at least $\frac{1}{2}H_1(U_1)$.

Step 3: The upper bound (Section 3.4). In this final step, we show that $V_1(U_1) = \text{OPT} \le 4H_1(U_1)$, which directly yield the performance guarantee stated in Theorem 2, since in step 2, we propose a policy that garners an expected revenue of at least $\frac{1}{2}H_1(U_1)$. We establish the aforementioned upper bound by considering the dual of the deterministic fluid approximation of our problem. Since the primal optimal solution upper bounds OPT (and it is a maximization), we know that any dual feasible solution is also an upper bound by weak duality. With this insight in-mind, we construct a dual feasible solution whose objective is no more than $4H_1(U_1)$.

3.2. The value function approximation

We approximate $V_t(U_t)$ using the following quadratic function of the period-t inventory levels:

$$H_t(U_t) = \sum_{i\in\mathcal{N}} \frac{u_{it}}{u_{i1}} \cdot (\gamma_i^t + \gamma_{ii}^t) + \sum_{i\in\mathcal{N}} \sum_{j\neq i} \frac{u_{it}}{u_{i1}} \frac{u_{jt}}{u_{j1}} \gamma_{ij}^t, \tag{7}$$

and so $H_1(U_1) = \sum_{i\in\mathcal{N}} (\gamma_i^t + \gamma_{ii}^t) + \sum_{i\in\mathcal{N}} \sum_{j\neq i} \gamma_{ij}^t$. We will shortly describe how to compute the quantities $\{\gamma_i^t\}_{i\in\mathcal{N},t\in[T]}$ and $\{\gamma_{ij}^t\}_{i,j\in\mathcal{N},t\in[T]}$, which we henceforth refer to as "tuning parameters". Before doing so, however, we first provide loose intuition on their respective interpretations. Namely, the tuning parameter γ_i^t can be interpreted as an estimate of the expected revenue of product i over periods $t, \dots, T$ coming from its sale as a transparent product. On the hand, γ_{ij}^t is an estimate of the expected revenue contribution from product j sold through the opaque product to a type-i customer over periods $t, \dots, T$. Next, we define

$$\theta_i^t = \gamma_i^t + \gamma_{ii}^t + \sum_{j\neq i} (\gamma_{ij}^t + \gamma_{ji}^t),$$

to be an estimate of the total expected revenue accrued from product i via sales of both transparent and opaque products over periods $t, \dots, T$. Here, we include the term $\sum_{j\neq i} \gamma_{ij}^t$, since the sale of product j through the opaque product to a type-i customer is not possible if i is stocked out. In the case of such a purchase event, revenue contribution should be attributed to both product j (since it was consumed) and product i (since it influenced the customer to select the opaque option). Noting that

$$H_t(U_t) - H_t(U_t - e_i) = \frac{1}{u_{i1}} \cdot \left(\gamma_i^t + \gamma_{ii}^t + \sum_{j \neq i} \frac{u_{jt}}{u_{j1}} \cdot (\gamma_{ij}^t + \gamma_{ji}^t) \right) \leq \frac{\theta_i^t}{u_{i1}}, \tag{8}$$

we see that $\frac{\theta_i^t}{u_{i1}}$ is a valid inventory-level-indepedent upper bound on the marginal value of single unit of product i.

The assortment sub-problem. Before delving into our approach for computing the tuning parameters, it useful to introduce the following period-t sub-problem, which plays a central role in our recursive approach for setting the tuning parameters. For an arbitrary vector of marginal value estimates $\theta = (\theta_1, \ldots, \theta_n)$, let

$$Z_t(\theta) = \max_{\substack{S \subseteq [n], \\ S_{\mathcal{O}} \subseteq S, q \in S_{\mathcal{O}}, \\ x \in \mathcal{X}(S, S_{\mathcal{O}})}} \sum_{i \in S} \lambda_{it} \cdot \left(\pi_i(x, S_{\mathcal{O}}) \cdot \left(p_i(x) - \frac{\theta_i}{u_{i1}} \right) + \pi_{\mathcal{O}}^i(x, S_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(x) - \frac{\theta_q}{u_{q1}} \right) \right).$$

Additionally, define

$$R_i(S, S_{\mathcal{O}}, q, x, \theta) = \pi_i(x, S_{\mathcal{O}}) \cdot \left(p_i(x) - \frac{\theta_i}{u_{i1}} \right) + \pi_{\mathcal{O}}^i(x, S_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(x) - \frac{\theta_q}{u_{q1}} \right) \tag{9}$$

to denote the product-i contribution, and so

$$Z_t(\theta) = \max_{\substack{S \subseteq [n], \\ S_{\mathcal{O}} \subseteq S, q \in S_{\mathcal{O}}, \\ x \in \mathcal{X}(S, S_{\mathcal{O}})}} \sum_{i \in S} \lambda_{it} R_i(S, S_{\mathcal{O}}, q, x, \theta). \tag{SUB-ASSORT}$$

The following lemma reveals the existence of a $\frac{1}{2}$-optimal solution to SUB-ASSORT in which the contribution of each product is non-negative. This approximate solution is a critical feature of our approach for setting the tuning parameters.

LEMMA 3. *For arbitrary period $t \in [T]$ and $\theta \in \mathbb{R}^n$, there exists a solution $(\hat{S}^t, \hat{S}_{\mathcal{O}}^t, \hat{q}^t, \hat{x}^t)$ to SUB-ASSORT that can be computed in polynomial time, and that satisfies*

(i) $R_i(\hat{S}^t, \hat{S}_{\mathcal{O}}^t, \hat{q}^t, \hat{x}^t, \theta) \geq 0$ *for each* $i \in \hat{S}^t$, *and*

(ii) $\sum_{i \in \hat{S}_t} \lambda_{it} R_i(\hat{S}^t, \hat{S}_{\mathcal{O}}^t, \hat{q}^t, \hat{x}^t, \theta) \geq \frac{1}{2} \cdot Z_t(\theta)$.

Algorithm 2 Computing the tuning parameters

Initialization: $\theta_i^{T+1} = \gamma_i^{T+1} = \gamma_{ii}^{T+1} = 0$ and $\gamma_{ij}^{T+1} = \gamma_{ji}^{T+1} = 0$

for $t \in T, T-1, \ldots, 1$ **do**

Compute $(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t)$ as described in Lemma 3 with $\theta = (\theta_1^{t+1}, \ldots, \theta_n^{t+1})$.

Compute period-t tuning parameters

$$\gamma_i^t = \lambda_{it} \pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) + \gamma_i^{t+1}$$

$$\gamma_{ij}^t = \mathbb{1}_{j=\hat{q}^t} \cdot \left(\lambda_{it} \pi_i^{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_j^{t+1}}{u_{j1}} \right) \right) + \gamma_{ij}^{t+1}$$

$$\theta_i^t = \gamma_i^t + \gamma_{ii}^t + \sum_{j \neq i} (\gamma_{ij}^t + \gamma_{ji}^t)$$

end for

The tuning parameters Our procedure for computing the tuning parameters is presented in Algorithm 2. In essence, the period-t tuning parameters are computed under the assumption that we follow the decision vector $(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t)$, and that the marginal value of any unit of product i can be approximated as $\frac{\theta^{t+1}}{u_{i1}}$. This high-level intuition is formalized in Section 3.3, where we establish the efficacy pf a roll-out policy based on this value function approximation.

We conclude this section by showing a critical characteristic of the tuning parameters, which is essential to establish the upper bound in Step 3 of our approach. Namely, we establish that the θ_i^t are decreasing in t, as is formalized in the following lemma. The proof of this Lemma reveals the significance of property (i) in Lemma 3.

LEMMA 4. *For any $i \in \mathcal{N}$ and $t \in [T]$, we have $\theta_i^t \geq \theta_i^{t+1} \geq 0$.*

Interestingly, in previous works that utilize a tuning-parameter-based approximation (Ma et al. 2020, Rusmevichientong et al. 2020), a similar property as stated in the above lemma falls naturally out of their respective assortment sub-problems, akin to SUB-ASSORT. At a high-level, the reason for this is that in these previous works, each unit of a particular resource contributes to its value only when it is consumed. In our setting, however, a unit of a particular resource can contribute to its value by serving as a "decoy" for the opaque product; for example, a unit of product i can influence a type-i customer to purchase the opaque product, who then ultimately is allocated some product $q \neq i$, thus leaving the original unit of product i untouched (and hence readily available to serve in this decoy role over and over again). In turn, with this phenomenon in-mind, it is easy to see that $\frac{\theta_i^t}{u_{i1}} >> \max_{\ell \in [L]} p_\ell$, which is the source of the novel technical difficulties that arise in our opaque selling setting.

3.3. The roll-out policy

In this section, we formalize our roll-out policy, and show that it garners an expected revenue of at least $H_1(U_1)$. In the process, we assume that all tuning parameters have been computed via Algorithm 2.

The policy. In each period $t \in [T]$, given current inventory levels of U_t, we compute

$$(\bar{S}^t, \bar{S}^t_{\mathcal{O}}, \bar{q}^t, \bar{x}^t) = \operatorname*{arg\,max}_{\substack{S \subseteq \mathcal{N}(U_t), \\ S_{\mathcal{O}} \subseteq S, q \in S_{\mathcal{O}}, \\ x \in \mathcal{X}(S, S_{\mathcal{O}})}} \sum_{i \in S} \lambda_{it} R_i(S, S_{\mathcal{O}}, q, x, \theta^{t+1}) \tag{10}$$

which is precisely problem (SUB-ASSORT) in which the set of offered transparent products is restricted to those with non-zero inventory, and θ is set to the period $t+1$ tuning parameters. Consequently, this problem can be solved optimally via the algorithm presented in Appendix C. We then implement this vector of decisions in period t. The expected revenue garnered by this policy across periods $t, \dots, T$, given period-t inventory levels of U_t, can be expressed recursively as

$$\begin{aligned}\bar{H}_t(U_t) = \sum_{i \in \bar{S}^t} \lambda_{it} \cdot \Big(& \pi_i(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \big(p_i(\bar{x}^t) - \big(\bar{H}_{t+1}(U_t) - \bar{H}_{t+1}(U_t - e_i)\big)\big) \\ & + \pi_i^{\mathcal{O}}(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \big(p_{\mathcal{O}}(\bar{x}^t) - \big(\bar{H}_{t+1}(U_t) - \bar{H}_{t+1}(U_t - e_{\bar{q}^t})\big)\big) \Big) + \bar{H}_{t+1}(U_t),\end{aligned}$$

and hence $\bar{H}_1(U_1)$ gives the expected revenue of our policy.

The performance guarantee. The following lemma provides a general relation between $\bar{H}_t(U_t)$ and $H_t(U_t)$, which when invoked at $t = 1$ establishes that $\bar{H}_1(U_1) \geq \frac{1}{2} H_1(U_1)$.

LEMMA 5. *For any $t \in [T]$ and $U_t \leq U_1$ (component-wise), we have that $\bar{H}_t(U_t) \geq \frac{1}{2} H_t(U_t)$.*

3.4. The upper bound

Having shown in the preceding section that the roll-out policy garners an expected revenue of at least $H_1(U_1)$, we conclude our analysis of the general setting by showing that $\text{OPT} \leq 4H_1(U_1)$, which confirms the performance guaranteed reported in Theorem 2. To do so, we first introduce the deterministic fluid approximation for our general setting, whose optimal objective value upper bounds OPT. From here, we show how to construct a feasible solution to its dual that is not larger than $4H_1(U_1)$.

The fluid linear program and its dual. Akin to the fluid approximation in the opaque only setting, the decision variables $\{h_t(S, S_{\mathcal{O}}, q, x) : S \subseteq \mathcal{N}, S_{\mathcal{O}} \subseteq S, q \in S_{\mathcal{O}}, x \in \mathcal{X}(S, S_{\mathcal{O}})\}$ denote the fraction of period t in which the set of offered transparent products is S, the opaque option is $S_{\mathcal{O}}$, q will be allocated if the opaque product is selected, and the pricing decisions are given by x. Letting $\mathcal{A} = \{(S, S_{\mathcal{O}}, q, x) : S \subseteq \mathcal{N}, S_{\mathcal{O}} \subseteq S, q \in S_{\mathcal{O}}, x \in \mathcal{X}(S, S_{\mathcal{O}})\}$, our linear program of interest is given by:

$$\begin{aligned}
\text{OPT}_{\text{fluid}} = \max & \sum_{t\in[T]} \sum_{A\in\mathcal{A}} h_t(A) \cdot \sum_{i\in S} \lambda_{it} R_i(A, 0) \\
\text{s.t. } (1) \quad & \sum_{t\in[T]} \sum_{A\in\mathcal{A}} u(A) \cdot h_t(A) \leq u_{i1} && \forall i \in \mathcal{N} \\
(2) \quad & \sum_{A\in\mathcal{A}} h_t(A) \leq 1 && \forall t \in [T] \\
(3) \quad & h_t(A) \geq 0,
\end{aligned}$$

where for $A = (S, S_{\mathcal{O}}, q, x)$, the term $R_i(A, 0)$ is defined as in (9) to represent the expected revenue garnered from product i, and we introduce $u(A) = \lambda_{it}\pi_i(x, S_{\mathcal{O}}) + \sum_{j\in\mathcal{N}} \mathbb{1}_{i=q} \cdot \lambda_{jt}\pi_j^{\mathcal{O}}(x, S_{\mathcal{O}})$ to denote the expected consumption of product i. Note that constraint (1), which encodes the capacity restriction for each product, now incorporates the fact that each product can either be consumed as a transparent or opaque product. Associating dual variables of $\{\alpha_i\}_{i\in\mathcal{N}}$ with the constraints of (1), and dual variables $\{\beta\}_{t\in[T]}$ with the constraints of (2), we arrive the following dual formulation of the fluid deterministic linear problem

$$\begin{aligned}
\text{OPT}_{\text{fluid}} = \min & \sum_{i\in\mathcal{N}} u_{i1}\alpha_i + \sum_{t\in[T]} \beta_t \\
\text{s.t. } (1) \quad & \beta_t + \sum_{i\in\mathcal{N}} u(A) \cdot \alpha_i \geq \sum_{i\in\mathcal{N}} \lambda_{it} R_i(A, 0) \ \forall t \in [T], A \in \mathcal{A} \\
(2) \quad & \alpha_i, \beta_t \geq 0.
\end{aligned}$$

Noting that the constraints in (1) can be re-written as

$$\beta_t = \max_{A\in\mathcal{A}} \sum_{i\in\mathcal{N}} (\lambda_{it} R_i(A, 0) - u(A) \cdot \alpha_i),$$

the dual can be more succinctly expressed as

$$\text{OPT}_{\text{fluid}} = \min_{\alpha_i \geq 0} \sum_{i\in\mathcal{N}} u_{i1}\alpha_i + \sum_{t\in[T]} \max_{A\in\mathcal{A}} \sum_{i\in\mathcal{N}} (\lambda_{it} R_i(A, 0) - u(A) \cdot \alpha_i).$$

The bound and final guarantee. To conclude this section and the proof of Theorem 2, we show the aforementioned upper bound on the dual's optimal objective value. We do so by bounding both terms that make-up $\text{OPT}_{\text{fluid}}$ by $2H_1(U_1)$.

LEMMA 6. $4H_1(U_1) \geq \text{OPT}_{\text{fluid}}$.

All-in-all, recalling that $\bar{H}_1(U_1)$ is the expected revenue garnered by our roll-out policy, we see that

$$\bar{H}_1(U_1) \geq \frac{1}{2} H_1(U_1) \geq \frac{1}{8} \text{OPT}_{\text{fluid}} \geq \frac{1}{8} \text{OPT}$$

where the first inequality follows by Lemma 5 and the second inequality is a result of Lemma 6.

4. Computational Experiments: Opaque-Only Setting

In this section, we present the details of an extensive set of computational experiments aimed at measuring the efficacy the approximation scheme outlined in Section 2 for the opaque-only setting. More specifically, we generate a large and diverse suite of test instances, and test three variants of our approach that differ in the manner in which the core randomized policy is derandomized.

4.1. Instance generator

We randomly generate problem instances with $n \in \{5, 10, 20\}$ products, where to ensure a sufficiently long time horizon, set $T = 10n$. Recalling that α determines the mix of risk averse and risk neutral customers, we vary $\alpha \in \{0, 0.5, 1\}$. The remaining parameters are generated using the following schema.

- *Prices:* We vary $L \in \{2, 5, 10\}$, and generate the prices uniformly from the interval from the interval $[0, 1]$.
- *MNL-based weights:* For each instance, we generate β uniformly from the interval $[0, 1]$, and w_i uniformly from the interval $[3, 4]$, and then we set $w_{i\ell} = w_i e^{-\beta p_\ell}$.
- *Initial inventories:* For each product $i \in [n]$, let $w_i^{\max} = \max_{\ell \in [L]} w_{i\ell}$ denote the weight of product i when priced at the lowest price level. We set

$$u_{i1} = \lceil \gamma * \frac{w_i^{\max}}{1 + \sum_{j \in [n]} w_j^{\max}} * T \rceil$$

  to be the initial inventory levels, where the parameter γ controls how these initial levels compare to the maximum possible total demand. We vary $\gamma \in \{0.1, 0.3, 1.1\}$.

Summarizing the discussion above, we characterize a single test case through the parameter combination $(n, L, \gamma, \alpha) \in \{5, 10, 20\} \times \{2, 5, 10\} \times \{0.1, 0.3, 1.1\} \times \{0, 0.5, 1\}$. For each test case, we generate 100 problem instances, ultimately reporting the results as various summary statistics over all instances for a fixed test case.

4.2. Implemented policies

The three policies that we implement are detailed below. The intent in testing these three approaches is to get a sense of the extent to which derandomization improves the performance of our randomized policy.

Randomized Policy (RP). This is the randomized policy prescribed by Algorithm 1 in Section 2. In the event that the focal product i_q has already stocked out when q is randomly chosen, we offer nothing in period t. The expected revenue of the randomized policy is derived by solving the dynamic program in (4).

Derandomized Policy (DRP) : This is the policy that results from the derandomization of RP described in Section 2.4. The expected revenue of this approach can be obtained through the dynamic program given in (5).

Generalized Derandomized Policy (G-DRP) : This policy is a generalization of DRP, which is derived from a modified version of (5), where we remove the restriction in the inner maximization that we must choose a $q \in [Q]$ such that $P_{qt} > 0$. More specifically, G-DRP is derived by solving dynamic program.

$$\hat{\mathcal{R}}_t^G(i,u) = \max \Bigg\{ \mathbb{1}_{u>0} \cdot \left(\max_{q\in[Q]:i_q=i} \left\{ \pi_q \cdot \left(p_{\ell_q} + \hat{\mathcal{R}}_{t+1}^G(i,u-1) - \hat{\mathcal{R}}_{t+1}^G(i,u) \right) + \hat{\mathcal{R}}_{t+1}^G(i,u) \right\} \right), \\ \max_{q\in[Q]:i_q>i} \left\{ \pi_q \cdot \left(p_{\ell_q} + \hat{\mathcal{R}}_{t+1}^G(i,u_{i_q1}-1) - \hat{\mathcal{R}}_{t+1}^G(i,u_{i_q1}) \right) + \hat{\mathcal{R}}_{t+1}^G(i,u_{i_q1}) \right\} \Bigg\},$$

A straightforward inductive arguments reveals that $\hat{\mathcal{R}}_t^G(i,u) \geq \hat{\mathcal{R}}_t(i,u)$, and hence G-DRP is guaranteed to improve upon DRP. In short, this modified dynamic program allows for deviations from the policy prescribed by the optimal solution to (OPA-LP), thus giving way to a more nuanced state-dependent policy.

4.3. Results

Table 1 provides a detailed summary of the results for the test cases in which $\alpha = 0.5$ For brevity, we omit the results for $\alpha \in \{0,1\}$, since they were qualitatively identical to those observed for $\alpha = 0.5$. The first four columns of Table 1 identify the test case, while the fifth and sixth column respectively denote the value of the average minimum inventory level over the 100 problem instances, and worst case optimality gap as established in Theorem 1. The remaining nine columns report the average, 25th percentile and minimum optimality gaps of the three implemented policies, where for each instance, the optimality gap is computed with respect to the optimal objective of (Reduced-OPA-LP).

There are a handful of salient trends to be immediately gleaned from Table 1. First, we see that all three policies perform significantly better than their worst-case guarantee, generally halving the optimality gap presented in column six. Additionally, we that the derandomization afforded to both DRP and G-DRP does indeed result in significant improvements over RP, which can be as high as 10%. Finally, as expected, we do see that the performance of all three policies dramatically improve as the initial inventories are scaled up. In particular, when $\gamma = 1.1$, G-DRP is within 1% of optimal, on average.

5. Expedia Case Study: Measuring the Benefits of Opaque Pricing

In this section, we use the general setting considered in Section 3 along with the publicly available data from Expedia's booking platform (Kaggle 2013), to measure the extent to which introducing an opaque product can improve revenues. In particular, we use the Expedia data set, which provides extensive accounts of the displayed hotels and subsequent purchasing decisions for thousands of

Table 1 **Results for test cases with** $\alpha = 0.5$

						RP			DRP			G-DRP		
T	N	L	γ	$u_{\min}$	$(1-\text{APX}(u_{\min}))\times 100$	mean	25%	min	mean	25%	min	mean	25%	min
50	5	2	0.1	1	36.79	17.86	17.14	21.62	8.3	7.34	11.44	8.15	7.24	11.08
50	5	2	0.3	2.78	23.43	11.41	11.01	13.12	7.03	6.44	8.52	6.73	6.24	8.28
50	5	2	1.1	8.88	13.32	5.44	5.19	6.48	4.09	3.72	5.53	1.11	0.79	2.1
50	5	5	0.1	1	36.79	18.71	17.46	22.43	9.02	8.36	11.56	8.82	8.2	11.26
50	5	5	0.3	2.97	22.54	11.69	11.36	13.18	7.35	6.86	9.82	6.97	6.62	9.08
50	5	5	1.1	9.29	13	5.56	5.25	6.49	4.18	3.82	5.49	1.29	0.96	2.16
50	5	10	0.1	1	36.79	19.04	17.62	24.31	9.31	8.47	13.69	9.06	8.2	13.08
50	5	10	0.3	3	22.4	11.83	11.43	13.49	7.44	6.9	9.75	7.02	6.58	9.37
50	5	10	1.1	9.65	12.75	5.6	5.3	6.64	4.31	3.96	5.8	1.31	1.03	2.29
100	10	2	0.1	1	36.79	17.56	17.04	18.68	8.4	7.94	10.41	8.32	7.89	10.21
100	10	2	0.3	2.69	23.85	11.18	10.94	12.31	6.79	6.31	8.28	6.64	6.17	8.08
100	10	2	1.1	8.61	13.53	5.24	5.05	6.02	3.96	3.76	4.74	0.87	0.7	1.55
100	10	5	0.1	1	36.79	17.78	17.15	21.1	8.74	8.36	11.8	8.64	8.3	11.57
100	10	5	0.3	2.97	22.54	11.47	11.19	12.53	7.15	6.8	8.56	6.94	6.62	8.2
100	10	5	1.1	9.39	12.93	5.48	5.31	6.33	4.16	3.92	5.11	1.05	0.91	1.57
100	10	10	0.1	1	36.79	17.85	17.17	20.03	8.75	8.42	10.16	8.64	8.32	9.96
100	10	10	0.3	3	22.4	11.49	11.21	12.34	7.2	6.9	8.32	6.97	6.72	7.99
100	10	10	1.1	9.65	12.75	5.58	5.4	6.26	4.23	4	5.15	1.15	1.01	1.75
200	20	2	0.1	1	36.79	17.17	16.88	18.49	8.1	7.73	9.6	8.07	7.71	9.5
200	20	2	0.3	2.77	23.47	11.08	10.9	11.83	6.72	6.4	7.51	6.64	6.36	7.37
200	20	2	1.1	8.69	13.48	5.32	5.19	6.1	4.01	3.92	4.7	0.72	0.59	1.11
200	20	5	0.1	1	36.79	17.38	17.04	19.44	8.51	8.24	9.63	8.46	8.19	9.53
200	20	5	0.3	2.98	22.49	11.26	11.06	11.84	6.96	6.68	7.78	6.86	6.6	7.64
200	20	5	1.1	9.2	13.07	5.46	5.34	6	4.13	4	4.74	0.88	0.76	1.28
200	20	10	0.1	1	36.79	17.49	17	19.02	8.69	8.38	9.73	8.62	8.33	9.63
200	20	10	0.3	3	22.4	11.36	11.21	12	7.06	6.86	7.71	6.94	6.75	7.57
200	20	10	1.1	9.56	12.81	5.55	5.43	5.93	4.2	4.1	4.62	0.95	0.86	1.3

searches on the platform, to derive a collection of practically reasonable instantiations of our dynamic pricing problem under the OS independent demand model. We then implement the pricing policy described in Section 3, and compare the revenue it generates to the revenue generated by a policy that sets $S_{\mathcal{O}} = \emptyset$, i.e. never offers an opaque option.

5.1. Model fitting

We fit the Expedia data [4] to an MNL-based choice model that conforms to our independent demand assumption. We also exclusively focus on 3-star and 4-star hotels in our estimation and experiments. In order for the purchase probabilities to be consistent with those observed in the data, we impose the independent demand assumption in our estimation process. More specifically, we assume each customer has an ideal product and she only chooses between her ideal product and the no-purchase option. To do so, we first define each observed customer's ideal product. This is simply done via the following procedure:

[4] See Appendix D for a description of the data set

(1) if a search results in a purchase of some hotel i, then the customer's ideal product is i;

(2) if a search does not result in a purchase of any products, then the first clicked hotel option i is regarded as the customer's ideal product.

Recall that a customer is called a type-i customer if her ideal product is i. After recovering each search customer's type as above, we then fit an MNL model by assuming each customer is only choosing between her ideal product and the no-purchase option. Furthermore, we simultaneously estimate the arrival probability of customer type i from the data with the MNL estimate. We assume that the arrival probability of customer type i is time-invariant, e.g., $\lambda_{i,t} = \lambda_i$. We simultaneously estimate λ and the MNL model parameter β by maximizing the probability of observing the outcome in the data, which we detail below.

We index the historical searches by $t \in \{1, ..., T\}$. For each search t, we use X_t to denote the logged feature information of the ideal product for t, $y_t \in \{0, 1\}$ to indicate the purchase outcome, and I_t to denote the type of the observed customer in search t. Let the feature coefficients in the MNL model be denoted by β. Given the observed outcome in the data, the joint probability of observing the outcome is

$$\prod_{t=1}^{T} (\sum_i \lambda_i \cdot \mathbb{1}_{\{I_t = i\}}) \frac{(1 - y_t) \cdot 1 + y_t \cdot e^{X_t \beta}}{1 + e^{X_t \beta}}$$

where $\sum_i \lambda_i \cdot \mathbb{1}_{\{I_t=i\}}$ is the arrival probability of customer type I_t. Hence the log-likelihood of the above joint probablity is

$$\begin{aligned} \mathcal{LL}(\lambda, \beta) &= \sum_{t \in [T]} \left[\log(\sum_i \lambda_i \cdot \mathbb{1}_{\{I_t=i\}}) + y_t X_t \beta - \log(1 + e^{X_t \beta}) \right] \\ &= \sum_i m_i \log(\lambda_i) + \sum_{t \in [T]} \left[y_t X_t \beta - \log(1 + e^{X_t \beta}) \right] \end{aligned} \tag{11}$$

where $m_i = \sum_{t \in [T]} \mathbb{1}_{\{I_t=i\}}$ is the total number of arrivals of customer type i in the T periods. We estimate $(\hat{\lambda}, \hat{\beta}) = \arg\max \mathcal{LL}(\lambda, \beta)$, subject to $\sum_i \lambda_i = 1$. Note that $\sum_i m_i \log(\lambda_i)$ is a concave function of λ, and the estimation of λ, β can be separated in $\mathcal{LL}(\lambda, \beta)$. Applying a standard convex optimization approach will reveal that $\hat{\lambda}_i$ is estimated by $\hat{\lambda}_i = \frac{m_i}{T}$, and $\hat{\beta}$ can be estimated via a standard logistic regression.

Estimation results. We apply the above-mentioned estimation approach to the search destination level data and generate feature coefficients for each search destination. We show the estimation results for two top-searched destinations and conduct our experiments mainly for these two destinations in the next section. We first summarize the feature information used in our model in Table 3. We include an intercept term that captures the average effect of all other features not included in our model. Note that features (1) - (5) are hotel-dependent and features (7) - (11) are search-dependent, while feature (6) is both hotel-dependent and search-dependent (time-dependent). We

		Hotel Count	Estimated $\hat{\lambda}$					Price History				Price Menu		
site	destination	3 & 4 stars	min	25th	50th	75th	max	median	mean	10th	90th	Min	Max	L
5	8347	256	0.18	0.18	0.35	0.53	1.58	120	141	69	220	70	220	16
5	9402	177	0.25	0.25	0.49	0.74	3.45	179	197	110	299	110	300	20

Table 2 Estimated Arrival Prob ($\hat{\lambda}$) and Price Menu

β	Feature Name	Description
(0)	intercept	Dummy variable that captures the effects of all features not listed below
(1)	price_usd	Hotel price per night per room (in USD)
(2)	prop_starrating_3	Dummy variable indicating whether the hotel is 3-star (otherwise 4-star)
(3)	prop_review_score	Customer review score of a hotel
(4)	prop_location_score1	Location score of a hotel
(5)	prop_brand_bool	Boolean variable indicating whether the hotel belongs to a major hotel chain
(6)	prop_log_historical_price	Average log historical price of a hotel over a period of time preceding the search
(7)	srch_booking_window	Number of days before the start of the trip at the time of a search
(8)	srch_length_of_stay	Duration of stay (in hotel nights)
(9)	srch_adults_count	Number of adults
(10)	srch_children_count	Number of children
(11)	srch_saturday_night_bool	Boolean variable indicating whether the trip includes a Saturday night

Table 3 Feature Summary

$\hat{\beta}$	Feature Name	estimate	std.err	p value
(0)	(Intercept)	-2.5240	2.0114	0.2095
(1)	price_usd	-0.0104	0.0054	0.0569
(2)	prop_starrating_3	-0.4418	0.4197	0.2925
(3)	prop_review_score	0.3285	0.3693	0.3738
(4)	prop_location_score1	0.1399	0.2811	0.6188
(5)	prop_brand_bool	0.2025	0.4489	0.6519
(6)	prop_log_historical_price	-0.0104	0.1156	0.9281
(7)	srch_booking_window	-0.0012	0.0043	0.7758
(8)	srch_length_of_stay	-0.3427	0.1496	0.0220
(9)	srch_adults_count	-0.0824	0.2232	0.7119
(10)	srch_children_count	0.3103	0.1645	0.0592
(11)	srch_saturday_night_bool	0.6952	0.4356	0.1105

(a) site = 5, dest = 8347

$\hat{\beta}$	Feature Name	estimate	std.err	p value
(0)	(Intercept)	-8.7345	3.3708	0.0096
(1)	price_usd	-0.0096	0.0049	0.0516
(2)	prop_starrating_3	-0.4522	0.5228	0.3871
(3)	prop_review_score	1.5574	0.8104	0.0546
(4)	prop_location_score1	0.1082	0.1826	0.5534
(5)	prop_brand_bool	0.4736	0.6995	0.4983
(6)	prop_log_historical_price	0.1073	0.1468	0.4649
(7)	srch_booking_window	-0.0012	0.0051	0.8078
(8)	srch_length_of_stay	-0.0714	0.179	0.6898
(9)	srch_adults_count	0.0926	0.3252	0.7759
(10)	srch_children_count	0.6786	0.279	0.015
(11)	srch_saturday_night_bool	0.2733	0.4458	0.5399

(b) site = 5, dest = 9402

Table 4 Estimated Coefficients ($\hat{\beta}$)

will specify how we set the values of the search-dependent features when we conduct our experiments.

We provide $(\hat{\lambda}, \hat{\beta})$ estimates for two top searched destinations, both of which are from the site 5. These two destinations have the most statistically significant price coefficients among the top destinations. Because pricing is a critical component of our work, we choose to further conduct our experiments for these two destinations. We provide a summary of our estimation results for $\hat{\lambda}$ in Table 2 and for $\hat{\beta}$ in Table 4.

5.2. Instance generator

We conduct independent experiments for the two search destinations to test our proposed approach for the general assortment and pricing problem. For each destination, we generate independent problem instances that we control by overall assortment size, number of time periods, and starting inventory. Below we explain how we construct the price menu, the preference weights, and control the problem instances in our experiments.

Price menu. For each destination, we set the minimum and maximum price on the price menu by (Min, Max) = (10th, 90th) quantile of hotel prices observed in our MNL model data, which are then rounded to the nearest 10s. We then construct a price menu starting from the Min with an incremental of $10, which is then capped by the Max. This has resulted in a menu of L price levels, where $L = 16$ and 20 for the two destinations of interest (see Table 2).

Preference weights. For a regular product i priced at price level $l \in [L]$, we set its preference weight (as a regular product) by

$$w_{il} = e^{\hat{\beta}^{(0)} x_{t,i}^{(0)} + \hat{\beta}^{(2)} x_{t,i}^{(2)} + \ldots + \hat{\beta}^{(11)} x_{t,i}^{(11)}}$$

For an opaque product $S_{\mathcal{O}} \subseteq [N]$ (with $|S_{\mathcal{O}}| \geq 2$) priced at price level $l \in [L]$, we set its preference weight for customer type i by

$$w_{il}^{\mathcal{O}} = \mathbb{1}_{\{i \in S_{\mathcal{O}}\}} \cdot \frac{w_{il}}{|S_{\mathcal{O}}|}$$

For each i and l, w_{il} can be calculated with given feature values $[1, x_{t,i}^{(1)}, ..., x_{t,i}^{(11)}]$. We assume that the feature values are time-invariant in our experiments except for $x_{t,i}^{(1)}$ (*price_usd*) and hence need to fix other features. We detail how we choose the feature values below.

(i) $x_{t,i}^{(1)}$. We set $x_{t,i}^{(1)} \triangleq p_l$ for some $l \in [L]$, which is our decision variable for pricing hotel options.

(ii) $x_{t,i}^{(2)}, ..., x_{t,i}^{(5)}$. These feature values are inherited from the observed data and are invariant.

(iii) $x_{t,i}^{(6)}, ..., x_{t,i}^{(11)}$. We use the median of these features' values for each hotel i from the observed data. We apply the median values rather omitting these features (which is equivalent to setting their values to 0) because omission would inflate/deflate preference weights and hence may make purchase probabilities in our experiments less consistent with the observed data.

Problem instances. We generate independent problem instances by varying the combination of the following parameters:

(a) N: the set of all products in a problem instance. We vary $N \in \{10, 20\}$ and randomly draw N hotels from the entire hotel list used in the β-estimation data.

(b) T: the number of time periods. We vary $T \in \{50, 100\}$.

(c) γ: inventory inflator. We control initial product inventory by setting $\gamma \in \{0.25, 0.5, 1\}$. First, our model assumes there is exactly one customer arriving in each time period, and we have estimated the arrival rate of customers as $\hat{\lambda}$. Since we draw a subset of N hotels to constitute the product set in a single problem instance, we adjust the arrival rate of customer type i by setting $\lambda_i = \frac{\hat{\lambda}_i}{\sum_{i \in [N]} \hat{\lambda}_i}$. We then set the initial inventory of product i as $\lceil \gamma \cdot \lambda_i T \rceil$, where $\lceil \cdot \rceil$ is the ceiling operator.

Summary. Combining the above discussion, each problem instance in our experiments is characterized by a tuple (N, T, γ). Furthermore, for each of the tuples, we randomly draw the set of hotels (of size N) 5 times and indicate the j-th draw by $m = 1, ..., 5$. Hence our single problem instance is denoted by (N, T, γ, m). For each problem instance, we generate 500 sample paths of customer arrivals and simulate their purchase choices through the random draw of a customer's idiosyncratic utilities for the transparent, opaque, and no-purchase options so that the revenue performance of each approach is not biased by distinct samples of customer arrivals. We report the expected revenue for a problem instance by averaging the revenues from the 500 sample paths.

5.3. Results.

For each problem instance, we implement the approximation scheme outlined in Section 3 (referred to as OPA), which is benchmarked against a policy that never offers an opaque product (referred to as BMK). It is not difficult to see that when we fix $S_{\mathcal{O}} = \emptyset$, our dynamic pricing problem of interest decomposes into independent pricing problems for each product, which can be solved optimally via standard dynamic programming ideas. Consequently, the benchmark approach is afforded an optimal algorithm (given no opaque offerings), while our approach is guaranteed only a $\frac{1}{8}$-th fraction of the optimal expected revenue.

Due to space limitations, we only present a detailed look at the results for destination 8347 with $n = 20$, which are given in Table 5. The first four columns give the problem instance, while the proceeding eight columns report the summary statistics of the revenue performance of OPA and BMK. The final columns report the revenue gain of our approach over the benchmark, computes as

$$\text{Revenue gain } (\%) = \frac{\text{OPA} - \text{BMK}}{\text{BMK}} \times 100\%$$

Quite surprisingly, we observed that the expected revenue of our approach outperformed that of the benchmark for every problem instance. This trend is a strong indicator that opaque pricing can be quite beneficial, and it also likely points to the fact that our approach performs far better than its worst case guarantee. Overall, averaging over all test cases, we found that our approach outperformed the benchmark by 7.22% and 5.27% for n = 10 and 20 in search destination 8347, and by 5.56% and 4.21% for n = 10 and 20 in search destination 9402.

6. Future Work

Our work opens the door for extensions and potential improvements along multiple dimensions. First, it is natural to wonder if it is possible to improve upon the performance guarantee of $\frac{1}{8}$-th presented for the general setting using a new set of tools. Alternatively, it would be beneficial to pursue the hardness of approximation results that establish an upper bound on the best-possible

				OPA				BMK				
N	T	γ	m	mean	sd	CI (95%)		mean	sd	CI (95%)		Gain(%)
20	50	0.25	1	3359	463	3318	3400	3170	412	3134	3206	5.96
20	50	0.25	2	3434	435	3396	3472	3166	373	3133	3198	8.46
20	50	0.25	3	3302	437	3264	3340	3105	403	3069	3140	6.34
20	50	0.25	4	3605	442	3566	3644	3307	392	3273	3341	9.01
20	50	0.25	5	3354	449	3315	3394	3164	420	3127	3201	6.01
20	50	0.5	1	3871	525	3825	3917	3690	476	3648	3731	4.91
20	50	0.5	2	4585	519	4539	4630	4303	486	4261	4346	6.55
20	50	0.5	3	4123	534	4076	4170	3968	499	3924	4012	3.91
20	50	0.5	4	4742	539	4695	4789	4473	513	4428	4518	6.01
20	50	0.5	5	4444	553	4395	4492	4237	546	4189	4285	4.89
20	50	1	1	4909	694	4849	4970	4760	647	4704	4817	3.13
20	50	1	2	5598	674	5539	5657	5379	656	5321	5436	4.07
20	50	1	3	4884	650	4826	4941	4733	643	4677	4789	3.19
20	50	1	4	5949	710	5887	6011	5766	693	5705	5826	3.17
20	50	1	5	5106	711	5044	5168	4905	667	4846	4963	4.10
20	100	0.25	1	5642	502	5598	5686	5242	426	5205	5280	7.63
20	100	0.25	2	6267	510	6222	6311	5773	422	5736	5810	8.56
20	100	0.25	3	6042	523	5996	6088	5679	452	5640	5719	6.39
20	100	0.25	4	6473	575	6423	6524	6006	475	5964	6047	7.78
20	100	0.25	5	6700	560	6651	6749	6327	509	6283	6372	5.90
20	100	0.5	1	8542	774	8474	8609	8182	724	8118	8245	4.40
20	100	0.5	2	9631	730	9567	9695	9092	674	9033	9151	5.93
20	100	0.5	3	8416	751	8350	8482	8067	706	8005	8128	4.33
20	100	0.5	4	10051	768	9984	10119	9434	706	9372	9496	6.54
20	100	0.5	5	9267	886	9189	9345	8874	862	8798	8949	4.43
20	100	1	1	10183	1001	10095	10270	9860	921	9779	9940	3.28
20	100	1	2	11319	1043	11228	11410	10912	993	10825	10999	3.73
20	100	1	3	10135	1071	10041	10229	9834	1007	9745	9922	3.06
20	100	1	4	11831	954	11748	11915	11539	959	11455	11623	2.53
20	100	1	5	10190	1041	10099	10281	9816	981	9730	9902	3.81

(a) dest = 8347, n = 20

Table 5 Experiment results

approximation factor. Finally, there remains the question of whether non-trivial approximation schemes exist for the general setting under a more nuanced demand model in which customers value the opaque option based on its contents, similar to the model under consideration in Section 2.

References

Adelman D (2007) Dynamic bid prices in revenue management. *Operations Research* 55(4):647–661.

Aggarwal G, Goel G, Karande C, Mehta A (2011) Online vertex-weighted bipartite matching and single-bid budgeted allocations. *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, 1253–1264 (SIAM).

Alaei S, Hajiaghayi M, Liaghat V (2012) Online prophet-inequality matching with applications to ad allocation. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 18–35.

Anderson CK, Xie X (2012) A choice-based dynamic programming approach for setting opaque prices. *Production and Operations Management* 21(3):590–605.

Aouad A, Saban D (2022) Online assortment optimization for two-sided matching platforms. *Management Science* .

Elmachtoub AN, Hamilton ML (2021) The power of opaque products in pricing. *Management Science* 67(8):4686–4702.

Elmachtoub AN, Wei Y, Zhou Y (2015) Retailing with opaque products. *Available at SSRN 2659211* .

Fay S, Xie J (2008) Probabilistic goods: A creative way of selling products and services. *Marketing Science* 27(4):674–690.

Feldman J, Segev D (2022) The multinomial logit model with sequential offerings: Algorithmic frameworks for product recommendation displays. *Operations Research* 70(4):2162–2184.

Feng Y, Niazadeh R, Saberi A (2019) Linear programming based online policies for real-time assortment of reusable resources. *Chicago Booth Research Paper* (20-25).

Gallego G, Li A, Truong VA, Wang X (2015a) Online resource allocation with customer choice. *arXiv preprint arXiv:1511.01837* .

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

Gong XY, Goyal V, Iyengar GN, Simchi-Levi D, Udwani R, Wang S (2022) Online assortment optimization with reusable resources. *Management Science* 68(7):4772–4785.

Huang T, Yu Y (2014) Sell probabilistic goods? a behavioral explanation for opaque selling. *Marketing Science* 33(5):743–759.

Jerath K, Netessine S, Veeraraghavan SK (2010) Revenue management with strategic customers: Last-minute selling and opaque selling. *Management science* 56(3):430–448.

Kaggle (2013) Personalize expedia hotel searches - ICDM 2013. https://www.kaggle.com/c/expedia-personalized-sort/data.

Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, 352–358.

Liu Q, Van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2):288–310.

Liu Q, Xiao Y, Zhang D (2022) Opaque selling of multiple substitutable products with finite inventories. *Naval Research Logistics (NRL)* 69(4):529–549.

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

Ma W, Simchi-Levi D, Zhao J (2021) Dynamic pricing (and assortment) under a static calendar. *Management Science* 67(4):2292–2313.

Ma Y, Rusmevichientong P, Sumida M, Topaloglu H (2020) An approximation algorithm for network revenue management under nonstationary arrivals. *Operations Research* 68(3):834–855.

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

Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5):22–es.

Post D, Spann M (2012) Improving airline revenues with variable opaque products: "blind booking" at germanwings. *Interfaces* 42(4):329–338.

Ren H, Huang T (2022) Opaque selling and inventory management in vertically differentiated markets. *Manufacturing & Service Operations Management* 24(5):2543–2557.

Rusmevichientong P, Sumida M, Topaloglu H (2020) Dynamic assortment optimization for reusable products with random usage durations. *Management Science* 66(7):2820–2844.

Sinha P, Zoltners AA (1979) The multiple-choice knapsack problem. *Operations Research* 27(3):503–515.

Talluri K, Van Ryzin G (1998) An analysis of bid-price controls for network revenue management. *Management science* 44(11-part-1):1577–1593.

Talluri K, Van Ryzin G (1999) A randomized linear programming method for computing network bid prices. *Transportation science* 33(2):207–216.

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

Topaloglu H (2009) Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. *Operations Research* 57(3):637–649.

Appendix A: Missing Proofs of Section 2

A.1. Proof of Claim 1

For any opaque product $S \subseteq \mathcal{N}$ such that $i \in S$, it is straightforward to see that $\pi(\ell, S) \le \pi(\ell, S \cap [i,n])$, since removing any product indexed lower than i from the opaque product will not decrease the weight of the opaque option for either risk-averse or risk-neutral customers. Along the same lines, it is straightforward to see that adding product n (i.e. the highest-weight product) to any assortment to which it previously did not belong will only increase the choice probability of said assortment. Consequently, it is sufficient to show that

$$S_{\mathcal{O}}(\ell, i) = \underset{S \subseteq [i+1, n-1]: i, n \in S}{\arg\max} \pi(\ell, S).$$

To show the result above, we assume by way of contradiction that there exists some $S \subseteq \mathcal{N}$ such that $i \in S$ and $S \neq \{i\} \cup [k,n]$ for some $k \in [i,n]$, and whose choice probability satisfies $\pi(\ell, S) > \pi(\ell, S_{\mathcal{O}}(\ell, i))$. Given this supposed structure of S, there must exist products $q, j \in \mathcal{N}$ that satisfy $i < q < j$ such that $q \in S$ and $j \notin S$. We establish the contradiction by showing that the choice probability of S cannot decrease if either we remove product q or add product j. If we iteratively apply this argument until no such q and j exist, the resulting opaque product S will be a candidate for $S_{\mathcal{O}}(\ell, i)$.

To establish that one of these two changes must result in at least as high a choice probability, we first consider the risk-averse customers and show that either of these changes leaves the weight of the opaque option unchanged, meaning the choice probability is unchanged as well. To see this, observe that

$$w_\ell^{\mathrm{RA}}(S) = w_\ell^{\mathrm{RA}}(S \cup \{j\}) = w_\ell^{\mathrm{RA}}(S \setminus \{q\}) = w_{i\ell},$$

since risk-averse customers value the opaque product at the smallest weight offered. Moving to the risk-neutral customers, we have that

$$\begin{aligned} w_\ell^{\mathrm{RN}}(S) &= \frac{\sum_{j \in S} w_{j\ell}}{|S|} \\ &= \frac{|S|-1}{|S|} \cdot \frac{\sum_{j \in S \setminus \{q\}} w_{j\ell}}{|S|-1} + \frac{1}{|S|} \cdot w_{q\ell} \\ &= \frac{|S|-1}{|S|} \cdot w_\ell^{\mathrm{RN}}(S \setminus \{q\}) + \frac{1}{|S|} \cdot w_{q\ell}. \end{aligned}$$

Consequently, if $w_\ell^{\mathrm{RN}}(S \setminus \{q\}) \ge w_{q\ell}$, then we get that $w_\ell^{\mathrm{RN}}(S \setminus \{q\}) \ge w_\ell^{\mathrm{RN}}(S)$, and hence removing product q from S cannot decrease its choice probability. Otherwise, it must be the case $w_\ell^{\mathrm{RN}}(S \setminus \{q\}) \le w_\ell^{\mathrm{RN}}(S) \le w_{q\ell}$, and so

$$\begin{aligned} w_\ell^{\mathrm{RN}}(S \cup \{j\}) &= \frac{|S|}{|S|+1} \cdot w_\ell^{\mathrm{RN}}(S) + \frac{1}{|S|+1} \cdot w_{j\ell} \\ &\ge \frac{|S|}{|S|+1} \cdot w_\ell^{\mathrm{RN}}(S) + \frac{1}{|S|+1} \cdot w_\ell^{\mathrm{RN}}(S) \\ &= w_\ell^{\mathrm{RN}}(S), \end{aligned}$$

and hence adding product j to S can only increase its choice probability. The lone inequality above follows because, in this case, we must have that $w_{j\ell} \ge w_{q\ell} \ge w_\ell^{\mathrm{RN}}(S)$, since $j > q$.

A.2. Proof of Claim 1

We prove the result by showing that if there exists $h^*(\ell, S_{\mathcal{O}}, i) > 0$ where $S_{\mathcal{O}} \neq S_{\mathcal{O}}(\ell, i)$, then we can produce an alternative feasible solution in which $h^*(\ell, S_{\mathcal{O}}, i) = 0$ and the objective remains unchanged. Specifically, we simply set $h^*(\ell, S_{\mathcal{O}}, i) = 0$ and $h^*(\ell, S_{\mathcal{O}}(\ell,i), i) = h^*(\ell, S_{\mathcal{O}}(\ell,i), i) + \frac{\pi(\ell, S_{\mathcal{O}})}{\pi(\ell, S_{\mathcal{O}}(\ell,i))} \cdot h^*(\ell, S_{\mathcal{O}}, i)$. Straightforward algebra reveals that this alternative solution achieves the same objective, while also keep the left-hand-side of constraint (1) unchanged. Finally, since $\frac{\pi(\ell, S_{\mathcal{O}})}{\pi(\ell, S_{\mathcal{O}}(\ell,i))} \leq 1$, the left-hand-side of constraint (2) does not increase with this update.

A.3. Proof of Claim 2

We prove the result by showing that Reduced-OPA-LP is precisely a relaxation of a multiple-choice knapsack problem (MCKP), whose exact nature is formalized shortly. Once we show the reduction, the claim will follow immediately from known results regarding the structure of the optimal solution for the linear programming relaxations of the MCKP.

The MCKP. In the MCKP, there are n groups of L distinct items. The reward and capacity consumption of item $i \in [n]$ from group $\ell \in [L]$ is $r_{i\ell}$ and $c_{i\ell}$, respectively. The goal is to select at most a single item from each group to maximize the total reward of all selected items, while not exceeding the total capacity restriction of C. The exact integer programming formulation of MCKP is given below, where the binary decision variable $x_{i\ell}$ indicates whether item i in group ℓ is selected:

$$
\begin{aligned}
\max \ & \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} r_{i\ell} x_{i\ell} \\
\text{s.t.} \ & (1) \quad \sum_{\ell \in [L]} x_{i\ell} \leq 1 \qquad \forall i \in \mathcal{N} \\
& (2) \quad \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} c_{i\ell} x_{i\ell} \leq C \\
& (3) \quad x_{i\ell} \in \{0,1\}.
\end{aligned}
$$

The linear programming relaxation of MCKP simply replaces $x_{i\ell} \in \{0,1\}$ with $x_{i\ell} \geq 0$. Proposition 3 of Sinha and Zoltners (1979), establishes that the optimal solution to this linear programming relaxation has at most two fractional variables that must correspond to items from the same group. The proof of this result is a simple accounting argument involving the number of tight constraints required to make a basic feasible solution.

The reduction. We show that Reduced-OPA-LP is exactly an instance of the linear programming relaxation of the MCKP, where the groups are the products and the items within each group correspond to the L price points. The reduction follows quite smoothly by issuing the change of variable $h'(\ell, S_{\mathcal{O}}(\ell,i), i) = \frac{\pi(\ell, S_{\mathcal{O}}(\ell,i))}{u_{i1}} \cdot h(\ell, S_{\mathcal{O}}(\ell,i), i)$ within Reduced-OPA-LP. The end result is the following linear program

$$
\begin{aligned}
\max \ & \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} p_\ell u_{i1} h'(\ell, S_{\mathcal{O}}(\ell,i), i) \\
\text{s.t.} \ & (1) \quad \sum_{\ell \in [L]} h'(\ell, S_{\mathcal{O}}(\ell,i), i) \leq 1 \qquad \forall i \in \mathcal{N} \\
& (2) \quad \sum_{i \in \mathcal{N}} \sum_{\ell \in [L]} \frac{u_{i1}}{\pi(\ell, S_{\mathcal{O}}(\ell,i))} \cdot h'(\ell, S_{\mathcal{O}}(\ell,i), i) \leq T \\
& (3) \quad h'(\ell, S_{\mathcal{O}}(\ell,i), i) \geq 0,
\end{aligned}
$$

which matches a relaxation of MCKP with $r_{i\ell} = p_\ell u_{i1}$, $w_{i\ell} = \frac{u_{i1}}{\pi(\ell, S_{\mathcal{O}}(\ell,i))}$, and $C = T$.

A.4. Proof of Lemma 2

Intermediate results. Our proof critically relies on the following three results, all of which concern properties of binomial random variables. The proof of the claim can be found at the end of this section. Throughout this section, we use the shorthands Binom(n, λ) and Bern(λ) to respectively represent a binomial random variable with n trials, each with success probability λ, and a Bernoulli trial with success probability λ.

CLAIM 4. *Let $X = \sum_{t\in[T]} B_t$, where $B_t \sim \text{Bern}(\lambda_t)$ are independent Bernoulli trials with success probabilities $\sum_{t\in[T]} \lambda_t = \lambda$. For arbitrary integer $m \geq 0$, we have*

$$\mathbb{E}\left[\min\left\{X, m\right\}\right] \geq \mathbb{E}\left[\min\left\{\text{Binom}\left(\frac{\lambda}{T}, T\right), m\right\}\right].$$

LEMMA 7 (**Ma et al. (2021) Lemma EC.5**). *Let c be any real-valued positive number, T any positive integer. The function*

$$f(\lambda) = \frac{\mathbb{E}\left[\min\{c, \text{Binom}(T, \lambda/T)\}\right]}{\lambda}$$

is non-increasing in λ over the interval $[0, T]$.

LEMMA 8 (**Alaei et al. (2012) Lemma 5.3**). *For positive integer T and $\lambda \in [0, T]$, let $X \sim \text{Binom}(T, \frac{\lambda}{T})$, then $\mathbb{E}[\min\{X, b\}] \geq \text{APX}(b) \cdot \lambda$ for any $b \geq \lambda$.*

Proof. For each product $i \in \mathcal{N}$ such that $\mathcal{L}_i \neq \emptyset$, each price level $\ell \in \mathcal{L}_i$, we let $\tau_{i\ell}$ denote the random number of periods in which we offer the opaque product $S_{\mathcal{O}}(\ell, i)$ at price level ℓ, when i is the focal give-away product. For the remainder of the section, we use $h^*_{i\ell} = h^*(\ell, S_{\mathcal{O}}(\ell, i), i)$ and $\pi_{i\ell} = \pi(\ell, S_{\mathcal{O}}(\ell, i))$ as shorthands. Via Figure 3, which provides a visualization of our randomize policy, it is straightforward to see that there exists probabilities $\lambda_{1\ell}, \lambda_{2\ell} \geq 0$ and integer $t_{i\ell} \leq \lfloor h^*_{i\ell} \rfloor$ that satisfy $h^*_{i\ell} = t_{i\ell} + \lambda_{1\ell} + \lambda_{2\ell}$ and

$$\tau_{i\ell} \sim t_{i\ell} + \text{Bern}(\lambda_{1\ell}) + \text{Bern}(\lambda_{2\ell})$$

We prove the lemma by considering two cases based on whether $|\mathcal{L}_i| = 1$ or $|\mathcal{L}_i| = 2$.

- Case 1 - $|\mathcal{L}_i| = 1$: Define Y to be the random number of times opaque product $S_{\mathcal{O}}(\ell, i)$ is selected. Based on the preceding discussion we have that

$$Y \sim \text{Binom}(\tau_{i\ell}, \pi_{i\ell}) \sim B_0 + B_1 + B_2$$

where $B_0 = \text{Binom}(t_{i\ell}, \pi_{i\ell})$, $B_1 = \text{Bern}(\lambda_{1\ell}\pi_{i\ell})$, and $B_2 = \text{Bern}(\lambda_{2\ell}\pi_{i\ell})$. Moreover, define $X_1, \ldots, X_{t_{i\ell}+2}$, where

$$X_{i\ell} = \begin{cases} \text{Bern}(\pi_{i\ell}) & \text{if } i \leq t_{i\ell} \\ \text{Bern}(\lambda_{1\ell}\pi_{i\ell}) & \text{if } i = t_{i\ell} + 1 \\ \text{Bern}(\lambda_{2\ell}\pi_{i\ell}) & \text{if } i = t_{i\ell} + 2, \end{cases}$$

and let $X = \sum_{i=1}^{t_{i\ell}+2} X_i$, noting that $\mathbb{E}[X] = \pi_{i\ell} \cdot (t_{i\ell} + \lambda_{1\ell} + \lambda_{2\ell}) = \pi_{i\ell} h^*_{1\ell}$. It is straightforward to see that $Y \sim X$, and so, recalling that $D_{i\ell}$ is the random sales of opaque product $S_{\mathcal{O}}(\ell, i)$, we get hat $\mathbb{E}[D_{i\ell}] = \mathbb{E}[\min\{Y, u_{i1}\}] = \mathbb{E}[\min\{X, u_{i1}\}]$. From here, observe that

$$\begin{aligned} \frac{R_i}{\text{OPT}_i} &= \frac{p_\ell \cdot \mathbb{E}\left[D_{i\ell}\right]}{p_\ell \pi_{i\ell} h^*_{i\ell}} \\ &= \frac{\mathbb{E}[\min\{X, u_{i1}\}]}{\pi_{i\ell} h^*_{i\ell}} \\ &\geq \frac{\mathbb{E}[\min\{\text{Binom}(t_{i\ell} + 2, \frac{\pi_{i\ell} h^*_{i\ell}}{t_{i\ell}+2}), u_{i1}\}]}{\pi_{i\ell} h^*_{i\ell}}, \end{aligned} \qquad (12)$$

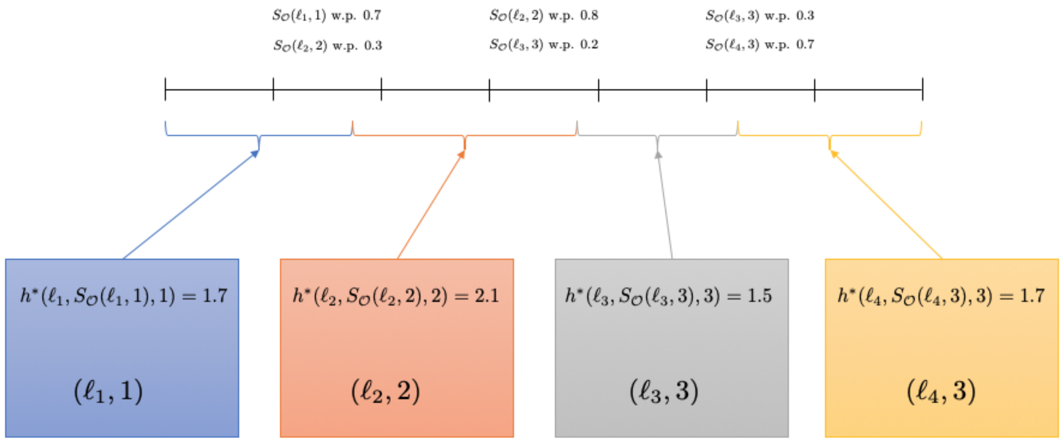

Figure 3 **Depiction of our pricing policy, in which we randomize in periods 2,4, and 6. When $i=2$, for example, we have that $t_{2\ell_2}=1, \lambda_{1\ell_2}=0.3, \lambda_{2\ell_2}=0.8$.**

where the lone inequality follows from Claim 4. Now, if $t_{i\ell}+2 \leq u_{i1}$, we see that

$$\mathbb{E}[\min\{\text{Binom}(t_{i\ell}+2, \frac{\pi_{i\ell}h^*_{i\ell}}{t_{i\ell}+2}), u_{i1}\}] = \pi_{i\ell}h^*_{i\ell},$$

and hence $\frac{R_i}{\text{OPT}_i}=1$. Otherwise. we get that

$$\begin{aligned}(12) &\geq \frac{\mathbb{E}[\min\{\text{Binom}(t_{i\ell}+2, \frac{u_{i1}}{t_{i\ell}+2}), u_{i1}\}]}{u_{i1}} \\ &\geq \text{APX}(u_{i1}).\end{aligned}$$

The first inequality follows from Theorem 7, while the second uses Lemma 8.

- Case 2 - $|\mathcal{L}_i|=2$: In this second case, we will assumes $\mathcal{L}_i=\{\ell,\ell'\}$ and that $p_\ell \geq p_{\ell'}$, so price level ℓ is offered first within our randomized policy. Define Y and Y' to be the random number of times opaque products $S_{\mathcal{O}}(\ell,i)$ and $S_{\mathcal{O}}(\ell',i)$ are purchased under our policy. We have that

$$\begin{aligned}R_i &= p_\ell \cdot \mathbb{E}\left[D_{i\ell}\right] + p_{\ell'} \cdot \mathbb{E}\left[D_{i\ell'}\right] \\ &= p_\ell \cdot \mathbb{E}\left[\min\{Y,u_{i1}\}\right] + p_{\ell'} \cdot \mathbb{E}\left[\min\{Y',(u_{i1}-Y)^+\}\right] \\ &= p_{\ell'} \cdot \mathbb{E}\left[\min\{Y,u_{i1}\} + \min\{Y',(u_{i1}-Y)^+\}\right] + (p_\ell - p_{\ell'}) \cdot \mathbb{E}\left[\min\{Y,u_{i1}\}\right] \\ &= p_{\ell'} \cdot \mathbb{E}\left[\min\{Y+Y',u_{i1}\}\right] + (p_\ell - p_{\ell'}) \cdot \mathbb{E}\left[\min\{Y,u_{i1}\}\right].\end{aligned}$$

Moreover, we have that

$$\begin{aligned}\text{OPT}_i &= p_\ell \pi_{i\ell} h^*_{i\ell} + p_{\ell'} \pi_{i\ell'} h^*_{i\ell'} \\ &= p_{\ell'} \cdot (\pi_{i\ell} h^*_{i\ell} + \pi_{i\ell'} h^*_{i\ell'}) + (p_\ell - p_{\ell'}) \cdot \pi_{i\ell} h^*_{i\ell}.\end{aligned}$$

Consequently, to prove the lemma, it is sufficient to show that

$$\frac{\mathbb{E}\left[\min\{Y+Y',u_{i1}\}\right]}{\pi_{i\ell}h^*_{i\ell} + \pi_{i\ell'}h^*_{i\ell'}} \geq \text{APX}(\text{u}_{\text{i1}})$$

and

$$\frac{\mathbb{E}\left[\min\{Y, u_{i1}\}\right]}{\pi_{i\ell} h^*_{i\ell}} \geq \mathrm{APX(u_{i1})}$$

Both of these bounds can be established by repeating the same set of arguments as presented for Case 1.

Proof of Claim 4. The claim is trivially true if $\lambda_1 = \lambda_2 = \ldots = \lambda_T = \frac{\lambda}{T}$, and hence we proceed under the assumption that there exists at least one pair $t_1, t_2 \in [T]$ such that $\lambda_{t_1} > \frac{\lambda}{T}$ and $\lambda_{t_2} < \frac{\lambda}{T}$. Let $\epsilon = \min\{\lambda_{t_1} - \frac{\lambda}{T}, \frac{\lambda}{T} - \lambda_{t_2}\}$, and define a new collection of Bernoulli random variables $\{\hat{B}_t\}_{t\in[T]}$, where for $t \in [T] \setminus \{t_1, t_2\}$ we set $\hat{B}_t = B_t$, while $\hat{B}_{t_1} = \mathrm{Bern}(\lambda_{t_1} - \epsilon)$ and $\hat{B}_{t_2} = \mathrm{Bern}(\lambda_{t_2} + \epsilon)$. By construction, it must be the case that at least one of $\hat{B}_{t_1}$ and $\hat{B}_{t_1}$ has success probability of $\frac{\lambda}{T}$. Consequently, letting $\hat{X} = \sum_{t\in[T]} \hat{B}_t$, if we can show

$$\mathbb{E}\left[\min\{X, m\}\right] \geq \mathbb{E}\left[\min\left\{\hat{X}, m\right\}\right] \tag{13}$$

then the proof is complete, since the above reasoning can be applied iteratively until all success probabilities are $\frac{\lambda}{T}$.

To show (13), note that

$$\begin{aligned}
&\mathbb{E}\left[\min\{X, m\}\right] - \mathbb{E}\left[\min\left\{\hat{X}, m\right\}\right] = \\
&\underbrace{\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right]}_{\text{term 1}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right] + \\
&\underbrace{\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right]}_{\text{term 2}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right] + \\
&\underbrace{\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m-1\right]}_{\text{term 3}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m-1\right].
\end{aligned}$$

We proceed to show that each of the three terms is non-negative:

- term 1: For the first term, we have that

$$\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right] = 0,$$

  since both X and $\hat{X}$ will be larger than m with probability 1.

- term 2: In this case, we have

$$\begin{aligned}
&\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right] \\
&= \mathbb{E}\left[X - \hat{X}\right] \\
&= \mathbb{E}\left[\sum_{t\in\{t_1,t_2\}} \left(B_t - \hat{B}_t\right)\right] \\
&= (\lambda_{t_1} + \lambda_{t_2}) - (\lambda_{t_1} - \epsilon + \lambda_{t_2} + \epsilon) \\
&= 0.
\end{aligned}$$

- term 3: Here, we have

$$\begin{aligned}
&\mathbb{E}\left[\min\{X,m\}-\min\left\{\hat{X},m\right\} \;\middle|\; \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m-1\right] \\
&\quad = \mathbb{E}\left[\min\{B_{t_1}+B_{t_2},1\}\right]-\mathbb{E}\left[\min\{\hat{B}_{t_1}+\hat{B}_{t_2},1\}\right] \\
&\quad = (\lambda_{t_1}+(1-\lambda_{t_1})\cdot\lambda_{t_2})-(\lambda_{t_1}-\epsilon+(1-\lambda_{t_1}+\epsilon)\cdot(\lambda_{t_2}+\epsilon)) \\
&\quad = \epsilon\cdot\underbrace{(\lambda_{t_1}-\lambda_{t_2})}_{\geq\epsilon}-\epsilon^2 \\
&\quad \geq 0
\end{aligned}$$

A.5. Proof of Claim 3

Proof of (i). If product i is allocated in period t, then there must $q^*\in[Q]$ such that $i_{q^*}=i$ and $P_{q^*t}>0$. As such, based on our construction of the offer probabilities that guide the randomization, it must be the case that $T_{q^*-1}<t$, since otherwise, it is straightforward to see that $P_{q^*t}=0$. In the statement of (i), we focus on some $q\in[Q]$ such that $i_q<i=i_{q^*}$, and hence it must be the case that $q<q^*$, which implies that $T_q\leq T_{q^*-1}<t$. Consequently, we get that $[T_{q-1},T_q]\cap[t,n]=\emptyset$, and hence $P_{q\tau}=0$ for any $\tau>t$, as desired.

Proof of (ii). If product i is allocated in period t, then there must $q^*\in[Q]$ such that $i_{q^*}=i$ and $P_{q^*t}>0$. As such, based on our construction of the offer probabilities that guide the randomization, it must be the case that $T_{q^*}\geq t-1$, since otherwise, it is straightforward to see that $P_{q^*t}=0$. In the statement of (ii), we focus on some $q\in[Q]$ such that $i_q>i=i_{q^*}$, and hence it must be the case that $q>q^*$, which implies that $T_{q-1}\geq T_{q^*}\geq t-1$. Consequently, we get that $[T_{q-1},T_q]\cap[1,t-1]=\emptyset$, and hence $P_{q\tau}=0$ for any $\tau\leq t-1$. Moreover, since product $i\neq i_q$ was purchased in period t, we also know that product i_q was not consumed in period t.

Appendix B: Missing Proofs of Section 3

B.1. Proof of Lemma 3

Preliminaries. In Appendix C, we show how to derive the optimal solution to problem SUB-ASSORT in a running time of $O(n^{O(1)}L)$, which we refer to as $(S^t,S^t_{\mathcal{O}},q^t,x^t)$. Throughout this proof, we will almost exclusively manipulate only S^t and $S^t_{\mathcal{O}}$, and hence for ease of notation, we drop every parameter's dependence on x^t and q^t. Moreover, in referencing these latter two parameters as well as θ^{t+1}, we drop their dependence on t, since it plays no role in the proof. [5]

Intermediate claims. Our proof will make us of the following claims, which establish structural properties of the optimal solution $(S^t,S^t_{\mathcal{O}},q,x)$. The proofs of these two claims can be found at the end of the section.

Claim 5. *If $S^t_{\mathcal{O}}\neq\emptyset$, then $p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}}\geq 0$.*

Claim 6. *If $S^t_{\mathcal{O}}\neq\emptyset$, then $R_q(S^t,S^t_{\mathcal{O}})\geq 0$.*

[5] When products are dropped from S^t, we assume that x^t is implicitly updated to reflect this change. Moreover, if we do make a non-trivial update to the prices dictated by x^t, then we explicitly note how these changes effect the various parameters.

Proof of Lemma 3. To begin, observe that if $R_i(S^t, S^t_{\mathcal{O}}, \theta) \geq 0$ for each $i \in S^t$, then the proof is complete, since conditions (i) and (ii) of the lemma statement are clearly met. Otherwise, we make the following simple update. Let $D = \{i \in S^t : R_i(S^t, S^t_{\mathcal{O}}, \theta) < 0\}$ denote all products whose contribution is negative, and consider the assortments $\hat{S}^t = S^t \setminus D$ and $\hat{S}^t_{\mathcal{O}} = S^t_{\mathcal{O}} \setminus D$. If $\hat{S}^t_{\mathcal{O}} = \emptyset$, then it must have been the case that $S^t_{\mathcal{O}} = \emptyset$. To see this, note that Claim 6 implies that $q \notin D$ if $S^t_{\mathcal{O}} \neq \emptyset$, which in turn implies that $q \in \hat{S}^t_{\mathcal{O}}$ and hence $\hat{S}^t_{\mathcal{O}}$ must at least contain q. In this scenario, i.e. $\hat{S}^t_{\mathcal{O}} = \emptyset$, it is straightforward to see that the $\hat{S}^t$ and $\hat{S}^t_{\mathcal{O}} = \emptyset$ will satisfy the two properties of the lemma by using the definition of D. Below we analyze the cases where $|\hat{S}^t_{\mathcal{O}}| = 1$ and $|\hat{S}^t_{\mathcal{O}}| \geq 2$, and make necessary adjustments for $(\hat{S}^t, \hat{S}^t_{\mathcal{O}})$ in each case.

- Case 1 - $|\hat{S}^t_{\mathcal{O}}| = 1$: In this case, we must have that $\hat{S}^t_{\mathcal{O}} = \{q\}$, but since we cannot offer an opaque product of size one, we set $\hat{S}^t_{\mathcal{O}} = \emptyset$. In this case, observe that

$$\begin{aligned}
Z_t(\theta) &= \lambda_{qt} R_q(S^t, S^t_{\mathcal{O}}, \theta) + \sum_{i \in D} \lambda_{it} R_q(S^t, S^t_{\mathcal{O}}, \theta) + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(S^t, S^t_{\mathcal{O}}, \theta) \\
&\leq \lambda_{qt} R_q(S^t, S^t_{\mathcal{O}}, \theta) + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(S^t, S^t_{\mathcal{O}}, \theta) \\
&= \underbrace{\lambda_{qt} R_q(\hat{S}^t, S^t_{\mathcal{O}}, \theta) + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta)}_{\text{term (i)}},
\end{aligned}$$

where the inequality follows by definition of D, and the second equality follows by observing that for each $i \in S^t \setminus (D \cup \{q\})$ we know that $i \notin \hat{S}^t_{\mathcal{O}}$, and hence the contribution of these products is unchanged when the opaque option is removed. From here, we abuse notation and use $\pi_q(p_q, S^t_{\mathcal{O}})$ and $\pi^{\mathcal{O}}_q(p_{\mathcal{O}}, S^t_{\mathcal{O}})$ to respectively denote the choice probability of transparent product q at price p_q and the opaque option at price $p_{\mathcal{O}}$ to get

$$\begin{aligned}
\text{term (i)} &= \lambda_{qt} \cdot \left(\pi_q(p_q, S^t_{\mathcal{O}}) \cdot \left(p_q - \frac{\theta_q}{u_{q1}} \right) + \pi^{\mathcal{O}}_q(p_{\mathcal{O}}, S^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}} \right) \right) \\
&\qquad + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta) \\
&\leq \lambda_{qt} \cdot 2 \left(\max \left\{ \pi_q(p_q, S^t_{\mathcal{O}}) \cdot \left(p_q - \frac{\theta_q}{u_{q1}} \right), \pi^{\mathcal{O}}_q(p_{\mathcal{O}}, S^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}} \right) \right\} \right) \\
&\qquad + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta) \\
&\leq 2\lambda_{qt} \cdot \left(\max \left\{ \pi_q(p_q, \emptyset) \cdot \left(p_q - \frac{\theta_q}{u_{q1}} \right), \pi_q(p_{\mathcal{O}}, \emptyset) \cdot \left(p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}} \right) \right\} \right) \\
&\qquad + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta) \\
&= 2\lambda_{qt} R_q(\hat{S}^t, \emptyset, \theta) + \sum_{i \in S^t \setminus (D \cup \{q\})} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta) \\
&\leq 2 \sum_{i \in \hat{S}^t} \lambda_{it} R_i(\hat{S}^t, \emptyset, \theta),
\end{aligned}$$

as desired for property (ii). The first inequality follows from the fact that $a+b \le 2\max\{a,b\}$ for any $a,b \ge 0$, the second inequality follows from Assumption 2, which implies that $\pi_q(p_q,\emptyset) \ge \pi_q(p_q,S^t_{\mathcal{O}})$ and $\pi_q(p_{\mathcal{O}},\emptyset) \ge \pi^{\mathcal{O}}_q(p_{\mathcal{O}},S^t_{\mathcal{O}})$. Here $\pi_q(p_q,\emptyset) \ge \pi_q(p_q,S^t_{\mathcal{O}})$ is trivial, and $\pi_q(p_{\mathcal{O}},\emptyset) \ge \pi^{\mathcal{O}}_q(p_{\mathcal{O}},S^t_{\mathcal{O}})$ holds since

$$\pi^{\mathcal{O}}_q(p_{\mathcal{O}},S^t_{\mathcal{O}}) = \frac{w^{\mathcal{O}}_q(p_{\mathcal{O}})}{1+w_q(p_q)+w^{\mathcal{O}}_q(p_{\mathcal{O}})} \le \frac{w_q(p_{\mathcal{O}})}{1+w_q(p_q)+w_q(p_{\mathcal{O}})} \le \pi_q(p_{\mathcal{O}},\emptyset),$$

where, the first inequality above holds because $w^{\mathcal{O}}_q(p_{\mathcal{O}}) \le w_q(p_{\mathcal{O}})$ by Assumption 2. For the second equality in bounding term (i), we set the price of the transparent product q to

$$\hat{p}_q = \operatorname*{arg\,max}_{p\in\{p_q,p_{\mathcal{O}}\}} \pi_q(p,\emptyset)\cdot(p-\frac{\theta_q}{u_{q1}})$$

and hence we achieve the max. We also assume this pricing update is reflected in $\hat{x}^t$. The last inequality holds due to the fact that for each $i \in S^t \setminus (D \cup \{q\})$, we know its contribution is non-negative by definition of D.

Finally, we note that it is easy to see that $R_i(\hat{S}^t,\emptyset,\theta) \ge 0$ for each $i \in \hat{S}^t$, as is required for property (i) of the lemma. To see this, note that for $i \in \hat{S}^t \setminus \{q\} = S^t \setminus (D \cup \{q\})$, we have that $R_i(\hat{S}^t,\emptyset,\theta) = R_i(S^t,S^t_{\mathcal{O}},\theta) \ge 0$ by definition of D. Moreover, we clearly have that $R_q(\hat{S}^t,\emptyset,\theta) \ge 0$, since $p_q - \frac{\theta_q}{u_{q1}} \ge p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}} \ge 0$ by Claim 5.

- Case 2: $|\hat{S}^t_{\mathcal{O}}| \ge 2$. In this case, we have

$$\begin{aligned}
Z_t(\theta) &= \sum_{i\in S^t\setminus D} \lambda_{it} R_i(S^t,S^t_{\mathcal{O}},\theta) + \sum_{i\in D} \lambda_{it} R_i(S^t,S^t_{\mathcal{O}},\theta) \\
&\le \sum_{i\in S^t\setminus D} \lambda_{it} R_i(S^t,S^t_{\mathcal{O}},\theta) \\
&= \sum_{i\in \hat{S}^t} \lambda_{it} R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta) \\
&= \underbrace{\sum_{i\in \hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it} R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta) + \sum_{i\in \hat{S}^t\setminus\hat{S}^t_{\mathcal{O}}} \lambda_{it} R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)}_{\text{term (ii)}},
\end{aligned}$$

where the first inequality follows by definition of the set D. The last equality follows from the observation that for any $i \in \hat{S}^t \setminus \hat{S}^t_{\mathcal{O}}$, it must hold that $i \notin S^t_{\mathcal{O}}$. Therefore the contents of $S^t_{\mathcal{O}}$ do not affect $R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta)$ as long as $i \notin S^t_{\mathcal{O}}$, and hence $S^t_{\mathcal{O}}$ can be replaced with any $\hat{S}^t_{\mathcal{O}}$ where $\hat{S}^t_{\mathcal{O}} \subseteq S^t_{\mathcal{O}}$. Next we show that the first of the two summations in term (ii) can be upper bounded as follows, where we let $S_1 = \{i \in \hat{S}^t \cap \hat{S}^t_{\mathcal{O}} : p_i - \frac{\theta_i}{u_{i1}} \ge 0\}$.

$$\begin{aligned}
&\sum_{i\in \hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it} R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta) \\
&\quad= \sum_{i\in \hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it} \cdot \left(\pi_i(S^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}}) + \pi^{\mathcal{O}}_i(S^t_{\mathcal{O}})\cdot(p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}})\right) \\
&\quad= \sum_{i\in \hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it} \cdot \left(\left(\pi_i(S^t_{\mathcal{O}})-\pi_i(\hat{S}^t_{\mathcal{O}})+\pi_i(\hat{S}^t_{\mathcal{O}})\right)\cdot(p_i-\frac{\theta_i}{u_{i1}}) + \pi^{\mathcal{O}}_i(S^t_{\mathcal{O}})\cdot(p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}})\right)
\end{aligned}$$

$$
\begin{aligned}
&\le \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\cdot\Big(\pi_i(S^t_{\mathcal{O}})-\pi_i(\hat{S}^t_{\mathcal{O}})\Big)\cdot(p_i-\frac{\theta_i}{u_{i1}})^+ \\
&\qquad + \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\cdot\left(\pi_i(\hat{S}^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}})+\pi_i^{\mathcal{O}}(S^t_{\mathcal{O}})\cdot(p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}})\right) \\
&\le \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\cdot\pi_i(S^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}})^+ \\
&\qquad + \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\cdot\left(\pi_i(\hat{S}^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}})+\pi_i^{\mathcal{O}}(\hat{S}^t_{\mathcal{O}})\cdot(p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}})\right) \\
&\le 2\max\left\{\sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\pi_i(S^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}})^+,\ \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\right\} \\
&\le 2\max\left\{\sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}\pi_i(\emptyset)\cdot(p_i-\frac{\theta_i}{u_{i1}})^+,\ \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\right\} \\
&= 2\max\left\{\sum_{i\in S_1} \lambda_{it}R_i(S_1,\emptyset,\theta),\ \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}} \lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\right\}.
\end{aligned}
$$

In the above analysis,the first inequality holds because $\pi_i(S^t_{\mathcal{O}})-\pi_i(\hat{S}^t_{\mathcal{O}})\ge 0$, the second inequality holds because $\pi_i^{\mathcal{O}}(S^t_{\mathcal{O}})\le\pi_i^{\mathcal{O}}(\hat{S}^t_{\mathcal{O}})$, the third inequality holds by again invoking $a+b\le 2\max\{a,b\}$, and the last inequality holds because $\pi_i(S^t_{\mathcal{O}})\le\pi_i(\emptyset)$. The final equality holds by the definition of S_1. By plugging in the above upper bound, we show that term(ii) is upper bounded by

$$
\begin{aligned}
\text{term (ii)} &\le 2\max\left\{\sum_{i\in S_1}\lambda_{it}R_i(S_1,\emptyset,\theta),\ \sum_{i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}}\lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\right\}+\sum_{i\in\hat{S}^t\setminus\hat{S}^t_{\mathcal{O}}}\lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta) \\
&\le 2\max\left\{\underbrace{\sum_{i\in S_2}\lambda_{it}R_i(S_2,\emptyset,\theta)}_{\text{term (a)}},\ \underbrace{\sum_{i\in\hat{S}^t}\lambda_{it}R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)}_{\text{term (b)}}\right\} \qquad (14)
\end{aligned}
$$

where $S_2=(\hat{S}^t\setminus\hat{S}^t_{\mathcal{O}})\cup S_1$. From here, we set $(\hat{S}^t,\hat{S}^t_{\mathcal{O}})$ as follows: if term (b) is larger in (14), then we set $\hat{S}^t=S^t\setminus D$ and $\hat{S}^t_{\mathcal{O}}=S^t_{\mathcal{O}}\setminus D$, and if term (a) is larger, we set $\hat{S}^t=S_2$ and $\hat{S}^t_{\mathcal{O}}=\emptyset$.

To conclude, we argue that property (i) of the lemma is satisfied by our choice of $(\hat{S}^t,\hat{S}^t_{\mathcal{O}})$ for Case 2. On one hand, if term (a) is larger, and we set $\hat{S}^t=S_2$ and $\hat{S}^t_{\mathcal{O}}=\emptyset$, then by definition of S_2 we get that $R_i(\hat{S}^t,\emptyset,\theta)\ge 0$ for each $i\in\hat{S}^t$. If, on the other hand, term (b) is larger, we show $R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\ge 0$ for all $i\in\hat{S}^t$. Suppose there is some i such that $R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)<0$. We note that this i must satisfy $i\in\hat{S}^t\cap\hat{S}^t_{\mathcal{O}}$ and $p_i-\frac{\theta_i}{u_{i1}}<0$. However, if this were the case, we get that

$$0>R_i(\hat{S}^t,\hat{S}^t_{\mathcal{O}},\theta)\ge R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta),$$

which leads to a contradiction, since $R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta)\ge 0$ by definition of D. The second inequality follows because $\hat{S}^t_{\mathcal{O}}\subseteq S^t_{\mathcal{O}}$, and so $\pi_i(\hat{S}^t_{\mathcal{O}})\le\pi_i(S^t_{\mathcal{O}})$ and $\pi_i^{\mathcal{O}}(\hat{S}^t_{\mathcal{O}})\ge\pi_i^{\mathcal{O}}(S^t_{\mathcal{O}})$. As a result,

$$R_i(\hat{S}^t,S^t_{\mathcal{O}},\theta)=\pi_i(S^t_{\mathcal{O}})\cdot(p_i-\frac{\theta_i}{u_{i1}})+\pi_i^{\mathcal{O}}(S^t_{\mathcal{O}})\cdot(p_{\mathcal{O}}-\frac{\theta_q}{u_{q1}})$$

$$\leq \pi_i(\hat{S}^t_{\mathcal{O}}) \cdot (p_i - \frac{\theta_i}{u_{i1}}) + \pi^{\mathcal{O}}_i(\hat{S}^t_{\mathcal{O}}) \cdot (p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}})$$
$$= R_i(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \theta).$$

Proof of Claim 5. Assume by contradiction that $p_{\mathcal{O}} - \frac{\theta_q}{u_{q1}} < 0$, and define $N_- = \{i \in S^t : p_i - \frac{\theta_i}{u_{i1}} < 0\}$. We have that

$$\begin{aligned} Z_t(\theta) &= \sum_{i \in N_-} \lambda_{it} R_i(S^t, S^t_{\mathcal{O}}, \theta) + \sum_{i \in S^t \setminus N_-} \lambda_{it} R_i(S^t, S^t_{\mathcal{O}}, \theta) \\ &\leq \sum_{i \in S^t \setminus N_-} \lambda_{it} R_i(S^t, S^t_{\mathcal{O}}, \theta) \\ &< \sum_{i \in S^t \setminus N_-} \lambda_{it} R_i(S^t, \emptyset, \theta) \\ &= \sum_{i \in S^t \setminus N_-} \lambda_{it} R_i(S^t \setminus N_-, \emptyset, \theta). \end{aligned}$$

The strict inequality follows since by setting $S^t_{\mathcal{O}} = \emptyset$, we have removed the negative contribution of the opaque option, while increasing the choice probabilities of all the transparent options, which each by constructions have net positive contributions. The above sequence of inequalities contradicts the optimality of $(S^t, S^t_{\mathcal{O}})$, since they imply that $(S^t \setminus N_-, \emptyset)$ strictly improves on this supposed optimal solution.

Proof of Claim 6. Exploiting Claim 5, it is straightforward to see that $R_q(S^t, S^t_{\mathcal{O}}) < 0$ only if $p_q - \frac{\theta_q}{u_{q1}} < 0$. However, if this latter inequality is true it must be the case that $p_q < p_{\mathcal{O}}$, which contradicts the feasibility of x^t, since transparent products must be priced above the opaque option.

B.2. Proof of Lemma 4

Basic algebra reveals that

$$\theta^t_i - \theta^{t+1}_i = \begin{cases} \lambda_{it} R_i(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t, \theta^{t+1}) + \mathbb{1}_{\hat{q}^t = i} \cdot \left(\sum_{j \neq i} \lambda_{jt} \pi^{\mathcal{O}}_j(\hat{S}^t_{\mathcal{O}}) \cdot (p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta^{t+1}_{\hat{q}^t}}{u_{\hat{q}^t 1}}) \right) & \text{if } i \in \hat{S}^t \\ 0 & \text{otherwise} \end{cases}$$

Due to Lemma 3, and Claim 5, we have that $\theta^t_i - \theta^{t+1}_i \geq 0$ in the first case. The fact that these tuning parameters are non-negative follows from the observation that they are decreasing in t and $\theta^{T+1}_i = 0$.

B.3. Proof of Lemma 5

We begin with an intermediate claim, which is essential to proving the lemma at-hand.

Intermediate claim. Consider the following updated version of the product-i contribution defined in (9), which we now express as a function of the inventory vector U_t in period t

$$\begin{aligned} \hat{R}_i(S, S_{\mathcal{O}}, q, x, \theta; U_t) = \mathbb{1}_{u_{it} > 0} \cdot \Bigg(\pi_i(x, S_{\mathcal{O}}) \cdot \left(p_i(x) - \frac{\theta_i}{u_{i1}} \right) \\ + \mathbb{1}_{u_{qt} > 0} \cdot \pi^i_{\mathcal{O}}(x, S_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(x) - \frac{\theta_q}{u_{q1}} \right) \Bigg). \end{aligned}$$

This updated product-i contribution function reflects the notion that out-of-stock products can be offered, but the newly added indicators ensure that these out-of-stock products have zero contributions. Note that, if problem (SUB-ASSORT) is formulated with these updated product-i contributions, then it might be the case that its optimal objective exceeds that of (10). In other words, there is a subtle difference between enforcing that we can only offer in-stock products (problem (10)), and adding indicators to ensure that we only pick up contributions from in-stock products. To see why this is the case, consider a scenario with two products, where one is out-of-stock. In problem (10), it must be the case that $\bar{S}^t_{\mathcal{O}} = \emptyset$ since the opaque product must contain at least two products, and so the opaque option cannot add to the total contribution. In this scenario, when we consider the indicator version of (SUB-ASSORT), the out-of-stock product can be included so that we can gain a contribution from the opaque option (setting the lone in-stock product to be the give-away product). With this in-mind, we present the following intermediate claim, which serves to relate the two aforementioned problem settings, and whose proof appears at the end of this section.

Claim 7. *For any period $t \in [T]$, we have that*

$$\sum_{i \in \bar{S}_t} \lambda_{it} \hat{R}_i(\bar{S}^t, \bar{S}^t_{\mathcal{O}}, \bar{q}^t, \bar{x}^t, \theta^{t+1}) \geq \frac{1}{2} \sum_{i \in \bar{S}_t} \lambda_{it} \hat{R}_i(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t, \theta^{t+1}; U_t),$$

where $(\bar{S}^t, \bar{S}^t_{\mathcal{O}}, \bar{q}^t, \bar{x}^t)$ is the optimal solution to (10) when the period-t inventory is U_t and $(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t)$ is as defined in Lemma 3.

Proof of lemma. We prove the result by induction over t. The base case of $t = T+1$ holds trivially, and hence we move to the general case case of $t \leq T$. Starting with the induction hypothesis, we have

$$\begin{aligned}
\bar{H}_t(U_t) \geq & \sum_{i \in \bar{S}^t} \lambda_{it} \cdot \left(\pi_i(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_i(\bar{x}^t) - \frac{1}{2}\left(H_{t+1}(U_t) - H_{t+1}(U_t - e_i)\right) \right) \right. \\
& \left. + \pi_i^{\mathcal{O}}(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\bar{x}^t) - \frac{1}{2}\left(H_{t+1}(U_t) - H_{t+1}(U_t - e_{\bar{q}^t})\right) \right) \right) + \frac{1}{2} H_{t+1}(U_t), \\
\geq & \sum_{i \in \bar{S}^t} \lambda_{it} \cdot \left(\pi_i(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_i(\bar{x}^t) - \frac{1}{2} \frac{\theta_i^{t+1}}{u_{i1}} \right) \right. \\
& \left. + \pi_i^{\mathcal{O}}(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\bar{x}^t) - \frac{1}{2} \frac{\theta_{\bar{q}^t}^{t+1}}{u_{\bar{q}^t 1}} \right) \right) + \frac{1}{2} H_{t+1}(U_t), \\
\geq & \sum_{i \in \bar{S}^t} \lambda_{it} \cdot \left(\pi_i(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_i(\bar{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) \right. \\
& \left. + \pi_i^{\mathcal{O}}(\bar{x}^t, \bar{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\bar{x}^t) - \frac{\theta_{\bar{q}^t}^{t+1}}{u_{\bar{q}^t 1}} \right) \right) + \frac{1}{2} H_{t+1}(U_t), \\
\geq & \frac{1}{2} \cdot \left(\sum_{i \in \mathcal{N}} \mathbb{1}_{u_{it} > 0} \cdot \lambda_{it} \cdot \left(\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) \right. \right. \\
& \left. \left. + \sum_{j \in \mathcal{N}} \mathbb{1}_{j = \hat{q}^t} \cdot \mathbb{1}_{u_{jt} > 0} \cdot \pi_i^{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_j^{t+1}}{u_{j1}} \right) \right) + H_{t+1}(U_t) \right),
\end{aligned} \tag{15}$$

where the second inequality follows by (8), the third inequality follows from the fact that the θ-values are non-negative (Lemma 4), and the fourth inequality follows by Claim 7. Continuing the string of inequality, we have

$$(15) \geq \frac{1}{2} \cdot \left(\sum_{i \in \mathcal{N}} \frac{u_{it}}{u_{i1}} \cdot \lambda_{it} \cdot \left(\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) \right. \right.$$

$$
\begin{aligned}
&\qquad + \sum_{j\in\mathcal{N}} \mathbb{1}_{j=\hat{q}^t} \cdot \frac{u_{jt}}{u_{j1}} \cdot \pi_i^{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_{\hat{q}^t}^{t+1}}{u_{\hat{q}^t 1}} \right) \Bigg) + H_{t+1}(U_t) \Bigg) \\
&= \frac{1}{2} \cdot \left(\sum_{i\in\mathcal{N}} \frac{u_{it}}{u_{i1}} \cdot \lambda_{it} \cdot \left(\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) + \pi_i^{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}} \right) \right) \right. \\
&\qquad \left. + \sum_{i\in N} \sum_{j\neq i} \frac{u_{it}}{u_{i1}} \frac{u_{jt}}{u_{j1}} \cdot \mathbb{1}_{j=\hat{q}^t} \cdot \lambda_{it} \pi_i^{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}}) \cdot \left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_{\hat{q}^t}^{t+1}}{u_{\hat{q}^t 1}} \right) \right) + H_{t+1}(U_t) \Bigg) \\
&= \frac{1}{2} \cdot \left(\sum_{i\in\mathcal{N}} \frac{u_{it}}{u_{i1}} \cdot \left(\gamma_i^t - \gamma_i^{t+1} \right) + \sum_{i\in\mathcal{N}} \frac{u_{it}}{u_{i1}} \cdot \left(\gamma_{ii}^t - \gamma_{ii}^{t+1} \right) \right. \\
&\qquad \left. + \sum_{i\in N} \sum_{j\neq i} \frac{u_{it}}{u_{i1}} \frac{u_{jt}}{u_{j1}} \cdot \left(\gamma_{ij}^t - \gamma_{ij}^{t+1} \right) + H_{t+1}(U_t) \right) \\
&= \frac{1}{2} H_t(U_t).
\end{aligned}
$$

The first inequality follows because $\mathbb{1}_{x>0} \geq \frac{x}{y}$ for any $x, y \geq 0$. The third equality uses the definition of the tuning parameters are given in Algorithm 2, and the last inequality uses (7)

Proof of Claim 7. Define $S^t = \hat{S}^t \cap \mathcal{N}(U_t)$ and $S^t_{\mathcal{O}} = \hat{S}^t_{\mathcal{O}} \cap \mathcal{N}(U_t)$. The proof then proceeds by replicating the arguments from the proof of Lemma 3 for the cases of $|S^t_{\mathcal{O}}| = 1$ and $|S^t_{\mathcal{O}}| \neq 1$. For brevity, and to limit redundancies in the arguments we present, we omit the remaining details of the proof.

B.4. Proof of Lemma 6

To begin, observe that for any $A = (S, S_{\mathcal{O}}, q, x) \in \mathcal{A}$ and $t \in [T]$, basic algebra reveals that

$$
\begin{aligned}
\sum_{i\in\mathcal{N}} \lambda_{it} R_i(A, 0) - u(A) \cdot \alpha_i &= \sum_{i\in\mathcal{N}} \pi_i(x, S_{\mathcal{O}}) \cdot (p_i(x) - \alpha_i) + \pi^i_{\mathcal{O}}(x, S_{\mathcal{O}}) \cdot (p_{\mathcal{O}}(x) - \alpha_q) \\
&= \sum_{i\in\mathcal{N}} \lambda_{it} R_i(A, \alpha(u)),
\end{aligned}
$$

where $\alpha(c) = (u_{11}\alpha_1, \ldots, u_{n1}\alpha_n)$, and so can rewrite the optimal dual objective as

$$
\mathrm{OPT}_{\text{fluid}} = \min_{\alpha_i \geq 0} \sum_{i\in\mathcal{N}} u_{i1}\alpha_i + \sum_{t\in[T]} \max_{A\in\mathcal{A}} \sum_{i\in\mathcal{N}} \lambda_{it} R_i(A, \alpha(c)).
$$

From here, consider the dual solution $\alpha_i = \frac{\theta_i^1}{ui1}$, which is clearly feasible since $\theta_i^1 \geq 0$ for each $i \in \mathcal{N}$, and which results in a dual objective value of

$$
\hat{Z}_{\text{dual}} = \underbrace{\sum_{i\in\mathcal{N}} \theta_i^1}_{\text{term (i)}} + \underbrace{\sum_{t\in[T]} \max_{A\in\mathcal{A}} \sum_{i\in\mathcal{N}} \lambda_{it} R_i(A, \theta)}_{\text{term (ii)}}.
$$

By weak duality, we have $\hat{Z}_{\text{dual}} \geq \mathrm{OPT}_{\text{fluid}}$, and hence it suffices to show that $4H_1(U_1) \geq \hat{Z}_{\text{dual}}$. We do so by showing that both term (i) and term (ii) are upper bounded by $2H_1(U_1)$.

Bound for term (i). We have

$$
\begin{aligned}
\sum_{i\in\mathcal{N}} \theta_i^1 &= \sum_{i\in\mathcal{N}} \left(\gamma_i^1 + \gamma_{ii}^1 + \sum_{j\neq i} \gamma_{ij}^1 + \gamma_{ji}^1 \right) \\
&= \sum_{i\in\mathcal{N}} \left(\gamma_i^1 + \gamma_{ii}^1 \right) + \sum_{i\in\mathcal{N}} \sum_{j\neq i} \left(\gamma_{ij}^1 + \gamma_{ji}^1 \right) \\
&\leq \sum_{i\in\mathcal{N}} \left(\gamma_i^1 + \gamma_{ii}^1 \right) + 2 \cdot \sum_{i\in\mathcal{N}} \sum_{j\neq i} \gamma_{ij}^1 \\
&\leq 2H_1(U_1),
\end{aligned}
$$

where the final inequality follows by (7).

Bound for term (ii). We have

$$\begin{aligned}
&\sum_{t\in[T]} \max_{A\in\mathcal{A}} \sum_{i\in\mathcal{N}} \lambda_{it} R_i(A,\theta) \\
&\le 2\cdot \sum_{t\in[T]}\sum_{i\in\mathcal{N}} \lambda_{it} R_i(\hat{S}^t, \hat{S}^t_{\mathcal{O}}, \hat{q}^t, \hat{x}^t, \theta) \\
&= 2\cdot \sum_{t\in[T]}\sum_{i\in\mathcal{N}} \lambda_{it}\cdot \left(\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_i(\hat{x}^t) - \frac{\theta_i^1}{u_{i1}}\right) + \pi^i_{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta^1_{\hat{q}^t}}{u_{\hat{q}^t 1}}\right)\right) \\
&\le 2\cdot \sum_{t\in[T]}\sum_{i\in\mathcal{N}} \lambda_{it}\cdot \left(\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}}\right) + \pi^i_{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta^{t+1}_{\hat{q}^t}}{u_{\hat{q}^t 1}}\right)\right), \qquad (16)
\end{aligned}$$

where the first inequality uses Lemma 3, and the second uses Lemma 4. Using the definitions of the tuning parameters as given in Algorithm 2, we have

$$\begin{aligned}
(16) &= 2\cdot\left(\sum_{t\in[T]}\sum_{i\in\mathcal{N}} \lambda_{it}\pi_i(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_i(\hat{x}^t) - \frac{\theta_i^{t+1}}{u_{i1}}\right)\right. \\
&\qquad\qquad \left. + \sum_{t\in[T]}\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}} \mathbb{1}_{j=q}\cdot\lambda_{it}\pi^i_{\mathcal{O}}(\hat{x}^t, \hat{S}^t_{\mathcal{O}})\cdot\left(p_{\mathcal{O}}(\hat{x}^t) - \frac{\theta_j^{t+1}}{u_{j1}}\right)\right) \\
&= 2\cdot\left(\sum_{t\in[T]}\sum_{i\in\mathcal{N}}(\gamma_i^t - \gamma_i^{t+1}) + \sum_{t\in[T]}\sum_{i\in\mathcal{N}}(\gamma_{ii}^t - \gamma_{ii}^{t+1}) + \sum_{t\in[T]}\sum_{i\in\mathcal{N}}\sum_{j\ne i}(\gamma_{ij}^t - \gamma_{ji}^{t+1})\right) \\
&= 2\cdot\left(\sum_{i\in\mathcal{N}}(\gamma_i^1 + \gamma_{ii}^1) + \sum_{i\in\mathcal{N}}\sum_{j\ne i}(\gamma_{ij}^1 + \gamma_{ji}^1)\right) \\
&= 2H_1(U_1),
\end{aligned}$$

where the third equality follows since these sums telescope.

Appendix C: Finding the Optimal Pricing/Assortment Choice

Our approach for recovering the period-t optimal solution of SUB-ASSORT, referred to as $(S^t, S^t_{\mathcal{O}}, q^t, x^t)$, unfolds over two steps. First, there is a guessing step, where guess components of the optimal solution $(S^t, S^t_{\mathcal{O}}, q^t, x^t)$ that allow for the problem to be decoupled by customer type. Next, exploiting this decoupling, we formulate a simple dynamic program that allows for efficient recover of the remaining components of the optimal solution that were not guessed in the first step.

Step 1: Guessing. To begin, we guess q^t, $|S^t_{\mathcal{O}}|$ and $p_{\mathcal{O}}(x^t)$ (henceforth referred to as $p_{\mathcal{O}}$) using complete enumeration over all candidate combinations. Indeed, there are $O(n)$ candidates for the former two quantities and $O(L)$ for the latter, and hence complete enumeration requires considering $O(n^2L)$ possible guesses. Importantly, with both $|S^t_{\mathcal{O}}|$ and $p_{\mathcal{O}}$ in-hand, we know $w_i^{\mathcal{O}}(x^t, S^t_{\mathcal{O}})$ for each customer type $i \in \mathcal{N}$. As such, for the remainder of this section, we use $w_i^{\mathcal{O}}$ as a shorthand for $w_i^{\mathcal{O}}(x^t, S^t_{\mathcal{O}})$.

Step 2: The dynamic program. Next, we present a dynamic program that sequentially progresses over the customer types, and for each, decides whether or not to include their focal product as a transparent product (while also choosing its price in this case) and/or whether to include this focal product within the opaque option. The value functions $R(i,k)$ denote the optimal "reward" that can garnered from types

$i,\ldots,n$, given that k products among $\{1,\ldots,i-1\}$ have been included within the opaque option. Formally, we have

$$R(i,k) = \max\Bigg\{ \underbrace{R(i+1,k)}_{\text{do not offer } i}, \underbrace{\max_{\ell\in[L]} \lambda_{it}\cdot\left(p_\ell - \frac{\theta_i}{u_{i1}}\right)\cdot\frac{w_{i\ell}}{1+w_{i\ell}} + R(i+1,k)}_{\text{Only offer } i \text{ as transparent}}, \underbrace{\max_{\substack{\ell\in[L]:\\ p_\ell\ge p_{\mathcal{O}}}} \lambda_{it}\cdot\left(p_\ell - \frac{\theta_i}{u_{i1}}\right)\cdot\frac{w_{i\ell}}{1+w_{i\ell}+w_i^{\mathcal{O}}} + \lambda_{it}\cdot\left(p_{\mathcal{O}} - \frac{\theta_{q^t}}{u_{q^t1}}\right)\cdot\frac{w_i^{\mathcal{O}}}{1+w_{i\ell}+w_i^{\mathcal{O}}} + R(i+1,k+1)}_{\text{Offer } i \text{ as transparent and within opaque}} \Bigg\},$$

with base cased of

$$R(n+1,k) = \begin{cases} 0 & \text{if } k = |S^t_{\mathcal{O}}| \\ -\infty & \text{otherwise} \end{cases}$$

to enforce that the final opaque option abides by our guess of its cardinality. We must also enforce that when $i = q^t$, we offer q^t as both a transparent and opaque product. It is straightforward to see that following the decisions of this dynamic program from the initial state of $(1,0)$ will yield $S^t, S^t_{\mathcal{O}}$ and x^t.

The overall running time. There are a total of $O(n^2)$ values functions, which can each be computed in a running time of $O(L)$ by enumerating over the L possible pricing options. Consequently, including the initial guessing step, the final running to time to compute $(S^t, S^t_{\mathcal{O}}, q^t, x^t)$ is $O(n^4L^2)$.

Appendix D: Expedia data description

The Expedia data set contains customers' search and purchase data on various Expedia-owned websites in different countries. For each logged customer search, the data records the search's time stamp, the searched destination, hotel listings returned for the search including a rich set of search-dependent and hotel-dependent features, and the customers' response actions including click and purchase decisions. The search-dependent features include those that only depend on the trip information of a customer, for example, the duration of stay, the number of adults, and the number of children, and they do not vary across hotels. The hotel-dependent features include those that only depend on the hotel and do not vary with customer searches, for example, hotels' star ratings, review scores, and location scores. Both types of features clearly have an impact on the final purchase decisions of customers.

We follow the approach of Feldman and Segev (2022) to clean the data for our experiments. According to their work, a subset of the data entries are deemed as outliers and are therefore omitted. These entries mainly include those with unrealistic hotel prices (less than \$10 or higher than \$5000 USD per room per night) and those with a duration of stay greater than 10 days. We also only include searches whose results are displayed in random order. For a complete set of criteria applied to the data cleaning procedure, we refer readers to Appendix C.2.1 of Feldman and Segev (2022). We provide a brief summary of the key statistics of the cleaned data in Table 6 for the three biggest sites. We summarize the total number of searches, the total number of bookings, the conversion rate (or the purchase probability, calculated as the ratio of the number of bookings to the number of searches), and the average number of hotels returned in a search ($|S|$) on the three sites (5, 14, 15).

Table 6 Summary Statistics

Site	Searches	Bookings	Conversion	$\|S\|$
5	63383	9045	14.27%	25.1
14	10544	1401	13.29%	24.6
15	6674	861	12.90%	24.3