

A Data-Driven Approach to Personalized Bundle Pricing and Recommendation

Markus Ettl, Pavithra Harsha

TJ Watson Research Center, 1101 Kitchawan Rd, Yorktown Heights, NY, 10598, msettl@us.ibm.com, pharsha@us.ibm.com

Anna Papush, Georgia Perakis

Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA, 02139, apapush@mit.edu, georgiap@mit.edu

Problem Definition: The growing trend in online shopping has sparked the development of increasingly more sophisticated product recommendation systems. We construct a model that recommends a personalized discounted product bundle to an online shopper that considers the trade-off between profit maximization and inventory management, while selecting products that are relevant to the consumer's preferences. **Methodology:** We focus on simultaneously balancing personalization through individualized functions of consumer propensity-to-buy, inventory management for long-run profitability, and tractability for practical business implementation. We develop two classes of approximation algorithms, multiplicative and additive, in order to produce a real-time output for use in an online setting.

Academic and Practical Relevance: We provide analytical performance guarantees that illustrate the complexity of the underlying problem, which combines assortment optimization with pricing. We implement our algorithms in two separate case studies on actual data from a large U.S. e-tailer and a premier global airline. **Results:** Our computational results demonstrate significant lifts in expected revenues over current industry pricing strategies on the order of 2-7% depending on the setting. We find that on average our best algorithm obtains 92% of the expected revenue of a full-knowledge clairvoyant strategy across all inventory settings, and in the best cases this improves to 98%. **Managerial Implications:** We compare the algorithms and find that the multiplicative approach is relatively easier to implement and on average empirically obtains expected revenues within 1-6% of the additive methods when both are compared to a full-knowledge strategy. Furthermore, we find that the greatest expected gains in revenue come from high-end consumers with lower price sensitivities and that predicted improvements in sales volume are dependent on product category and they are a result of providing relevant recommendations.

Key words: pricing and revenue management, retailing, OM Practice, inventory theory and control, dynamic programming

1. Introduction

The online market as a whole has grown enormously over the past decade. According to a forecast released by Forrester Research in 2016, U.S. e-commerce retail sales are expected to grow from \$373 billion in 2016 to more than \$500 billion in 2020, an increase of over 34%; furthermore, the online sector alone impacts over \$1.5 trillion of total retail sales in the United States. This surge in online shopping has led to an increased availability of data regarding consumer preferences, which

can be leveraged by businesses across all industries in order to improve operations, revenue and consumer satisfaction. As a direct example of this, strategic recommendation services that utilize such data effectively are currently undergoing massive and rapid expansion. StitchFix, which offers personalized clothing styling for its customers and only became cash-flow positive in 2014, achieved \$250M in revenue in 2015 that nearly tripled to \$730M by the end of 2016. The travel industry has also undergone dramatic growth over the past decade due to the increasing availability of online services. As a result, travel products are becoming increasingly commoditized. Consumers are not willing to pay exorbitant fees for these generic services, resulting in a great deal of competition across industries such as airlines. However, travelers are both interested in and willing to pay for customized experiences. Therefore, businesses in the travel and hospitality industries have begun providing recommendations at the time of reservation by offering ancillary services that are unbundled from ticket cost or room rate. In the case of airlines, supplementary products to improve the travelers' experience before, during and after their ticketed trip such as VIP lounge access, priority boarding, seat upgrades, in-flight wi-fi, and destination-relevant deals are now provided at their own prices and offered throughout the online purchase process, which was previously not the case. As a result of these industry trends in the online sector, the development of a more sophisticated product recommendation system can provide the necessary competitive edge for any online seller, making the difference on the order of millions in profits.

As demonstrated by these industry examples, the majority of businesses with an online component now utilize recommendation systems. However, these methods are often primarily based on historical purchase trends across segments of the online population when there is also a wealth of individualized consumer information. Motivated by this increasingly prevalent cross-selling problem and current industry practices, the goal of our work is the development of a personalized model that selects, prices, and recommends a bundle of related products to a consumer during their online session. Having dynamically received this offer while browsing a particular item or ticket itinerary, the consumer can then choose to accept this discounted offer, or purchase any combination of items at their full prices, or simply exit the online marketplace without making any purchase at all. This dynamic bundle offer is constructed using a new model that combines diverse recommendations with personalized discounts by leveraging consumer profiles and in-session context, while considering the trade-off between myopic current profit with long-term profitability under inventory constraints. Note that because we consider the possibility that consumers may choose to purchase products at their full prices, we must additionally incorporate the upper-level problem of determining the time-dependent trajectories of full prices over the course of the selling horizon. Thus, the novelty of this work consists of *simultaneously* incorporating personalization,

bundle assortment selection, bi-level pricing, and inventory-balancing within this particular online bundle offer setting. These challenges have not yet been explored jointly in the existing literature.

In order to construct our bundle pricing and recommendation model we focus on incorporating all of the above components simultaneously. We aim to make relevant offers by solving a personalized bundle assortment selection and pricing problem that uses individualized propensity-to-buy models based on consumer profiles and online context. We integrate this personalized online offer setting within the goal of long-run profitability by additionally considering future demand through an inventory balancing function in our model, which improves expected profits by mitigating costs associated with overstocking and lost sales. Balancing all of these factors is novel to the analytical problem and practically crucial to sellers from an operational perspective, but also gives rise to several challenges with respect to both the analytical problem structure and its implementation. The combination of all of these components results in a complex dynamic programming problem that is highly intractable in an online setting. Furthermore, focusing on inventory-constrained products leads to the additional difficulty of incorporating upper-level dynamic pricing schemes that affect the full prices of products as the selling horizon progresses. Thus, our resulting model simultaneously addresses personalization, multiple levels of pricing, bundle assortment selection, demand forecasting, and inventory management. We develop approximation algorithms and provide analytical guarantees that improve in tightness as the problem becomes less inventory constrained. Furthermore, we analyze the performance of our algorithms through two case studies: (i) using point-of-sale transaction data from a major U.S. e-tailer that includes personalized features such as customer IDs and loyalty information, and, (ii) using ticket transaction data from a premier global airline that includes consumer-specific information such as tier level, miles balance and previous flight history at the time of ticket purchase. These case studies demonstrate that our approaches result in significant improvement in expected revenue over existing industry practices. In industries that operate on razor-thin margins, these gains can scale up to several millions of dollars in revenue.

1.1. Contributions

We analyze the problem of personalized online bundle recommendation, which lies at the intersection of several branches of revenue management literature. Our main contributions consist of,

1. **Two classes of approximation algorithms that provide real-time bundle recommendations and simultaneously incorporate personalization, inventory balancing and tractability.** We develop multiplicative and additive methods to implement our model in real-time in an online setting. These heuristics capture personalization as well as the trade-off between myopic profit maximization and long-run profitability under inventory constraints. We also coordinate the dynamic lower-level personalized bundle prices with an upper-level global pricing strategy that periodically determines the time-dependent trajectories of the full prices of all items.

2. Analytical guarantee on the performance of the multiplicative algorithm and empirical comparisons of both classes. We provide a bound on the ratio of the expected revenue of the multiplicative approach relative to a full-knowledge strategy that knows the entire consumer arrival sequence in advance. This becomes even tighter as the problem becomes less constrained by inventory and falls on average within 15% of the algorithm’s actual empirical performance ratio on data. We further compare the empirical performance of both algorithm classes and show that on average, the overall best heuristic is an additive benchmark that obtains 90-98% of the expected revenue of the full-information benchmark across various initial inventory settings. Furthermore, we show that the multiplicative approach is **easier to implement compared to the additive methods** and on average obtains an expected revenue that is within 1-6% of that achieved by additive methods, relative to the full-knowledge strategy.

3. Two detailed case studies on actual data from the retail and airline travel industries that demonstrate significant improvement in expected revenue on the order of 2-7% on average over existing practices depending on the setting. In the retail case, we observed that our algorithms provide output in real-time with predicted gains of up to 14% in revenue over current pricing schemes in the most unconstrained discounting settings. In the airline case, our model predicted improvements in sales volume and revenue as high as 7-8% over current strategies in settings when a fraction of the online population is unaware of the existence of ancillary services. The greatest gains in expected revenue were a result of *personalized pricing* targeted at consumers with lower price sensitivities, who are easily incentivized to make additional purchases through smaller personalized discounts. Conversely, the largest growth in predicted sales volume was dependent on product category and primarily a result of *relevant recommendations*, resulting in lifts on the order of up to 10% over current practices.

1.2. Literature Review

We consider two bodies of literature most closely related to our work: constrained assortment optimization and dynamic pricing. The first line of literature pertains to the assortment planning problem under capacity constraints. Initial works such as Mahajan and van Ryzin (1999) consider a single-period stochastic model under which the retailer selects a profit-maximizing assortment of substitutable products and determines their initial stock prior to the selling period, under the assumption that consumers choose products according to a multinomial logit model, which was extended in Mahajan and van Ryzin (2001) to incorporate dynamic substitution effects when a consumer’s product of choice may be stocked out. Kök et al. (2009) presents a summary of the initial works on the single-period assortment planning setting under inventory or budget constraints, which captures extensions to other consumer choice models and various dynamic substitution

effects. Later works such as Aouad et al. (2015) provide provably efficient algorithms under stochastic demand and dynamic substitution and show that these approximations are order optimal, or near-optimal as in Goyal et al. (2016), under general random-utility choice models. Gallego and Topaloglu (2014) consider both assortment cardinality and display space constraints, showing that the assortment problem is efficiently solvable while the space problem is NP hard. Another body of works has built on the one-period problem by considering dynamic assortment optimization that is solved distinctly for each consumer arrival. Rusmevichientong and Topaloglu (2012) and Rusmevichientong et al. (2014) study this problem when the parameters of the consumer demand functions are unknown. In Bernstein et al. (2015) they formulate this as a dynamic program to identify which optimal assortment of substitutable goods to offer each consumer and develop inventory threshold policies for determining this. By contrast, Golrezaei et al. (2014) propose an index-based inventory balancing approach for determining the optimal personalized assortments, which motivates our multiplicative algorithm that extends this setting by also incorporating pricing. Topaloglu (2009) and Jaillet and Lu (2012) present more generalized approaches to inventory balancing and demonstrate the value of duality-based approximations of DP formulations in the context of the network revenue management problem and dynamic resource allocation problems, respectively. In Gallego et al. (2015) they develop an asymptotically optimal policy for this dynamic setting, and in Li et al. (2015) and Chen and Jiang (2017) they consider further extensions under the d-nested logit choice model and the MNLD choice model, respectively. In this work, we consider an inventory-constrained assortment planning problem to dynamically determine the composition of personalized product bundles. However, the bundle recommendation system we propose sets this work apart from the existing literature in constrained assortment planning primarily because we extend the problem to incorporate dynamic pricing. Furthermore, we do not limit our analysis to any specific consumer choice model, nor do we assume that the assortment consists of only substitutable goods.

The second body of work is related to dynamic pricing and cross-selling. Dynamic pricing literature, which initially focused primarily on single products, is very well summarized in Bitran and Caldentey (2003). However, there is a vast body of more recent literature on this topic across a wide variety of consumer utility models and considering multiple products. In Agrawal et al. (2014) and Shamsi et al. (2014) the authors consider the online resource allocation problem under different settings and develop efficient algorithms for maximizing long-term system revenue based on dual price updates. By contrast, Kesselheim et al. (2014) develop an algorithm based on a scaled version of the partially known primal problem (as opposed to dual prices) to obtain and round fractional solutions and develop integral allocations for all requests. Tying online pricing to learning with a multi-armed bandit framework, Ferreira et al. (2016) develop an algorithm based on

Thompson sampling for dynamically pricing multiple products under inventory constraint in order to maximize long-run profitability. By contrast to all of these works, dynamic cross-selling grew as its own field from economics and initially did not incorporate pricing. Cavusgil and Zou (1994) provides an overview of the growth and expansion of this early literature. Netessine et al. (2006) is a key pivotal work that combined these fields by analyzing models with stochastic arrivals for the joint problem of cross-selling and pricing in which a consumer has one primary product of interest and is offered one additional complementary product at a discounted price for both; in Rapti et al. (2014) the authors extend this dynamic programming setting by proposing a rule-based approach to the joint bundle selection and pricing problem. New directions such as Xue et al. (2015) consider request for quote (RFQ) models where consumers interactively participate in the pricing process. Recent work in Gallego et al. (2016) presents a dynamic approach to product pricing and framing to determine optimal product displays on webpages for consumers. We consider a setting in which the bundle offer is presented as an additional option that the consumer can choose not to purchase in favor of any other combination of non-discounted products. Thus, since all products are also available for purchase at their full prices, we must consider the upper-level problem of determining the time-dependent full product prices over the course of the selling horizon. Thus, our work addresses a two-level pricing problem and aligns: (i) the lower-level personalized bundle prices offered dynamically to each consumer, with, (ii) the upper-level full price trajectories for each product. To the best of our knowledge, this simultaneous bi-level pricing problem is not addressed in the cross-selling literature.

2. Problem Setting and Model Formulation

We consider a monopolist online seller that makes a dynamic bundle offer to each arriving consumer who may choose to accept the offer, purchase individual items separately at full price, or choose to purchase nothing at all, as shown in Figure 1 below.




Figure 1 This is an example of a personalized bundle recommendation when a consumer is offered the set A, B and C at a 20% discount, otherwise they can choose from any combination of A, B, C and D at full price.

If the consumer chooses to purchase either the bundle or some other collection of items at their full prices, we assume that they only purchase one unit of each item. Let us consider a set of items $i = 1, \dots, n$ denoted by \hat{S} . These items' prices may affect one another and they can be complementary, substitutable, or even independent as is often the case in the travel industry. Given a captive online consumer considering products within \hat{S} , or a specific ticket itinerary for which \hat{S} is the set of ancillary goods, our model offers a relevant bundle of products from \hat{S} . We are

interested in cases where \hat{S} contains inventory-constrained products that we leverage to maximize expected long-run profitability by accounting for future demand. Therefore, we consider a finite selling horizon with a fixed number of periods T with no replenishments. Each arriving consumer is uniquely described by a combination of categorical and continuous features related to preferences, demographics, purchase history, loyalty, and online shopping context. Thus, we do not consider a discrete set of consumer types as is traditionally done in segmentation and instead assume that there is an infinite set of continuous consumer types. Furthermore, since we address a bi-level pricing problem, we index consumers within a given period t by (k, t) , where $k = 1, \dots, K^t$ and the total number of arrivals K^t in each period can differ. We define the full price of item i in period t as \bar{p}_i^t ; thus, the full price $\bar{p}_{S_{k,t}}$ of a bundle $S_{k,t}$ offered to consumer (k, t) is defined by,

$$\bar{p}_{S_{k,t}} = \sum_{i \in S_{k,t}} \bar{p}_i^t \quad (1)$$

The full prices \bar{p}_i^t are not necessarily fixed throughout the horizon and may follow some dynamic trajectory, summarized in each period by vector $\bar{\mathbf{p}}^t = [\bar{p}_1^t, \bar{p}_2^t, \dots, \bar{p}_n^t]$. We thus define price vector,

$$\mathbf{p}_{S_{k,t}} = [\bar{\mathbf{p}}^t, p_{S_{k,t}}] = [\bar{p}_1^t, \bar{p}_2^t, \dots, \bar{p}_n^t, p_{S_{k,t}}], \quad (2)$$

in which we append the discounted price of the personalized bundle for consumer (k, t) to the vector of full price settings for period t . It is common in business practice for sellers to consider discrete price ladders. Therefore, we make the assumption that we have a fixed set of price levels for every product i from which we can choose to construct bundle offers. We define the individual consumer propensity-to-buy $\xi_S^{k,t}(\mathbf{p}_{S_{k,t}})$ as the probability that consumer (k, t) will purchase the combination of products S when their personalized bundle $S_{k,t}$ is offered at price $p_{S_{k,t}}$ and all products $i \notin S_{k,t}$ are offered at their *full prices*, where it is possible that $S_{k,t} \subset S$. We will refer to $\mathbf{e}_{S_{k,t}}$ as the bundle unit vector that takes the value 1 for all $i \in S_{k,t}$ and 0 otherwise. Finally, we define $\mathbf{I}^{k,t}$ as the vector of inventory levels of all $i \in \hat{S}$ at the time when consumer (k, t) arrives, written explicitly as $\mathbf{I}^{k,t} = [I_1^{k,t}, I_2^{k,t}, \dots, I_n^{k,t}]$. This leads to the following decision variables for any given consumer (k, t) : the optimal bundle to recommend $S_{k,t} \in \hat{S}$, and, its personalized price $p_{S_{k,t}} \leq \bar{p}_{S_{k,t}}$.

2.1. Dynamic Programming Formulation

We formulate this personalized bundle offer problem ideally as a dynamic program, as is traditional in the revenue management literature. This results in a complex model that is difficult to solve, as we discuss below in Section 2.2. This DP approach leads us to the following formulation (3), defined by $\{\text{Dynamic}\}_{\forall(k,t)}$:

$$\begin{aligned} & \underset{S_{k,t} \subset \hat{S}, p_{S_{k,t}}, \bar{p}_i^t}{\text{maximize}} && V_{k,t}(\mathbf{I}^{k,t}) \\ & \text{subject to} && V_{k,t}(\mathbf{I}^{k,t}) = \sum_{S \subset \hat{S}} \xi_S^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \left(\left(\sum_{i \in S} \bar{p}_i^t \right) + (p_{S_{k,t}} - \bar{p}_{S_{k,t}}) \cdot \mathbb{1}_{\{S_{k,t} \subset S\}} + V_{k+1,t}(\mathbf{I}^{k,t} - \mathbf{e}_S) \right) \quad (3) \\ & && V_{K^t+1,t}(\mathbf{I}) = V_{1,t+1}(\mathbf{I}) \quad \forall K^t, t = 1, \dots, T \\ & && (1 - \epsilon)\bar{p}_{S_{k,t}} \leq p_{S_{k,t}} \leq \bar{p}_{S_{k,t}} \quad \forall(k, t), S_{k,t} \subset \hat{S} \end{aligned}$$

We solve this problem for every consumer $k = 1, \dots, K^t$ who arrives within each period $t = 1, \dots, T$ and connect the periods t through the forward-looking inventory cost-to-go functions $V(\cdot)$. Note that in addition to the offer for each consumer, the full price trajectories \bar{p}_i^t over all products in all periods are also decision variables in this model. However we note here that unlike the consumer-level bundle offer decisions, these prices are calculated periodically in an upper-level pricing problem at the conclusion of each period t , then held fixed for that period and updated with the most currently inventory for period $t + 1$.

The objective function consists of several terms: the first term captures the probability $\xi_S^{k,t}(\mathbf{p}_{S_{k,t}})$ with which a consumer (k, t) purchases some set of products $S \subset \hat{S}$, summed over all possible sets S (note that this captures the cases in which S is a superset that encompasses the personalized offer $S_{k,t}$); the second set of terms account for the expected revenue from consumers purchasing individual items at full price, as well as from accepting the bundle offer (at which point the discounted bundle at full price $\bar{p}_{S_{k,t}}$ is removed from the summation of \bar{p}_i^t). Note that we do not include the probability of a “no-buy” because this is innately captured in the set of collections of products $S \subset \hat{S}$ that also includes the null set \emptyset , corresponding to the consumer’s decision to make no purchase. This is a complex dynamic programming problem because it relies on knowledge of future demand and inventory levels to utilize $V_{k+1,t}(\cdot)$ in making bundle offers. The first constraint accounts for the recursive transition of the inventory revenue-to-go function $V(\cdot)$ between periods and the second constraint limits the depth of the bundle discount $\bar{p}_{S_{k,t}}$ and ensures that the bundle offers remain attractive. This DP formulation is intractable for the online setting due to the forward-looking nature of the functions $V(\cdot)$ and the need to calculate the full prices \bar{p}_i^t for all products in all periods.

If we were given the full price trajectories \bar{p}_i^t for every i , along with the values for the functions $V(\cdot)$ at all possible inventory levels and bundle combinations, then solving formulation (3) would be an enumeration over all the discrete prices and bundle combinations. We remark that even if provided with all of these values, this enumeration problem could potentially suffer from the curse of dimensionality. However, appropriate limitations to ensure a small size of \hat{S} would reasonably bound the number of bundle combinations and thus make this problem tractable. Therefore, if we were provided with the full price trajectories \bar{p}_i^t and all the precise values of $V(\cdot)$ (which become a constant independent of the current period’s decisions in the absence of inventory constraints), we could use the above model to make optimal individually tailored offers of bundles $S_{k,t}$ at prices $p_{S_{k,t}}$ for each consumer in a tractable manner.

2.2. Challenges

While we are interested in optimally solving $\{\text{Dynamic}\}_{V(k,t)}$ for every consumer in real-time, realistically the consumer arrival sequence, full price trajectories, and values of $V(\cdot)$ are unknown. This results in three fairly sizable challenges: (i) how to estimate a personalized propensity-to-buy, (ii) how to determine the upper-level prices \bar{p}_i^t and align them with offers $(S_{k,t}, p_{S_{k,t}})$, and, (iii) how to estimate the values of $V(\cdot)$ while maintaining tractability in an online setting.

Developing a personalized bundle recommendation at an individually tailored price requires the most granular possible estimate of a consumer's propensity-to-buy. Traditional methods lack distinctive information that distinguishes a customer from others in their segment. Therefore, to address (i), we use machine learning methods to fit high-dimensional models that capture all of these features through covariates as described in detail in Section 4. Considering an inventory-constrained problem with a finite horizon, results in challenge (ii) of determining and incorporating an upper-level pricing strategy into our model that alters the full prices of individual products over time. Thus, we propose a method (described in Section 3.1) for determining these price trajectories within our problem framework as follows: at the beginning of each t we calculate the full prices \bar{p}_i^t across all products i and fix them for the duration of that period, after which we update them at the beginning of the next period $t+1$ using current inventory levels after consumer demand is realized. This rolling approach coordinates the upper-level full price trajectories with the lower-level bundle offers (which are based on the values of \bar{p}_i^t) made to individual consumers within a given period. Finally, we address (iii) by developing various approximation approaches to the forward-looking inventory balancing functions $V(\cdot)$. A common linear programming approximation, in which we solved a series of LPs to estimate the values of $V(\cdot)$ at various inventory levels (as is commonly done in the revenue management literature, see Talluri and van Ryzin (2006)), runs far too slowly. Thus, we propose two classes of approximation algorithms in Section 3, multiplicative and additive methods, that are practically tractable and therefore applicable in an online setting.

2.3. Clairvoyant Formulation

Before presenting any approximation algorithms, we first generalize the dynamic programming problem to a "full-knowledge" model to establish a benchmark against which we can compare any algorithm's performance. We assume that the entire consumer arrival sequence $\{k, t\}_{\forall k=1, \dots, K^t}^{t=1, \dots, T}$ is known in advance, as well as the full price trajectories \bar{p}_i^t for all products i in all periods t , which

we assume are provided to us by an oracle. In order to model this perfect information setting we propose the following formulation that we refer to as the {Clairvoyant} problem:

$$\begin{aligned}
& \underset{y_{S_{k,t}}^{k,t}}{\text{maximize}} && \sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{S_{k,t} \subset \hat{S}} \left(\left[\sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \right] + \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (p_{S_{k,t}} - \bar{p}_{S_{k,t}}) \right) \cdot y_{S_{k,t}}^{k,t} \\
& \text{subject to} && \sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{S_{k,t} \subset \hat{S}} \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \right] \cdot y_{S_{k,t}}^{k,t} \leq I_i^0 \quad \forall i \\
& && \sum_{S_{k,t} \subset \hat{S}} y_{S_{k,t}}^{k,t} = 1 \quad \forall (k,t) \\
& && y_{S_{k,t}}^{k,t} \geq 0 \quad \forall (k,t), S_{k,t} \subset \hat{S}
\end{aligned} \tag{4}$$

The decision variables $y_{S_{k,t}, p_{S_{k,t}}}^{k,t}$ ¹ correspond to the probability with which bundle $S_{k,t}$ is offered at price $p_{S_{k,t}}$ to consumer (k,t) when the full product prices are set at \bar{p}_i^t . The discrete price setting allows us to relax these initially binary decisions to continuous variables, resulting in the above linear programming formulation. For this formulation we define the individual consumer propensity-to-buy $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$ as the probability that consumer (k,t) will purchase item i if their personalized bundle $S_{k,t}$ is offered at price $p_{S_{k,t}}$. Unlike the prior exclusive definitions of $\xi(\cdot)$ for formulation (3), these propensities-to-buy are defined as follows: $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$ captures all of the combinations in which product i can be purchased when only the bundle offer is discounted and all other products remain at full price \bar{p}_i^t . For example, if consumer (k,t) is offered bundle $S_{k,t}$ at price $p_{S_{k,t}}$, we can define the probability they purchase product i in some combination of other products S at full price as $\phi_S^{k,t}(\bar{p}_S^t, p_{S_{k,t}})$. Thus, the complete probability of (k,t) purchasing i is given by,

$$\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) = \sum_{S \subset \hat{S}: S \ni i} \phi_S^{k,t}(\bar{p}_S^t, p_{S_{k,t}}) \tag{5}$$

We similarly define the probability $\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}})$ with which a personalized bundle $S_{k,t}$ is purchased as,

$$\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) = \sum_{S \subset \hat{S}: S \supset S_{k,t}} \phi_S^{k,t}(\mathbf{p}_{S_{k,t}}) \tag{6}$$

Note that the above probability includes the scenarios in which the consumer purchases full-priced items in addition to the bundle $S_{k,t}$, meaning that their purchase set S contains $S_{k,t}$. Thus, formulation (4) is the offline version of formulation (3) in which the entire sample path of consumer arrivals and full price trajectories are both known before the start of the horizon. Therefore, the objective function is an expectation of the total revenue taken over the consumer purchase decisions for a specific known sample path $\{k,t\}_{\forall k=1, \dots, K^t}^{t=1, \dots, T}$, and is thus an upper bound on the expected revenue of any online algorithm. Notice that this problem is also subject to inventory constraints defined through initial stock levels I_i^0 for all products $i = 1, \dots, n$, because it allocates bundle offers according to expected consumer behavior over the entire horizon. By definition of the Clairvoyant,

¹ We define $y_{S_{k,t}, p_{S_{k,t}}}^{k,t}$ as being dependent on both the bundle composition and price. However, under discrete pricing, we can enumerate the collection \hat{S} of all bundles at all prices so each $S_{k,t}$ is inherently defined by its corresponding price. Thus, for ease of notation we neglect the additional subscript of $p_{S_{k,t}}$ and write the summations over $\forall S_{k,t} \subset \hat{S}$.

this problem eliminates the need for forward-looking inventory functions because it identifies the optimal bundle $S_{k,t}$ for every consumer (k,t) utilizing its full knowledge of all future arrivals. This benchmark is not actually attainable because it relies on precise future knowledge that is never available to any practical online model. However, since this formulation provides an upper bound on the expected profit for our setting, we will use its objective value $\{\text{Clairvoyant}\}$ as an “optimal” best-case benchmark against which we measure the performance of all proposed algorithms.

3. Approximation Algorithms

The primary source of complexity in our model stems from the calculation of the expected future revenue as a function of the inventory levels. Therefore, our main goal in this section is to develop methods that approximate the $V(\cdot)$ terms in formulation (3). We also aim to address the second challenge that arises from the consideration of products with limited stock, which is the incorporation of inventory-based dynamic pricing strategies that optimize the full product prices \bar{p}_i^t over the course of the selling horizon. Thus, we also aim to develop approximation algorithms within a framework that aligns: (i) the bundle offer pricing in our lower-level recommendation system for individual consumers within each period t , with, (ii) the global upper-level full price trajectories.

3.1. Multiplicative Approximation Algorithm

We first consider the following approach to our lower-level bundle recommendation problem, which incorporates the value of inventory through a multiplicative penalty on the bundle terms from the objective function of $\{\text{Dynamic}\}_{\forall(k,t)}$. This multiplicative penalty can be viewed as an approximation to the negative counterpart of the dual variables corresponding to the inventory constraints in formulation (4) of the Clairvoyant problem. Using this multiplicative approach allows us to maintain the previous trade-offs captured in the objective function of the DP problem in formulation (3), but include inventory balancing through a tractable calculation that does not require demand forecasting. In formulation (7) we present the general formulation for this multiplicative approximation algorithm, denoted by $\{\text{MultAlg}\}_{\forall(k,t)}$, which requires the full price trajectories \bar{p}_i^t as inputs (the procedure for computing the values of \bar{p}_i^t is below).

$$\begin{aligned} \{\text{MultAlg}\}_{\forall(k,t)} = & \underset{S_{k,t} \subset \hat{S}, p_{S_{k,t}}}{\text{maximize}} && \left[\sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] + \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (p_{S_{k,t}} - \bar{p}_{S_{k,t}}) \cdot \min_{i \in S_{k,t}} \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \\ & \text{subject to} && (1 - \epsilon) \bar{p}_{S_{k,t}} \leq p_{S_{k,t}} \leq \bar{p}_{S_{k,t}} \quad \forall(k,t), S_{k,t} \in \hat{S} \end{aligned} \quad (7)$$

The above formulation has a similar structure to the objective function of $\{\text{Clairvoyant}\}$. However, the key difference between this algorithm and formulation (4) lies in the use of the multiplicative inventory penalty $\psi(\cdot)$ to determine bundle composition and pricing using an approach that requires no estimation of future consumer behavior, as was previously the case in the use of the $V(\cdot)$ functions in formulation (3). This penalty $\psi(\cdot)$ is a twice-differentiable, monotone increasing, concave

function on the interval $[0, 1]$ and takes as input the fraction of remaining inventory $I_i^{k,t}/I_i^0$ at the time of arrival of consumer (k, t) . We consider several different forms for this function including linear, $\psi(x) = x$; polynomial, $\psi(x) = \sqrt{x}$; and exponential, $\psi(x) = (1 - e^{-x})$. In this work we consider the joint problem of bundle composition and pricing, and therefore introduce the minimization of $\psi(\cdot)$ over all i in bundle $S_{k,t}$, which penalizes bundles composed of items with low stock and instead promotes products with excess inventory. We also introduce a corresponding set of $\psi(\cdot)$ functions to the individual product purchases to approximately account for the corresponding $V(\cdot)$ functions that would influence those terms in the original dynamic programming formulation.² By avoiding any demand forecasting, this approach reduces the recommendation problem to an enumeration over all possible bundles and prices, which is typically small in scale given limitations on the size of \hat{S} . However, note that the choice of functional form for the multiplicative penalty is important to the implementation of this algorithm. We show in Section 4.3 that a choice of polynomial $\psi(\cdot)$, which is often a good approach when approximating a function whose true form is unknown, may result in lower empirical performance on the order of up to 6% when compared to the performance of a more sophisticated exponential $\psi(\cdot)$, relative to the full-knowledge Clairvoyant strategy. Nevertheless, we also find that our results are fairly robust to the choice of $\psi(\cdot)$ and still capture 81% of the expected profit of the full-knowledge approach, even in the most inventory-constrained cases.

We further improve on this algorithm by introducing a rolling extension that augments formulation (7) to use $\psi(I_i^{k,t}/\max\{I_i^t - 1, 1\})$ instead of $\psi(I_i^{k,t}/I_i^0)$ for all consumers $k = 1, \dots, K^t$ arriving during period t . This periodic approach increasingly emphasizes the difference between products with highly depleted stock and those with great excess as the horizon progresses. Note that we use $I_i^t - 1$ as opposed to I_i^t . Intuitively, if we set the denominator of $\psi(\cdot)$ to I_i^t at the start of each period t , the inventory levels of all available products initialize to 100% and become equivalent in terms of $\psi(\cdot)$ for first-arriving consumers ($k = 1, t$). Therefore, subtracting one unit consistently differentiates the fractions $I_i^{k,t}/\max\{I_i^t - 1, 1\}$. This extension is equally tractable and has improved empirical performance over the approach in formulation (7), as shown in Section 4.3.

Calculating Upper-Level Full Price Trajectories

To address our second challenge of aligning upper and lower level pricing strategies, note that the multiplicative algorithm only requires individual full product prices \bar{p}_i^t to make bundle offers. Due to lack of demand forecasting, formulation (7) cannot adequately calculate these full price

² If we directly replaced each $V(\cdot)$ term from the DP in formulation (3) with a minimization over $\psi(\cdot)$ we would get: $\sum_{S \subset \hat{S}} \xi_S^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (\sum_{i \in S} \bar{p}_i^t) \cdot \min_{i \in S} \psi(I_i^{k,t}/I_i^0)$ instead of the first term in formulation (7). However, we instead consider a good upper bound on this term without the minimization, which provides us with this approximation algorithm.

trajectories. Thus, we propose an upper-level method for calculating the time-dependent full price trajectories, denoted by \hat{p}_i^t as they are now estimated quantities, using the following formulation:

$$\begin{aligned} \max_{\hat{\mathbf{p}}^t, \forall t} \quad & \sum_{t=1}^T \sum_{i=1}^n D_i^t(\hat{\mathbf{p}}^t) \cdot \hat{p}_i^t \\ \text{subject to} \quad & \sum_{t=1}^T D_i^t(\hat{\mathbf{p}}^t) \leq I_i^0 \quad \forall i \end{aligned} \quad (8)$$

We define $D_i^t(\hat{\mathbf{p}}^t)$ as the expected demand for product i during period t based on all current product prices $\hat{\mathbf{p}}^t$, which can be calculated using expected future consumer arrival rates based on historical transactions. Based on the prior definition of $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$ from Eq. (5), $D_i^t(\hat{p}_i^t)$ is the expected aggregate demand for product i and incorporates all combinations S in which i is purchased along with other products at full price. Thus, formulation (8) is a tractable linear programming problem, as shown in Talluri and van Ryzin (2006). As is commonly done in practice, we implement this using a rolling approach by periodically re-solving the above LP at the beginning of each period t using updated inventory levels. The output of this upper-level problem provides us with a set of full price trajectories \hat{p}_i^t for all products i in all periods t . By holding these fixed for the duration of a given period t , we can now easily solve the lower-level bundle recommendation problem using the multiplicative algorithm.

3.1.1. Performance Ratios of the Multiplicative Approximation Algorithm The strength of the multiplicative approximation algorithm lies in the fact that it only assumes broad conditions on the structure of $\psi(\cdot)$ and $\phi(\cdot)$. This eliminates the need for demand forecasting and is thus applicable to the majority of possible demand groups \hat{S} and models $\phi(\cdot)$. However, note that the full price trajectories utilized by our algorithm from the upper-level problem may differ significantly from those selected by the full-knowledge Clairvoyant strategy. Therefore, we define:

$$\alpha_i^t = \frac{\hat{p}_i^t}{\bar{p}_i^0} \quad \forall i, t, \quad \text{where } \hat{p}_i^t \text{ are full price trajectories chosen by formulation (8), and,}$$

$$\beta_i^t = \frac{\bar{p}_i^t}{\bar{p}_i^0} \quad \forall i, t, \quad \text{where } \bar{p}_i^t \text{ are optimal price trajectories provided by an oracle to the Clairvoyant in formulation (4).} \quad (9)$$

In this multi-period setting, for a given sample path $\{k, t\}_{\forall k=1, \dots, K^t}^{t=1, \dots, T}$, we show the following result.

THEOREM 1. *Given a fixed adversarial sequence of consumer arrivals (k, t) and time-dependent trajectories of full product prices \hat{p}_i^t from formulation (8) defined through α_i^t for all products i and periods t , the worst-case competitive ratio of our multiplicative algorithm $\{\text{MultAlg}\}_{\forall(k,t)}$ relative to the full-knowledge strategy of $\{\text{Clairvoyant}\}$ when choosing personalized bundle composition and prices **as well as** the global full prices \bar{p}_i^t is bounded by,*

$$1 \geq \frac{\{\text{MultAlg}\}_{\forall(k,t)}}{\{\text{Clairvoyant}\}} \geq \min_{(I_{\min}^0, x): x \leq 1 - \frac{1}{I_{\min}^0}} \frac{\frac{l}{I_{\max}^0} \sum_{t=1}^T \sum_{k=1}^{K^t} R_{\min}^{k,t} - \frac{l}{I_{\min}^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\beta_{\max} \cdot \frac{l}{I_{\min}^0} \cdot \int_{x=1+\frac{1}{I_{\min}^0}}^1 \psi(y) dy + 1 - \psi(x) - \frac{l}{I_{\max}^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}$$

The parameters are explicitly defined as follows: I_{\min}^0 is the minimum initial inventory level across all products $i \in \hat{S}$, defined by $I_{\min}^0 = \min_{i=1, \dots, n} I_i^0$, I_{\max}^0 is similarly the maximum across all initial

inventory levels, $R_{min}^{k,t}$ is the minimum of the product of propensity-to-buy $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$ with the nominal price discount level α_i^t across all $i \in \hat{S}$ for consumer (k, t) , defined by $R_{min}^{k,t} = \min_i \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t$, M^t is defined as the maximum revenue loss from bundle discounting in period t that is defined explicitly in Proposition 1 of Appendix A as $M^t = \max_{S_{k,t} \subset \hat{S}, d_{S_{k,t}}} \sum_{k=1}^{K^t} \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\bar{\mathbf{p}}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (1 - d_{S_{k,t}})$, where $d_{S_{k,t}}$ is the bundle discount price ratio $p_{S_{k,t}}/\bar{p}_{S_{k,t}}$.

The proof of this theorem is provided in detail in Appendix A. Notice that if we were to remove pricing from formulation (7) the resulting problem would identify the most relevant bundle of products to offer, while still allowing consumers to purchase any other combination $S \neq S_{k,t}$ of products where all products (including the bundle $S_{k,t}$) are at full price. Similarly, if we were to remove bundling but continue pricing, the problem would reduce to identifying the best single-item discount offer from among products $i \in \hat{S}$, while allowing the consumer to purchase any other products at their full price. Thus, in the combined setting in which both bundling and pricing are removed from the problem, the above bound reduces to the result in Golrezaei et al. (2014). We empirically evaluate the performance of this ratio in realistic scenarios by using actual data from our case studies to generate the results in Figure 2 below.




Figure 2 This plot shows the empirical value of the bound as a function of the ratio between lowest initial inventory I_{min}^0 and the highest level setting I_{max}^0 across all products, when all else is held constant.

Notice that this bound depends on the choice of the inventory penalty $\psi(\cdot)$ and will vary depending on its functional form. Furthermore, the initial stock levels I_i^0 dictate the extent to which the problem is constrained by inventory, and thus the extremity of the lower bound. Remark that the larger the gap between I_{min}^0 and I_{max}^0 , the more conservative the lower bound, as shown in Figure 2. Intuitively, a less inventory-constrained problem with higher initial inventory settings will result in a significantly tighter bound as there is inherently less error in the multiplicative approach relative to the full-knowledge strategy due to the fact that the consideration of future consumer behavior becomes less critical. Overall, we find that more inventory-constrained instances with limited discounting opportunity due to initially low full price settings generate the most extremely conservative values of the above bound. However, as we demonstrate in Table 7 in Section 4.3, the empirical performance of the bound on the multiplicative algorithm on actual data falls within 7% of its actual performance relative to a full-knowledge strategy across all possible inventory cases.

We conduct an in-depth computational analysis with various functional forms for this multiplicative penalty algorithm, the results of which are summarized in Table 4 in Section 3.2 and Table

5 in Section 4.3 using real industry data from our case studies. We find that the empirically best multiplicative algorithm utilizes an exponential form for the penalty function and obtains 70% of the expected revenue achieved by the full-knowledge strategy in highly constrained inventory cases, which improves to 97% in the less constrained cases. We will refer to this approach as the exponential multiplicative penalty algorithm (EMPA) for the remainder of the work. By implementing the rolling version of the EMPA, we can improve these results to 72% and 97.4%, respectively.

3.2. Additive Approximation Methods

The multiplicative algorithm provides a tractable approach with analytical guarantees that does not need to account for demand forecasting. Therefore, in this section we develop a second class of approximation algorithms to use as benchmarks in order to evaluate the empirical value of considering future consumer behavior. Based on well-known methods from the DP literature such as problem decomposition and Lagrangian relaxation, for example as presented in Hawkins (2003), we construct two additive approaches to approximating the inventory functions $V(\cdot)$ in the lower-level bundle pricing and selection problem. While these approaches incorporate demand forecasting and are therefore more accurate, they suffer from the curse of dimensionality and rely on periodically re-optimized solutions of the upper-level LP problem given by formulation (8). Thus, additive methods are not as efficient to implement in real-time as the multiplicative approach, which is only marginally less efficacious in performance. In Section 4.3 we show that on average the multiplicative algorithm empirically performs within 1-6% of these additive approximations when compared to the full-information benchmark. Thus, the multiplicative approach provides a much easier implementation method at a very marginal cost in terms of expected revenue.

3.2.1. Separable Item Additive Algorithm (SIAA) We aim to estimate the functions $V(\cdot)$ efficiently by decomposing the expected future revenue function $V(\mathbf{I})$ of a given inventory state \mathbf{I} into the sum of the expected future revenues $f_i(I_i)$ for each of the items in demand group \hat{S} . Thus, we propose the following separable-by-item approximation:

$$V_{k+1,t}(\mathbf{I}) \approx \sum_{i \in \hat{S}} f_i^{k+1,t}(I_i), \text{ where,} \tag{10}$$

$$f_i^{k+1,t}(I_i) = \sum_{\tau=t}^T \bar{p}_i^\tau \cdot \min\{\delta^\tau D_i^\tau(\bar{p}_i^\tau), C_i^\tau\}, \text{ and } C_i^\tau = \{C_i^{\tau-1} - \delta^\tau D_i^\tau(\bar{p}_i^\tau)\}^+, \text{ initialized at } C_i^t = I_i.$$

The values of $D_i^t(\hat{p}_i^t)$ and \hat{p}_i^t in this expression are provided by output of the periodically re-solved upper-level LP problem defined by formulation (8) in Section 3.1. We define δ^t as the fraction of time remaining in the period t during which consumer (k, t) arrives, implying that $\delta^\tau = 1$ for all periods τ after the current one. Each of the terms in $f_i^{k+1,t}(I_i)$ from the second line of Eq. (10) considers the minimum between the expected demand for product i in that period and its expected available inventory. We capture this by defining inventory levels C_i^τ recursively for all periods τ

after the current one and initializing the inventory level at I_i . Thus, Eq. (10) provides us with an estimate of expected revenue $f_i^{k+1,t}(I_i)$ from product i over the remainder of the horizon given current inventory I_i . Leveraging this estimation for every product i in the approximation approach from the first line of Eq. (10), we develop a tractable algorithm that depends only on current inventory levels and fixed quantities that are known entirely in advance. We define this method as the SIAA, Separable-Item Additive Algorithm, which allows us to provide each individual customer with a personalized bundle offer in real-time at significantly less computational cost than the original dynamic programming formulation.

3.2.2. Additive Lagrangian Algorithm (ALA) Based on the above framework we now propose a more sophisticated approach, which we refer to as the Additive Lagrangian Heuristic (ALA). If we consider the SIAA more closely, we observe that the separable-by-item decomposition omits any terms related to the expected revenue from bundle purchases at possible bundle discounts. Therefore, we construct a second additive algorithm that incorporates these additional terms. Recalling the approximation framework from Eq. (10), we propose the following extension to include bundle purchases:

$$V_{k+1,t}(\mathbf{I}) \approx \sum_{S \subset \hat{S}} f_S^{k+1,t}(\mathbf{I}), \text{ where,} \quad (11)$$

$$f_S^{k+1,t}(\mathbf{I}) = \sum_{\tau=t}^T \bar{p}_S^\tau \cdot \min \left\{ \delta^\tau D_S^\tau(\bar{p}_S^\tau), \min_{i \in S} \{C_i^\tau\} \right\}, \text{ where } C_i^\tau = \left\{ C_i^{\tau-1} - \delta^\tau \left(\sum_{S \subset \hat{S}: S \ni i} D_S^\tau(\bar{p}_S^\tau) \right) \right\}^+$$

The extension here is the additional consideration of bundle terms at bundle prices, as captured by $f_S^{k+1,t}(\mathbf{I})$. Instead of solving the LP problem from formulation (8), which only outputs single-item trajectories and demand, we formulate an extension in which \bar{p}_S^t are also decision variables. To improve tractability, we further relax this problem by introducing the inventory constraints into the objective function using Lagrange multipliers and ultimately obtain an LP formulation extension of (8) that provides the distinct sets of trajectories \bar{p}_i^t , for individual products i , and \bar{p}_S^t , for combinations S . (In this extended setting $D_S(\bar{p}_S^t)$ is the bundle demand exclusively for bundle S .) Thus, for any consumer we can estimate the cost-to-go $V_{k+1,t}(\mathbf{I})$ using Eq. (11), which includes the expected revenue from both individual items *and* bundle purchases, accounting for the additional revenue not captured by the SIAA due to bundle discounts. This approximation is similarly tractable and in Section 4.3 we show that the ALA achieves empirical results on the order of up to 5% higher than the SIAA in average expected revenue relative to the full-knowledge strategy. Ultimately, we empirically demonstrate in Section 4 that on average these additive approximations capture 92-98% of the optimal Clairvoyant strategy across a range of inventory constrained instances. Therefore, these approaches serve as strong benchmarks against which we can measure the performance of the multiplicative algorithm as opposed to comparing only to the Clairvoyant strategy, which is unattainable in realistic business practice.

4. Industry Driven Case Studies and Computational Results

To test the performance of our proposed algorithms we conducted two extensive case studies using data sets from industry partners. In Section 4.1 we analyze online ticket transactions from a premier airline with a set of ancillary goods offered in addition to the ticket itself. Our models generated an expected increase in revenue and sales volume to be on the order of 2-7% over existing practices depending on the setting. Section 4.2 presents the analysis of data from a two year selling horizon of a large online e-tailer. We benchmark against their actual pricing strategies and observe expected revenue gains up to 14% in the most unconstrained discounting scenarios. Given that both industries operate on tight margins, these results are promising for practical business implementation. We conclude with an in-depth comparison of the two classes of approximation algorithms in Section 4.3 to demonstrate the trade-off between approximation accuracy and tractability.

4.1. Airline Case Study

According to the International Air Transport Association (IATA), the airline industry doubled in revenue from \$369B in 2004 to approximately \$746B in 2014. As a result, travel products such as airline tickets are becoming commoditized and therefore airlines must price competitively and operate on razor-thin margins. Despite this, consumers are increasingly willing to pay for unique experiences, which airlines have begun offering in the form of ancillary services that will customize and improve the journey for the traveler. These ancillary goods are offered to consumers at ticket purchase and customize their journey before, during and after the flight through services such as: seat selection and upgrades, VIP lounge access, in-flight wi-fi access, priority baggage handling and various destination-related deals. Offering these services as personalized recommendations for potential passengers can greatly increase traveler intimacy during their journey and improve satisfaction with the airline. Thus, in the context of this industry, our goal is to make personalized bundle offers consisting of ancillary goods that complement the ticket itinerary under consideration. Since airlines employ various revenue management methods for setting their ticket prices, we consider the ticket price to be externally fixed and aim to target a captive customer with a customized and discounted bundle of ancillary products.

4.1.1. Overview of Data, Modeling and Simulation Design We analyzed a one-month period of approximately 640,000 ticket transactions from a premier international airline. There are no repeat consumers in this short time frame and thus no details from previously purchased flight itineraries. Every transaction is described by a set of features categorized into two types: (i) personal consumer information including tier level, mileage balance, time since joining rewards, and number of previous business and economy flights taken; and, (ii) contextual itinerary booking data that includes transaction date, fare paid (USD), connection time, time to departure, day of

travel, and number of passengers. We used the purchase history information for the following ancillary products: in-flight wi-fi access, premium on-board entertainment, priority security, priority boarding, priority baggage handling, seat upgrades, checked excess baggage, VIP lounge access, gourmet in-flight meals, and 2,000 or 4,000 bonus miles. We were provided with the corresponding historical prices for these products that varied across flight itineraries and from which were able to estimate elasticities. These ancillary product prices are summarized in Table 8 in Appendix B.1. Note that in this data set the products are independent by definition since they correspond to distinct unrelated products that are neither substitutable nor complementary and are priced separately. Given all of this personalized information, our goal is to use the approximation algorithms to make personalized bundle offers consisting of relevant ancillary products for every consumer in the historical arrival sequence.

Developing Consumer Profiles and Demand Estimation

We used k-means clustering to analyze the personalized features in the data and develop distinct consumer profiles that we used to map the historical transactions for demand model estimations. We constructed 7 unique consumer personas ranging from premium business travelers to leisure individuals, the distribution of which is shown in Figure 9 in Appendix B.1. The clusters were developed using a combination of personalized consumer features and itinerary context features. For example, a premium business traveler is a single passenger with short time until departure on a week day, higher tier level, higher fare, and previous premium flights. By contrast, a leisure individual traveler is a single passenger with a lower tier level and longer time before departure on a Thursday, Friday or Saturday. We fit a model for every (persona profile, product) pair, as described in detail in Appendix B.1. The independence property between products resulted in defining the purchase probability of any bundle $\phi_S(\cdot)$ as the product of the purchase probabilities $\phi_i(\cdot)$ of all the products i in that bundle S . We estimated these models using logistic regressions on (persona profile, product) pairs and ultimately produced an exhaustive set of logistic MNL models, which had an out of sample weighted mean absolute percent error (WMAPE) of 0.12 on average across personas. We found that the coefficients for personalized features such as tier level, previous flights and miles balance were both significant and strongly positive across all (persona, product) pairs, indicating that personalization plays an important role in individualized propensities-to-buy. Note that in traditional segment-level models every consumer profile has the same *fixed* propensity-to-buy for any given product. By contrast, we automatically map each consumer to a persona profile for which we built a complete set of personalized (persona, product) models that are populated with this consumer's unique features to produce an *individualized* propensity-to-buy for any product.

Simulation Design

We designed a simulation to test our model and observe the effects of personalized pricing, product

recommendation and inventory management on expected revenues. We analyzed two settings for making bundle offers: (1) under unconstrained inventory as in the data set, and, (2) injecting inventory constraints through reasonable choices of ancillary products that could have restricted quantities. We considered each consumer in the sequence and used our offline clustering approach to automatically map this traveler to one of the persona profiles. Then, the recommendation model solves the lower-level bundle offer problem using the corresponding set of individualized demand models for all available ancillary products to decide which optimal bundle to offer.³ The model then recommends this bundle to the traveler who chooses to accept the offer, or purchase any combination of the ancillary products at their full prices, or make no purchase at all. We repeat this over the 640,000 historical consumer arrivals for 5,000 iterations and analyze the average performance ratios of our model relative to benchmark methods. In both settings we benchmarked our model against the baseline approach of current practice in which *no personalized pricing is offered* and all of the ancillary products are available at their full prices. Under the most current state of the art practice, the full prices of all ancillary products are now held constant throughout the selling horizon. Therefore, we do not solve the upper-level pricing problem in this case study and instead consider the full prices of all products to be fixed at the values provided in Table 8. Note that these prices are inputs to the online bundle pricing problem; therefore, treating them as fixed over the selling horizon does not impact our ability to solve the bundle pricing and recommendation problem.

4.1.2. Results The results are organized according to the two settings described above. Under (1) we analyze the effects of personalized bundle offers, then introduce the concept of lost sales to further enhance the impact of relevant product recommendations when consumers are unaware of the existence of ancillary products. We explain the design of (2) but discuss results in Section 4.3.

Value of Personalized Bundle Offers and Business Insights

In the first setting our initial goal was to analyze the effects of personalized pricing and the recommendation system. We considered a *baseline method that offers all of the ancillary products at their fixed full prices* from Table 8 in Appendix B.1. We implemented our model and observed the average expected gain in revenue over the baseline for all (persona, product) pairs, as summarized in Table 1 below. Furthermore, in this simulation scenario our model produced bundle offer outputs in 2.5ms on average. Note that under unconstrained inventory there is no need for the $V(\cdot)$ functions so our model reduces to the unconstrained Clairvoyant problem and selects a personalized myopic profit-maximizing bundle for each customer. Thus, the predicted improvements over the baseline

³ We note here that having analyzed the propensities-to-buy for all possible bundles $S \in \hat{S}$ we ultimately present results in which, for tractability, we chose to limit our bundles to at most three items due to the fact that all bundles over this size were practically negligible with respect to average contribution to expected revenue.

are a direct result of *personalized bundle offers*. From Table 1 we observe that on average the predicted gains in revenue over the baseline varied from 2% to 7% depending on the persona or ancillary product. The overall largest predicted relative improvements in revenues are generated by consumers with low price elasticities such as premium business and high end leisure travelers. Small discounts targeted at these frequent high loyalty consumers result in significantly more conversions and therefore the most expected revenue.

Average Expected Lift in Revenue Over No-Pricing Baseline

	Wi-fi Access	Prem. Entertain.	Prior. Secur.	Prior. Board.	Prior. Bag Handl.	Seat Up-grade	Excess Checked Baggage	VIP Lounge Access	In-Flight Meals	2,000 Bonus Miles	4,000 Bonus Miles	Total Avg
BusTravPREM	10.1%	3.2%	—	—	—	—	—	—	—	8.6%	3.5%	2.3%
BusTravECON	5.3%	2.3%	1.4%	2.6%	4.4%	8.7%	2.1%	9.7%	3.6%	4.3%	2.5%	4.3%
FamilyGroup	4.2%	2.5%	3.7%	4.0%	1.9%	3.7%	1.9%	5.2%	6.4%	3.9%	2.7%	3.6%
LastMinGroup	2.2%	5.0%	3.9%	2.5%	2.1%	3.1%	2.4%	4.3%	5.4%	4.5%	2.5%	3.4%
CoupleNormal	7.1%	4.2%	3.3%	2.1%	3.4%	2.5%	2.6%	3.3%	3.6%	4.5%	3.0%	3.6%
CoupleHighEnd	8.8%	7.3%	3.8%	2.1%	4.0%	5.3%	3.4%	2.9%	3.8%	7.4%	4.1%	4.8%
LeisNormal	6.5%	2.6%	2.4%	1.8%	2.6%	4.6%	3.6%	3.2%	3.8%	4.4%	2.8%	3.5%
LeisHighEnd	9.8%	8.0%	2.8%	1.9%	3.8%	9.1%	4.3%	6.5%	3.7%	8.8%	5.3%	5.8%
TotalAvg	6.7%	4.4%	2.7%	2.1%	2.8%	4.6%	2.5%	4.4%	3.8%	5.8%	3.3%	3.9%

Table 1 This table shows the lift in revenue from implementing our personalized pricing and recommendation model over the baseline benchmark in which all products are offered to all arriving consumers at full prices. (Note: some products are not offered to premium business travelers because they are included in their tier level benefits.)

By analyzing the corresponding counterpart table of gains in sales volume across all (persona, product) pairs relative to the baseline, we find product-dependent effects. High elasticity personas such as families and lower end leisure travelers see the greatest expected gains in sales volumes across cheaper travel convenience products such as priority security, boarding, and baggage handling, as well as in-flight meals. Conversely, higher-end personas have the highest predicted sales volume gains for more luxe products relevant to frequent travelers, such as VIP lounge access and bonus miles. Intuitively, cheaper products such as in-flight wi-fi have a relatively low elasticity and grow unanimously in predicted revenue across all persona types, particularly among personas with less price sensitivity that are easily converted with slightly discounted offers. These insights provide potential marketing and pricing strategies that could improve revenues and sales volume if used in the right combinations for (persona, product) pairs.

Impact of Context on Personalization

We also assess our model’s ability to distinguish between changes in personalized consumer features and itinerary contexts by analyzing the differences in the average offers made in the following two scenarios: (i) considering the same customer booking two different itineraries, and, (ii) considering two different customers interested in the same itinerary.

In scenario (i), due to lack of purchase history, we generate repeat consumers with similar personal features but different ticket itinerary contexts. The resulting pair of vectors has relatively

constant personal features such as tier level and miles balance, but itinerary features such as ticket fare, day of departure and time to departure vary. We find that our model recommends bundles with different compositions but similar discounts. For example in one such simulation, it categorizes the first context as a business trip and offers in-flight wi-fi and lounge access at a 5.2% discount, while it recognizes the second trip as leisure and offers in-flight entertainment and seat upgrades at a 6.1% discount. Since the personalized consumer features were held constant, the consumer's price elasticities stayed relatively constant over time and therefore the discount remained similar across this scenario. Thus, the primary benefit in this setting comes from the model's ability to *identify the significance of itinerary context* in the absence of major changes in personal features.

Under scenario (ii) we discover the converse effect. For example, we can compare one customer of high tier level with historical premium flights to another passenger with lower tier level traveling in a group. The first customer is offered in-flight wi-fi and 2,000 bonus miles at a 1.8% discount, whereas the second customer is offered in-flight meals and priority boarding at a 6.7% discount. The first customer has low price elasticities across all products (on average below -1) and is recommended business-related products, whereas the family traveler has much higher price elasticities (on average between -2 and -3) and receives a greater discount on products convenient for travel with a group. This demonstrates that the model not only identifies context-relevant items but also maximizes expected profit through personalized pricing.

Value of Relevant Product Recommendations

We also objectively analyze the enhanced effect of product recommendation by introducing a parameter α , which is defined as the proportion of consumers who are unaware of the existence of ancillary products and hence do not consider them at all. This is quite common in the travel industry, such as in cruise lines, where there is often an abundance of products that are not explicitly offered to consumers during their online browsing process resulting in loss of potentially interested consumers. Notice that $\alpha = 0$ corresponds to the setting in which all consumers are aware of all ancillary products and there are no lost sales, which is precisely the previous setting from Table 1. Thus, the results for any fixed setting across varying levels of α explicitly quantify the expected improvement from *relevant product recommendations*. We consider the *same baseline method as before* and implement our myopic profit maximizing recommendation model without inventory constraints. We summarize our predicted lifts in revenue over the baseline in Table 2 and note that in this expanded setting with the lost sales included, our model still produced output in under 3ms.

Each of the columns in Table 2 corresponds to a simulation setting in which we have imposed limitations on our recommendation model. For example, the column "20% Max" corresponds to comparing our recommendation model to the baseline in the case where no product in the bundle

offer is discounted by more than 20%; this definition similarly extends to the columns “No Constraints”, “15% Max”, and “10% Max”. In these scenarios we reasonably limit discounts for all products and observe that the predicted improvement in revenue over the baseline is on the order of 3-8% across all possible cases of lost sales, captured by the varying values of α . We found a similar trend in the corresponding results for expected lifts in sales volume on the order of 2-3%. The column “Discount by Item” imposes limitations depending on the full prices of the products; for example, cheaper products are only discounted up to 10%, but more expensive ones are discounted up to 15-20%. The “Bundle-Only” column is the case where the consumer is offered an optimal bundle by our model, but they can only purchase *any other subset of the bundle at full price*. This corresponds to the realistic setting where there is a vast number of ancillary products and the consumer only considers those displayed to them at checkout. Interestingly, the expected improvement in this case is comparable to the “15% Max” scenario because a consumer’s propensity-to-buy is typically highest for the set of products selected by our model; therefore, disregarding the products outside this relevant bundle does not heavily impact overall expected sales volume or revenue. Thus, we conclude that in reasonable discount-limiting scenarios the expected gain in revenue from our personalized pricing strategy is on the order of 5-6%. By definition, higher α values indicate that a greater proportion of the population is unaware of ancillary products.

Average Expected Lift in Revenue Over No-Pricing Baseline Method

α -level	No Constraints	20% Max	15% Max	10% Max	Discount by Item	“Bundle-Only”
$\alpha = 0$	8.04%	6.11%	4.21%	3.51%	4.83%	4.15%
$\alpha = 0.05$	9.48%	6.09%	5.27%	3.49%	6.18%	5.74%
$\alpha = 0.10$	13.23%	6.23%	5.30%	3.80%	5.66%	6.05%
$\alpha = 0.15$	14.92%	7.74%	5.89%	3.66%	6.89%	6.24%
$\alpha = 0.20$	18.83%	7.24%	6.64%	5.49%	7.94%	7.61%

Table 2 This table summarizes the lifts in revenue over the no-pricing benchmark in various scenarios of lost sales (α) ranging from 0 to 20%. When we compare across varying α levels we see the benefit of product recommendation, and as we compare across a fixed α we see the expected improvement from personalized pricing.

The results in Table 2 are robust to changes in α and by analyzing the symmetric results for sales volume we observe that these trends are consistent across both metrics. Therefore, the predicted improvements on the order of 2-3% over the baseline from the lowest $\alpha = 0$ level to the highest $\alpha = 0.2$ level are a direct result of exposing consumers to products of which they are otherwise unaware through *personalized and relevant product recommendations*.

Value of Inventory

We lastly consider setting (2) to assess the validity of the approximation algorithms presented in Section 3. While inventory is not inherent to this data set, we consider a subset of ancillary products that would reasonably be limited such as VIP lounge access, on-board wi-fi, gourmet meals, excess checked baggage and seat upgrades. We introduce initial inventory levels at quantities that are proportional to the length of the consumer arrival sequence. The simulation is identical

to setting (1), except that we consider an inventory-constrained problem across this smaller set of ancillary products. Instead of solving a myopic personalized profit-maximizing problem, we now solve our original DP problem using both classes of algorithms and observe the average percentage of Clairvoyant revenue they obtain. However, we still do not solve the upper-level problem of time-dependent full price trajectories due to the nature of the data and current industry practice. We present our results and a detailed discussion comparing the algorithms in Table 5 in Section 4.3.

4.2. Retail Case Study

In this second case study we analyzed data from a major U.S. e-tailer over the two year sales period from July 2011 to September 2013. We were provided with point-of-sale transaction data for electronic fulfillment orders across 312 departments, totaling approximately 13M customers and over 34M transactions. The data consisted of order information defined by customer IDs, transaction IDs, SKU numbers, prices and costs at time of purchase, dates and times of purchases, in addition to corresponding inventory data for the same period across the electronic fulfillment centers (EFCs) responsible for these online orders. In the context of this case study our goal was to use our algorithms to make personalized bundle offers to all of the consumers in the arrival sequence and evaluate the average performance of our methods against reasonable industry benchmarks. Note that in this case study the products are not independent and actual inventory levels are known. We believe that the resulting computational study additionally demonstrates the robust performance of our various algorithms.

4.2.1. Overview of Data, Modeling and Simulation Design While this data set contained information regarding inventory levels and time-dependent full price trajectories it lacked personalized features, such as those available in the airline case, outside of customer and transaction IDs. Thus, we had no individualized consumer information outside of historical purchases. To analyze the performance of our methods, we developed personalization metrics with which we could estimate individualized models of propensity-to-buy that provided the necessary basis for the implementation of our proposed algorithms.

Developing Personalization Metrics, Choosing Demand Groups & Demand Estimation

Using the only available personalized consumer features of customer and transaction IDs, we analyzed the data set and recorded each consumer's cumulative number of visits at the time of a given transaction, along with their corresponding total cumulative expenditure up until that time. By studying these metrics across all consumers, we developed a time-dependent loyalty mapping consisting of three distinct categories. Each consumer's transaction in the arrival sequence was automatically mapped into one of the following: (i) the low frequency group with no previous purchase history that accounted for 77% of all transactions; (ii) the medium frequency group with

at least one prior purchase but current cumulative expenditure below the population mean, which accounted for 17.5% of all transactions; or, (iii) the high frequency group with at least one historical purchase, but current cumulative spending over the mean population amount, which accounted for 5.4% of all transactions. The spending behavior of the medium and high frequency loyalty groups over time is visualized in Figures 3 and 4 in Appendix B.2. As an example, any consumer's second transaction would be mapped to at least the medium frequency group because their cumulative number of prior visits is greater than 0. Note that we develop time-dependent metrics instead of simply assigning every consumer statically to one loyalty category, so that our personalized demand estimation learns only from past purchases and cumulative history as it would in practice.

In addition to developing personalization metrics, we also had to narrow our focus and select specific demand groups to analyze. We chose the seasonal home decor department because these products had historical price trajectories with steep clearance periods for excess inventory at the end of their selling seasons. We further restricted our consideration set to the top 500 SKUs by historical purchase frequencies over the two year period. We utilized association rule learning, which is a branch of machine learning stemming from collaborative filtering, to extract meaningful combinations of related products. Association rule learning is primarily used for finding groups of items that are frequently purchased together and constructing probabilistic implications, known as association rules, based on historical transactions. By leveraging well-known algorithms in this field such as Apriori from Agrawal et al. (1994), and FP-Growth from Verhein (2008), we constructed a set of approximately 25 demand groups, each of which was united by a common holiday or seasonal theme such as Valentine's Day, St. Patrick's Day, Halloween, Patriotic, or Autumnal. However, as demonstrated by Figures 5 and 6 in Appendix B.2, these demand groups were highly interconnected through historical purchases. Therefore, we used association rule learning again to assess the strength of the connections between products; the resulting outputs of this analysis ultimately led to five distinct demand groups, which we used for algorithm testing. Finally, having established consumer and product groups of interest, we proceed with the demand estimation, which follows the approach in Harsha and Subramanian (2016) and is detailed in Appendix B.2.

Simulation Design

For every demand group, we simulated personalized bundle offers to each historical consumer based on our estimated demand models and the true inventory levels and full product prices at the time of their arrival. When presented with this discounted offer the consumer chose to accept it, or to purchase any combination of available products in the demand group at full price, or to purchase nothing at all. Unlike the airline case study, we did not consider any limitations on bundle sizes but partitioned consumers arrival sequences based on demand groups (the largest of which was $|\hat{S}| = 8$ products). By aggregating the averages over the historical arrival sequences for 10,000

iterations within each demand group, which ranged from 2,800 to 13,000 consumers, we ultimately measured the percentage of expected Clairvoyant revenue achieved by each algorithm, and their respective conversion rates (percentage of offers that resulted in a purchase). Due to the inventory-constrained products, we solve *both* the upper-level full pricing problem *and* the lower-level dynamic bundle recommendation problem in all the simulations presented in Section 4.2.2. Note that the Clairvoyant formulation presented in (4) used inputs \bar{p}_i^t as provided externally by an oracle. In the implementation for this case study we consider a Clairvoyant in expectation, which periodically re-solves formulation (4) at the conclusion of each period t using the known arrival sequence (within each demand group) for the remainder of the selling horizon. By comparison, the multiplicative algorithm and the SIAA use \hat{p}_i^t from the upper-level formulation (8) that is solved periodically in expectation over future consumer arrivals. While both upper-level methods using a rolling approach in implementation, the full price trajectory inputs \bar{p}_i^t and \hat{p}_i^t vary, therefore resulting in inherent gaps between the performance of our algorithms and the full-knowledge strategy.

4.2.2. Results Our empirical results are divided into two discussions: (1) the effects of personalization and dynamic pricing, and, (2) the value of inventory balancing.

Value of Personalization and Dynamic Pricing & Business Insights

We initially considered the bundle recommendation problem in the unconstrained inventory setting in order to objectively measure and *emphasize the impact of personalization* and dynamic pricing on expected revenue, while implementing the upper-level problem framework for determining the full prices of all products in all periods. To develop these results we introduce three relevant benchmarks: (i) the “actual” pricing strategy that parallels the airline case and offers every product in the demand group at its historical full price at the time of each consumer’s arrival; (ii) a rolling LP method that periodically re-optimizes the full prices for all products based on segment-level expected future demand, then offers all products at these optimized fixed prices in a given period t to all arriving consumers; and, (iii) an un-personalized version of our myopic recommendation model that uses segment-level consumer features to make bundle offers to each arrival. The results are presented as the average expected percentage of Clairvoyant revenue attained by each pricing strategy and are summarized in Table 3.

Effect of Personalization on Empirical Performance		
Model	Personalization	Percent of Clairvoyant Revenue
Actual Historical Prices	X	88.2%
Rolling LP Model	X	95.5%
Segment-Level Dynamic Model	X	96.8%
Personalized Dynamic Model	✓	98.5%

Table 3 This table summarizes the empirical performance of all the benchmarks for the unconstrained setting in the retail case as a percentage of the expected Clairvoyant revenue, averaged across all demand and loyalty groups.

The output of personalized bundle offers was produced on average in approximately 3ms. Note that our model does not achieve 100% of the Clairvoyant profit due to discrepancies in the pricing of

the upper-level problem. We find that on average employing a dynamic pricing strategy over a static approach (rolling LP) improves expected revenue by 1.3%. Furthermore, leveraging *personalized models* of propensity-to-buy to make bundle offers *increases the expected revenue by an additional 1.7%* over a generic segment-level strategy. The overall improvement over the current pricing strategy is on average on the order of 10%, which is very substantial in such a thin-margin setting.

Our loyalty analysis shows that majority of online consumers in a product category are one-stop shoppers ($\approx 40\%$ within our selected demand groups). Therefore, the real focus of online e-tailers should be on *improving* conversions of higher frequency customers. The results from Table 3 present the objective benefit in expected revenue from personalized dynamic pricing strategies. We found that higher loyalty consumers provided the greatest expected lifts in revenue over the “actual” pricing benchmark within each demand group. Furthermore, as shown in Figures 3 and 4 in Appendix B.2, these consumers spend substantially more than other customers and have a larger source of historical data from which our model can develop more tailored bundle offers. Thus, we conclude a similar result to the airline case: on average, higher frequency consumers respond the most effectively to personalized discounted prices, and thus should be the primary target audience for recommendation systems aiming to raise expected revenues and conversions.

Value of Inventory

We now expand our results to the inventory-constrained setting inherent in the data set in order to analyze the practicality and performance of our approximation algorithms from Section 3 relative to benchmark methods and the Clairvoyant strategy. The “actual” and rolling LP benchmarks remain the same in this setting. We additionally introduce the myopic heuristic benchmark, which offers the personalized myopically profit-maximizing bundle to each consumer as in the unconstrained case. We consider two metrics of performance for each method: (i) the average expected percentage of Clairvoyant revenue achieved across all loyalty and demand groups, and, (ii) the average conversion rate. The resulting empirical performance ratios are summarized in Table 4 below. In this setting where we implemented our approximation algorithms that depended on inventory levels, the output of personalized bundle recommendations was produced on average in 15ms, which is still very efficient, as is necessary for implementation in an online environment. These results show that on average all of the methods perform relatively well: the ALA obtains 97% of the expected revenue achieved by the full-knowledge strategy, the SIAA reaches 93% and the multiplicative algorithm reaches 91%. The 4% performance gap between the additive methods is precisely the estimation difference accounted for by the additional bundle terms included in the ALA at discounted prices. We also observe an average expected gain of 5.5% in revenue over the myopic approach by accounting for inventory balancing and future demand in the SIAA. Furthermore, the overall average improvement in expected revenue from our best algorithm compared

to the “actual” historical pricing strategy from the data set is 14% across these demand groups. Note that while the best multiplicative algorithm (EMPA) is slightly outperformed by the additive methods in this highly inventory-constrained setting with steep markdown periods, it still performs within 9% of the full-knowledge strategy and within 6% of the best additive approach. Furthermore, the EMPA is significantly easier to implement and maintains a very close empirical performance relative to the ALA even in this challenging setting.

Table 4 This table summarizes the empirical performance in expected revenue of all the algorithms for the retail case study as a percentage of the full-information Clairvoyant revenue, averaged across all demand and loyalty groups.

Model	Percent of Clairvoyant Profit	Conversion Rate
Actual Historical Prices	83.2%	1.5%
Rolling LP Model	84.6%	3.1%
Myopic Heuristic	87.9%	4.2%
Exponential Multiplicative Algorithm	91.5%	6.0%
Separable-Item Algorithm (SIAA)	93.4%	6.6%
Lagrangian Algorithm (ALA)	97.5%	7.8%
Clairvoyant Model	100%	8.6%

Note that Table 4 illustrates the fact that the SIAA and ALA are empirically effective benchmarks that perform well relative to the full-knowledge strategy, but are much more reasonable for comparison to the EMPA because the Clairvoyant is not actually attainable in any practical setting. Furthermore, as shown by the range of inventory constrained cases above, these results demonstrate the significant benefit of inventory management through personalized recommendations by bundling items at a lesser discount ahead of the markdown period in order to preserve already narrow margins.

4.3. Comparisons

We conclude by presenting comparisons between the relative performances of our algorithms, as well as the empirical behavior of their analytical guarantees under various inventory settings.

Comparison of Approximation Algorithms

The first set of comparisons, summarized using airline case data in Table 5 below, show the expected percentage of Clairvoyant revenue achieved on average by each algorithm. As described in Section 4.1.2, we introduce inventory constraints in the airline data on ancillary products for which this is realistic: VIP lounge access, in-flight wi-fi, gourmet meals, excess checked baggage and seat upgrades. Furthermore, we implement the bi-level framework and also determine the full price trajectories of all ancillary products, which are initialized at the values in Table 8. Each column in Table 5 corresponds to the initial inventory level of the ancillary products as a function of the total number of consumer arrivals in the data set as described in Section 4.1; furthermore, we define the column “unlimited” as having a higher initial stock of each product than there are consumer arrivals, meaning that none of the products can ever be consumed entirely. The “actual” prices benchmark corresponds to the prior baseline that offers all the available inventory-constrained

ancillary products at their full prices. We include two additional sets of hybrid benchmarks based on (i) a threshold policy, and, (ii) an automated procedure. In approach (i) the hybrid algorithm makes all recommendations based on the EMPA until one of the products' inventories is depleted by 20%, after which all recommendations are made using the SIAA. In hybrid approach (ii), which we consider with two parameter settings, we employ a variant of the hybrid method in Golrezaei et al. (2014) where γ is a multiplicative weighting factor applied to the objective function of the offer chosen by the SIAA when compared to the offer selected by the EMPA. At greater values of γ , the SIAA recommendation is made more frequently. We discuss the implications of these hybrid results in more detail below in conjunction with Table 6. For robustness, we conducted this set of simulations on the retail data and observed symmetrical results.

Algorithm Comparison on Airline Data as Percentage of Clairvoyant Revenue

Algorithm	Initial Inventory Level					
	Unlimited	100%	90%	80%	75%	50%
"Actual" Prices	85.1%	82.6%	78.4%	76.4%	73.1%	60.7%
Re-Optimized Rolling LP	90.4%	88.4%	85.7%	83.8%	81.5%	65.2%
Linear Multiplicative Penalty	94.2%	92.3%	87.6%	84.1%	78.7%	63.8%
SIAA (Separable-Item Additive Algorithm)	96.9%	96.1%	95.3%	92.2%	88.6%	76.6%
Polynomial Multiplicative Penalty	95.1%	92.3%	89.3%	84.8%	79.9%	64.7%
Exponential Multiplicative Penalty (EMPA)	97.0%	94.8%	92.4%	88.2%	84.2%	69.2%
Rolling Exponential Multiplicative Penalty	97.4%	95.6%	94.2%	91.1%	87.3%	72.8%
ALA (Additive Lagrangian Algorithm)	97.9%	97.4%	96.1%	93.5%	89.8%	79.5%
Threshold Hybrid (SIAA after 20% depletion)	97.3%	95.7%	94.4%	91.3%	88.9%	74.1%
Automated Hybrid: $\gamma = 1.5$	97.3%	95.7%	94.5%	91.7%	88.1%	73.7%
Automated Hybrid: $\gamma = 2$	97.4%	95.8%	94.6%	91.9%	88.4%	75.3%

Table 5 This table summarizes the performance gaps of the proposed algorithms, as well as some hybrid algorithms, in the airline case study in percent of expected revenue attained relative to the full-knowledge Clairvoyant strategy.

From Table 5 we can conclude that the rolling implementation of the EMPA performs within 3-6% of the full knowledge strategy and within 1-2% of the ALA in less constrained inventory settings; furthermore, it outperforms the EMPA that uses initial inventory levels I_i^0 by 1-3% across all inventory cases. We observe that the SIAA begins to slightly outperform the rolling EMPA in increasingly more inventory-constrained settings, which is fairly intuitive: as it becomes more important to avoid inventory-related costs, the difference in approximation accuracy between the additive approaches and the multiplicative penalty becomes increasingly greater. However, the multiplicative method requires no re-optimization and is easy to implement compared to the ALA, while only under-performing by a margin of up to 5.6% in worst cases in while still achieving on average at least 87% of the expected Clairvoyant revenue in reasonable inventory settings, and at least 73% in the most constrained case. This marginal trade-off in empirical performance is largely offset by the practicality of the multiplicative approach, as well as the fact that in less constrained settings it performs within 1% of the best additive method. Finally, Table 5 also demonstrates the value of choosing the correct functional form of the inventory penalty function $\psi(\cdot)$ depending on

the data. For example in this case, the polynomial form of the multiplicative penalty under performs on average by approximately 3% compared to the multiplicative algorithm using the exponential penalty function across all inventory scenarios. Furthermore, these gaps in performance grow from 1.9% to up to 4.6% as the problem becomes more inventory constrained; thus, using an increasingly sophisticated form of $\psi(\cdot)$ results in a multiplicative algorithm with stronger performance.

The results in Table 5 provide an empirical foundation for understanding the performance differences between the algorithms relative to the Clairvoyant “optimal”, but it is also important to gain an insight into which settings each algorithm is best suited for. Therefore, we introduced the hybrid methods in these results, which alternate between recommendations from both the EMPA and the SIAA. It is clear from Table 5 that under mildly constrained inventory settings, all of the hybrids behave essentially as the EMPA. It is in the more constrained inventory cases that we begin to see a greater gap between the performance of the hybrid methods and the EMPA, due to the incorporation of demand forecasting captured in the additive approach. To understand the magnitude of this effect, we provide a second set of hybrid results in Table 6.

Percentage of Offers Driven by Additive Algorithm (SIAA)

Algorithm	Initial Inventory Level					
	50%	20%	15%	10%	5%	2%
Threshold Hybrid	16.8%	52.6%	75.7%	87.1%	95.4%	98.3%
Automated Hybrid: $\gamma = 1.5$	31.6%	63.7%	81.9%	90.2%	94.8%	97.5%
Automated Hybrid: $\gamma = 2$	38.5%	68.8%	84.2%	92.8%	96.4%	98.6%

Table 6 This table summarizes the percentage of personalized bundle offers made by each of the hybrid algorithms that are selected by the SIAA (as opposed to the EMPA) in the airline case study.

The focus of Table 6 is to observe the percentage of personalized bundle offers that are made by the SIAA (as opposed to the EMPA) in the hybrid algorithms under increasingly more constrained inventory settings. Notice that in the 50% that column Tables 5 and 6 share in common, the SIAA outperforms the EMPA by 4% and accounts for up to 40% of the recommendations in the hybrid method with parameter $\gamma = 2$. Furthermore, the results of Table 5 demonstrate that the EMPA becomes less accurate than the SIAA by a growing margin as the problem becomes increasingly more inventory constrained. Supplementing this effect with the computations from Table 6, we find that in the most tightly constrained scenarios, the SIAA recommendations represent 75-99% of the offers made in the hybrid approaches. The joint results of these two table indicates that the solution quality of the EMPA deteriorates at a greater rate than the offers selected by the SIAA, and that in the most tightly constrained inventory scenarios, all hybrid methods ultimately resolve to implementing the SIAA. We can conclude that this indicates that the SIAA provide significantly better quality results under highly constrained inventory settings, whereas the multiplicative methods are more efficient and virtually indistinguishable in performance relative to additive methods when initial inventory levels are high.

In addition to algorithm performance and suitability, we also assess the difference between the actual composition and pricing of bundle offers made by the ALA and EMPA. We found that the

multiplicative approach typically recommends more expensive products at a 1-2% steeper discount than the ALA. Consider the following case of a business traveler flying in economy: the ALA recommends seat upgrades (\$50) and VIP lounge access (\$50) at an average of a 2.4% discount. By contrast, the EMPA recommends VIP lounge access (\$50) and 2,000 bonus miles (\$100) at an average discount of 4.8%. This indicates that rougher inventory estimates in EMPA generate bundles with a greater emphasis on myopic profit maximization that offset more expensive product offers with higher discounts to increase consumer propensity-to-buy. However in expectation, these two methods generate similar expected revenues that are on average within 10-13% of the full-knowledge approach in less inventory constrained cases. We also conduct a series of computational studies of the effects of marginal costs and competition on the algorithms relative to both one another as well as the Clairvoyant model, all of which are detailed in Appendix B.3.

Empirical Comparison of Analytical Guarantees

Finally, we discuss the difference between the empirical analytical guarantees and the practical algorithm performance in the multiplicative approach. Consider the following airline case results for the EMPA in Table 7 in which we now also implement the upper-level full pricing problem (similar tables for polynomial and linear multiplicative penalties echo the trends observed here).

Empirical Performance of the EMPA

Exp. Penalty Function	Initial Inventory Level				
	Unlimited	100%	90%	80%	75%
Performance Ratio from Data	91.6%	88.0%	84.3%	82.0%	76.6%
Posterior Bound on Ratio	84.8%	79.4%	75.8%	71.3%	65.6%
Prior Bound on Ratio	80.4%	74.7%	70.9%	65.6%	58.7%

Table 7 This table presents the average expected percentage of Clairvoyant revenue attained by the EMPA and the empirical values of its analytical guarantees that are path-dependent (posterior bound) and path-independent (prior bound).

The first row shows the average expected percentage of Clairvoyant revenue achieved by the EMPA from empirical data simulations. Note that these ratios slightly decrease from the results in Table 5 due to discrepancies between the upper-level solutions of formulation (8) and the Clairvoyant. The second row represents the corresponding data-driven values of a path-dependent lower bound on the performance ratio between $\{\text{MultAlg}\}_{v,k,t}$ and $\{\text{Clairvoyant}\}$, which is derived in the proof of Theorem 1 in Appendix A and depends on the trajectory of inventory levels I_i^t for a fixed consumer arrival sequence $\{k, t\}_{k=1, \dots, K}^{t=1, \dots, T}$. This intermediate bound is less conservative than the worst-case result of Theorem 1, which we refer to as the prior bound in the third row. The results in Table 7 demonstrate that the relative gaps between the algorithm's empirical performance ratios (first row) and the posterior bound (second row) grow increasingly as the problem becomes more inventory-constrained; we observe the same effect between the posterior bounds (second row) and the prior bounds (third row). Note that the empirical performance ratios relative to the Clairvoyant also decrease as a function of initial inventory levels, which confirms the insight that more limited inventory results in greater performance lag across all of the online algorithms relative to

the full-knowledge offline strategy. We found that this result was consistent across all scenarios in both case studies and all forms of penalty functions. However, the empirical performance of the algorithm on the actual data is significantly better than the worst-case analytical guarantees provided by the prior bound from Theorem 1. In the most inventory-constrained scenarios, the gap between the empirical ratio and the prior bound reaches up to 18%, while the actual performance of the algorithm is within at least 14% of the expected Clairvoyant revenue in reasonably inventory-constrained scenarios, improving to within 9% on average in the least constrained cases.

References

- Agrawal R, Srikant R, et al. (1994) Fast algorithms for mining association rules. *Proc. 20th int. conf. very large data bases, VLDB*, volume 1215, 487–499.
- Agrawal S, Wang Z, Ye Y (2014) A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62(4):876–890.
- Aouad A, Levi R, Segev D (2015) Approximation algorithms for dynamic assortment optimization models .
- Bernstein F, Kök AG, Xie L (2015) Dynamic assortment customization with limited inventories. *Manufacturing & Service Operations Management* 17(4):538–553.
- Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* 5(3):203–229.
- Cavusgil ST, Zou S (1994) Marketing strategy-performance relationship: an investigation of the empirical link in export market ventures. *The Journal of Marketing* 1–21.
- Chen R, Jiang H (2017) Capacitated assortment and price optimization for customers with disjoint consideration sets. *Operations Research Letters* 45(2):170–174.
- Ferreira KJ, Simchi-Levi D, Wang H (2016) Online network revenue management using thompson sampling .
- Gallego G, Li A, Truong VA, Wang X (2015) Online resource allocation with customer choice. *arXiv preprint arXiv:1511.01837* .
- Gallego G, Li A, Truong VA, Wang X (2016) Approximation algorithms for product framing and pricing. Technical report, Working paper, available online as SSRN report.
- Gallego G, Topaloglu H (2014) Constrained assortment optimization for the nested logit model. *Management Science* 60(10):2583–2601.
- Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-time optimization of personalized assortments. *Management Science* 60(6):1532–1551.
- Goyal V, Levi R, Segev D (2016) Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand. *Operations Research* 64(1):219–235.

- Harsha P, Subramanian S (2016) Hierarchical demand modeling in the presence of partially or completely unobserved lost sales. *INFORMS Revenue Management and Pricing Conference (New York, NY, United States, June 2016)*.
- Hawkins JT (2003) *A Lagrangian decomposition approach to weakly coupled dynamic optimization problems and its applications*. Ph.D. thesis, Massachusetts Institute of Technology.
- Jaillet P, Lu X (2012) Near-optimal online algorithms for dynamic resource allocation problems. *arXiv preprint arXiv:1208.2596* .
- Kesselheim T, Tönnis A, Radke K, Vöcking B (2014) Primal beats dual on online packing lps in the random-order model. *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, 303–312 (ACM).
- Kök AG, Fisher ML, Vaidyanathan R (2009) Assortment planning: Review of literature and industry practice. *Retail supply chain management*, 99–153 (Springer).
- Li G, Rusmevichientong P, Topaloglu H (2015) The d-level nested logit model: Assortment and price optimization problems. *Operations Research* 63(2):325–342.
- Mahajan S, van Ryzin G (1999) On the relationship between inventory costs and variety benefits in retail assortments. *Management Science* 45(11):1496–1509.
- Mahajan S, van Ryzin G (2001) Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49(3):334–351.
- Netessine S, Savin S, Xiao W (2006) Revenue management through dynamic cross selling in e-commerce retailing. *Operations Research* 54(5):893–913.
- Rapti E, Karageorgos A, Ntalos G (2014) Adaptive constraint and rule-based product bundling in enterprise networks. *WETICE Conference (WETICE), 2014 IEEE 23rd International*, 15–20 (IEEE).
- Rusmevichientong P, Shmoys D, Tong C, Topaloglu H (2014) Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* 23(11):2023–2039.
- Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research* 60(4):865–882.
- Shamsi D, Holtan M, Luenberger R, Ye Y (2014) Online allocation rules in display advertising. *arXiv preprint arXiv:1407.5710* .
- Talluri KT, van Ryzin GJ (2006) *The theory and practice of revenue management*, volume 68 (Springer Science & Business Media).
- Topaloglu H (2009) Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. *Operations Research* 57(3):637–649.
- Verhein F (2008) Frequent pattern growth (fp-growth) algorithm. *The University of Sydne. p1-10* .
- Xue Z, Wang Z, Ettl M (2015) Pricing personalized bundles: A new approach and an empirical study. *Manufacturing & Service Operations Management* 18(1):51–68.

Appendix A: Proof of Analytical Result

Proof of Theorem 1 We are interested in providing an analytical guarantee on the competitive ratio between the multiplicative algorithm and the optimal clairvoyant strategy. Specifically, we want to attain a lower bound on the performance of the following model, whose objective we will now refer to as $\{\text{MultAlg}\}_{\forall(k,t)}$:

$$\begin{aligned} \{\text{MultAlg}\}_{\forall(k,t)} = & \text{maximize}_{S_{k,t} \subset \hat{S}, p_{S_{k,t}}} \left[\sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] + \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (p_{S_{k,t}} - \bar{p}_{S_{k,t}}) \cdot \min_{i \in S_{k,t}} \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \\ \text{subject to} & \quad (1 - \epsilon) \bar{p}_{S_{k,t}} \leq p_{S_{k,t}} \leq \bar{p}_{S_{k,t}} \quad \forall(k,t), S_{k,t} \subset \hat{S} \end{aligned}$$

For any given sequence of customers $\{k,t\}_{\forall(k,t),t=1}^T$, we have the following primal $\{\text{Clairvoyant}\}$ problem that has full knowledge of all arrival types in advance, as presented in Section 2:

$$\begin{aligned} \text{maximize}_{y_{S_{k,t}}^{k,t}} & \sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{S_{k,t} \subset \hat{S}} \left(\left[\sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \right] + \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (p_{S_{k,t}} - \bar{p}_{S_{k,t}}) \right) \cdot y_{S_{k,t}}^{k,t} \\ \text{subject to} & \sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{S_{k,t} \subset \hat{S}} \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot y_{S_{k,t}}^{k,t} \right] \leq I_i^0 \quad \forall i \\ & \sum_{S_{k,t} \subset \hat{S}} y_{S_{k,t}}^{k,t} = 1 \quad \forall(k,t) \\ & y_{S_{k,t}}^{k,t} \geq 0 \quad \forall(k,t), S_{k,t} \subset \hat{S} \end{aligned} \quad (12)$$

By weak duality we aim to find the following lower bound on the competitive ratio between our algorithm and the clairvoyant primal problem:

$$\frac{\{\text{MultAlg}\}_{\forall(k,t)}}{\{\text{Clairvoyant}\}} \geq \frac{\{\text{MultAlg}\}_{\forall(k,t)}}{\{\text{Dual}\}}$$

We let the price of the bundle $S_{k,t}$ offered to consumer (k,t) be $p_{S_{k,t}}$, defined explicitly by the bundle discount price ratio $d_{S_{k,t}}$ as follows,

$$d_{S_{k,t}} = \frac{p_{S_{k,t}}}{\bar{p}_{S_{k,t}}}.$$

Thus, in order to derive the desired bound on the ratio of the primal problem using weak duality, we consider its dual given by $\{\text{Dual}\}_{\forall(k,t)}$ below,

$$\begin{aligned} \min_{\theta_i^t, \lambda^{k,t}} & \sum_{i=1}^n I_i^0 \cdot \theta_i + \sum_{t=1}^T \sum_{k \in K^t} \lambda^{k,t} \\ \text{subject to} & \lambda^{k,t} \geq \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) (\bar{p}_i^t - \theta_i) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \quad \forall(k,t), S_{k,t} \subset \hat{S} \\ & \theta_i \geq 0 \quad \forall i \end{aligned} \quad (13)$$

For the dual problem in Eq. (13), based on the choices of consumers in the sequence $\{k,t\}_{\forall(k,t),t=1}^T$, we utilize the result from Proposition 1 to consider the following dual feasible solution, where I_i^0 is the initial inventory of product i and \bar{p}_i^0 is the corresponding initial nominal price setting:

$$\begin{aligned} \hat{\theta}_i &= \bar{p}_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right), \quad \forall i \\ \hat{\lambda}^{k,t} &= \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right], \quad \forall(k,t) \end{aligned}$$

We now want to find the expected value of this dual feasible solution as it will give us an upper bound on the expected objective of the primal problem by weak duality. Since we have a fixed sequence $\{k,t\}_{\forall(k,t),t=1}^T$ the expectation is taken relative to each consumer's purchase decision, given the current state of inventory $\{I_1^{k,t}, I_2^{k,t}, \dots, I_n^{k,t}\}$. By Lemma 1, we obtain the following expression for the expectation over the dual feasible variables $\hat{\lambda}^{k,t}$:

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \hat{\lambda}^{k,t} \right] = \sum_{i=1}^n \left(\sum_{t=1}^T \sum_{l=I_i^t+1}^{I_i^{t-1}} \bar{p}_i^t \cdot \psi \left(\frac{l}{I_i^0} \right) \right) - \sum_{t=1}^T M^t,$$

We define the time-dependent constant M^t with the following expression:

$$M^t = \max_{S_{k,t} \subset \hat{S}, d_{S_{k,t}}} \sum_{k=1}^{K^t} \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (1 - d_{S_{k,t}})$$

We thus get the following form for our expected dual objective denoted {Dual}:

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \hat{\lambda}^{k,t} + \sum_{i=1}^n I_i^0 \cdot \hat{\theta}_i \right] = \sum_{i=1}^n \left[\sum_{t=1}^T \bar{p}_i^t \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \cdot \bar{p}_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right) \right] - \sum_{t=1}^T M^t$$

We want to now compare the expected objective values of the dual problem calculated above to the expected value of the proposed heuristic approach, which we defined as {MultAlg}_{V(k,t)}. The expected revenue can be written as follows from Proposition 2, denoted {MultAlg}_{V(k,t)}:

$$\sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t - \sum_{t=1}^T M^t$$

We can now revisit the original goal to use weak duality and finally derive the following desired ratio:

$$\frac{\{\text{MultAlg}\}_{V(k,t)}}{\{\text{Clairvoyant}\}} \geq \frac{\sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t - \sum_{t=1}^T M^t}{\sum_{i=1}^n \left[\sum_{t=1}^T \bar{p}_i^t \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \cdot \bar{p}_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right) \right] - \sum_{t=1}^T M^t}$$

However, this bound is path-dependent and relies on knowledge of the final inventory levels in order to calculate a value. We want to now develop a bound that depends solely on the initial conditions to compare our algorithm to the clairvoyant approach. We therefore work to bound it further to develop a worst-case analytical guarantee that is dependent only on initial inventory levels and expected demand (by using arrival rate estimates for consumer types to calculate $\sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$). We recall the time-dependent price trajectory definitions:

$$\begin{aligned} \alpha_i^t &= \frac{\bar{p}_i^t}{\bar{p}_i^0} \quad \forall i, t, \text{ as determined by formulation (8) of the upper-level problem in Section 3.1,} \\ \beta_i^t &= \frac{\bar{p}_i^t}{\bar{p}_i^0} \quad \forall i, t, \text{ as determined by formulation (12) of the Clairvoyant problem in Appendix A.} \end{aligned} \quad (14)$$

Based on the above expression, α_i^t and β_i^t are the extent of the discount on the full price of item i in period t from its initial setting at \bar{p}_i^0 , which is common to both the Clairvoyant and our upper-level method from formulation (8). Note that both algorithms are provided with these nominal price discounts in advance. Thus, we get the result below:

$$\begin{aligned} \frac{\{\text{MultAlg}\}_{V(k,t)}}{\{\text{Clairvoyant}\}} &\geq \frac{\sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t - \sum_{t=1}^T M^t}{\sum_{i=1}^n \left[\sum_{t=1}^T \bar{p}_i^t \cdot \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \cdot \bar{p}_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right) \right] - \sum_{t=1}^T M^t} \\ &= \frac{\sum_{i=1}^n \bar{p}_i^0 \cdot \sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\sum_{i=1}^n \bar{p}_i^0 \cdot \left[\sum_{t=1}^T \beta_i^t \cdot \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right) \right] - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}, \text{ by Eq. (14),} \\ &\stackrel{\min_{(I_i^0, I_i^T): I_i^T \leq I_i^0 - 1}}{\geq} \frac{\sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\sum_{t=1}^T \beta_i^t \cdot \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \left(1 - \psi \left(\frac{I_i^T}{I_i^0} \right) \right) - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}, \text{ Lem. 2.} \end{aligned}$$

By considering the discount factors α_i^t and β_i^t for each product i in period t , we are able to isolate the constant \bar{p}_i^0 and reduce the outer summation using a minimization, as a result of Lemma 2. We now introduce a change of variable by considering $x = \frac{l}{I_i^0}$ and get the following equivalent expressions.

$$\begin{aligned} &\min_{(I_i^0, x): x \leq 1 - \frac{1}{I_i^0}} \frac{\sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\sum_{t=1}^T \beta_i^t \cdot \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + I_i^0 \left(1 - \psi(x) \right) - \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t} \\ &= \min_{(I_i^0, x): x \leq 1 - \frac{1}{I_i^0}} \frac{\frac{l}{I_i^0} \sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t - \frac{l}{I_i^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\frac{l}{I_i^0} \sum_{t=1}^T \beta_i^t \cdot \sum_{l=I_i^t+1}^{I_i^{t-1}} \psi \left(\frac{l}{I_i^0} \right) + 1 - \psi(x) - \frac{l}{I_i^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t} \end{aligned}$$

In the second expression we scale all of the terms by $\frac{1}{I_i^0}$, so that we can apply the property below, which is the result of Lemma 3.

$$\frac{1}{I_i^0} \sum_{l=I_i^T+1}^{I_i^0} \psi\left(\frac{l}{I_i^0}\right) \leq \frac{1}{I_i^0} + \int_{\frac{I_i^T+1}{I_i^0}}^1 \psi(y) dy$$

By applying this to the previous expression we get the following result,

$$\min_{(I_i^0, x): x \leq 1 - \frac{1}{I_i^0}} \frac{\frac{l}{I_i^0} \sum_{t=1}^T \sum_{k=1}^{K^t} \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t - \frac{l}{I_i^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\frac{l}{I_i^0} + \sum_{t=2}^T \beta_i^t \cdot \int_{I_i^{t-1}}^{I_i^t} \psi(y) dy + 1 - \psi(x) - \frac{l}{I_i^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}$$

Finally we introduce $I_{\min}^0 = \min_i I_i^0$ (and symmetrically also I_{\max}^0 , and β_{\max}). We also define $R_{\min}^{k,t} = \min_i \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \alpha_i^t$ and by definition conclude the following result,

$$\min_{(I_i^0, x): x \leq 1 - \frac{1}{I_{\min}^0}} \frac{\frac{l}{I_{\max}^0} \sum_{t=1}^T \sum_{k=1}^{K^t} R_{\min}^{k,t} - \frac{l}{I_{\min}^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}{\beta_{\max} \cdot \frac{l}{I_{\min}^0} \cdot \int_{x=1+\frac{1}{I_{\min}^0}}^1 \psi(y) dy + 1 - \psi(x) - \frac{l}{I_{\max}^0} \cdot \frac{1}{\sum_{i=1}^n \bar{p}_i^0} \sum_{t=1}^T M^t}$$

This completes the proof of Theorem 1.

Proposition 1 For the dual problem presented in formulation (13), the following is a dual feasible solution, where I_i^0 is the initial inventory of product i and \bar{p}_i^0 is the initial nominal price setting for product i :

$$\begin{aligned} \hat{\theta}_i &= \bar{p}_i^0 \left(1 - \psi\left(\frac{I_i^T}{I_i^0}\right) \right) \\ \hat{\lambda}^{k,t} &= \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi\left(\frac{I_i^{k,t}}{I_i^0}\right) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \forall (k,t) \end{aligned}$$

Proof of Proposition 1 Given the formulation of $\{\text{Dual}\}_{\forall(k,t)}$ presented in Eq. (13), we want to show the following two conditions:

$$\begin{aligned} (1) \quad \lambda^{k,t} &\geq \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) (\bar{p}_i^t - \theta_i) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \forall k, t, S_{k,t} \subset \hat{S} \\ (2) \quad \theta_i &\geq 0 \quad \forall i \end{aligned}$$

Let us first focus on the more challenging condition (1). We define a new term as follows:

$$\theta_i^t = \bar{p}_i^t \left(1 - \psi\left(\frac{I_i^T}{I_i^0}\right) \right) = \bar{p}_i^t - \bar{p}_i^t \cdot \psi\left(\frac{I_i^T}{I_i^0}\right) \quad \forall i, t \quad (15)$$

Note that this new term θ_i^t is based on the nominal price \bar{p}_i^t for product i in period t . Thus, $\bar{p}_i^t \leq \bar{p}_i^0$, because all nominal prices follow a markdown trajectory over time. Therefore, $\theta_i^t \leq \theta_i^0 = \hat{\theta}_i$.

We can now show feasibility using this new terminology as follows:

$$\begin{aligned} \hat{\lambda}^{k,t} &= \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi\left(\frac{I_i^{k,t}}{I_i^0}\right) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &\geq \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi\left(\frac{I_i^T}{I_i^0}\right) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &= \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (\bar{p}_i^t - \theta_i^t) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &\geq \sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot (\bar{p}_i^t - \hat{\theta}_i) \right] + \sum_{i \in S_{k,t}} \left[\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \end{aligned}$$

We get the first inequality from the fact that $\psi(\cdot)$ is concave and increasing and $I_i^T \leq I_i^t \quad \forall t = 1, \dots, T$. The second equality comes directly from the definition of θ_i^t in Eq. (15); finally this leads to the last inequality by applying $\theta_i^t \leq \hat{\theta}_i$. For condition (2) concerning $\hat{\theta}_i$, showing feasibility is trivial. As stated, $\psi(\cdot)$ is a concave monotone increasing function defined on $[0, 1]$, so $\bar{p}_i^0 \cdot \psi(\cdot) \leq \bar{p}_i^0$. Thus, $\hat{\theta}_i = \theta_i^0 \geq 0 \quad \forall i$ by definition.

LEMMA 1. For a fixed arrival sequence $\{k, t\}_{\forall(k,t), t=1}^T$, the expected value of the expectation of the duals variables $\hat{\lambda}^{k,t}$ is defined by the expression:

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \hat{\lambda}^{k,t} \right] = \sum_{i=1}^n \left(\sum_{t=1}^T \sum_{l=I_i^t+1}^{I_i^{t-1}} \bar{p}_i^t \cdot \psi \left(\frac{l}{I_i^0} \right) \right) - \sum_{t=1}^T M^t, \quad (16)$$

where $M^t = \max_{S_{k,t} \subset \bar{S}, d_{S_{k,t}}} \sum_{k=1}^{K^t} \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (1 - d_{S_{k,t}})$.

Proof of Lemma 1 For a fixed arrival sequence $\{k, t\}_{\forall(k,t), t=1}^T$, we want to find the expected value of the objective function of the dual of the Clairvoyant problem. We define a binary variable $Q_i^{k,t} = 1$ if item i is purchased at time t , and is 0 otherwise. We first use this to consider the expectation $\mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \hat{\lambda}^{k,t} \right]$ over consumer choices below:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \hat{\lambda}^{k,t} \right] &= \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \left(\sum_{i=1}^n \left[\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] \right) + \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \left(\sum_{i=1}^n \left[Q_i^{k,t} \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] \right) + \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \left(\sum_{i=1}^n \left[(I_i^{k,t} - I_i^{k+1,t}) \cdot \bar{p}_i^t \cdot \psi \left(\frac{I_i^{k,t}}{I_i^0} \right) \right] \right) + \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (d_{S_{k,t}} - 1) \right] \\ &= \sum_{i=1}^n \left(\sum_{t=1}^T \sum_{l=I_i^t+1}^{I_i^{t-1}} \bar{p}_i^t \cdot \psi \left(\frac{l}{I_i^0} \right) \right) - \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (1 - d_{S_{k,t}}) \right] \\ &\leq \sum_{i=1}^n \left(\sum_{t=1}^T \sum_{l=I_i^t+1}^{I_i^{t-1}} \bar{p}_i^t \cdot \psi \left(\frac{l}{I_i^0} \right) \right) - \sum_{t=1}^T M^t, \text{ where } M^t = \max_{S_{k,t} \subset \bar{S}, d_{S_{k,t}}} \sum_{k=1}^{K^t} \sum_{i \in S_{k,t}} \phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \cdot (1 - d_{S_{k,t}}) \end{aligned} \quad (17)$$

Note that the first equality comes from the definition of Q_i^t as a Bernoulli variable that takes the value 1 if item i is purchased at time t by consumer k , and 0 otherwise. Thus, the conditional expectation on the binary variable $Q_i^{k,t}$ given the current inventory state \mathbf{I}^{t-1} is precisely $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$,

$$\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) = \mathbb{E} \left[Q_i^{k,t} \mid I_1^{t-1}, I_2^{t-1}, \dots, I_n^{t-1} \right],$$

because conditional on the current state of inventory \mathbf{I}^{t-1} (which drives the selection of $S_{k,t}$), the expected value of $Q_i^{k,t}$ is the probability that consumer k, t will purchase product i when offered $S_{k,t}$ at price $\mathbf{p}_{S_{k,t}}$, defined as: $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$. The second equality comes from the fact that $Q_i^{k,t}$ is a binary variable exactly defined by $I_i^t - I_i^{t-1}$.

Proposition 2 Given a fixed consumer arrival sequence $\{k, t\}_{\forall(k,t), k=1, \dots, K^t}^{t=1, \dots, T}$, the expected value of the objective function of $\{\text{MultAlg}\}_{\forall(k,t)}$ is given by the expression:

$$\{\text{MultAlg}\}_{\forall(k,t)} = \mathbb{E} \left[\sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \right] - \sum_{t=1}^T M^t$$

Proof of Proposition 2 Given a fixed consumer arrival sequence $\{k, t\}_{\forall(k,t), t=1}^T$, we can derive the objective value of the multiplicative approximation algorithm as follows:

$$\begin{aligned} \{\text{MultAlg}\}_{\forall(k,t)} &= \sum_{t=1}^T \sum_{k=1}^{K^t} \left(\left(\sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t \right) + \left(\phi_{S_{k,t}}^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_{S_{k,t}} (1 - d_{S_{k,t}}) \right) \right) \\ &= \sum_{t=1}^T \sum_{k=1}^{K^t} \sum_{i=1}^n \phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) \cdot \bar{p}_i^t - \sum_{t=1}^T M^t \end{aligned}$$

The second equality comes directly from the definition of M^t from Lemma 1.

LEMMA 2. Given a fixed set of constants $a_i \forall i = 1, \dots, n$ and corresponding variables x_i and y_i , the following property holds:

$$\frac{\sum_{i=1}^n a_i \cdot x_i}{\sum_{i=1}^n a_i \cdot y_i} \geq \min_i \frac{x_i}{y_i}$$

Proof of Lemma 2 We want to show that we can lower bound the ratio of two sums with the same weights a_i and different variable values x_i and y_i using a minimum over the ratios of all the variable pairs x_i, y_i . Let us first define the following term,

$$\hat{\alpha} = \min_i \frac{x_i}{y_i}$$

By the definition of α we know that $x_i \geq \hat{\alpha} \cdot y_i \forall i$. Therefore, we get the following result as desired,

$$\frac{\sum_{i=1}^n a_i \cdot x_i}{\sum_{i=1}^n a_i \cdot y_i} \geq \frac{\sum_{i=1}^n a_i \cdot (\hat{\alpha} \cdot y_i)}{\sum_{i=1}^n a_i \cdot y_i} = \hat{\alpha} = \min_i \frac{x_i}{y_i}.$$

LEMMA 3. Given a monotone increasing function $\psi(\cdot)$ and an increasing set of constants $x = x_0, \dots, x_N$, the following condition holds,

$$\sum_{x=x_0}^{x_N-1} \psi(x) \leq \int_{x=x_0}^{x_N} \psi(y)dy$$

Proof of Lemma 3 By definition of the values of x , we know that $x_0 \leq x_1 \leq \dots \leq x_N$. Since $\psi(\cdot)$ is a monotone increasing function we have that $\psi(x) \geq \psi(x_i) \forall x \in [x_i, x_i + 1]$. If we integrate this expression over $[x_i, x_i + 1]$ for a fixed value of x_i we get,

$$\int_{x_i}^{x_i+1} \psi(x) dx = \psi(x_i) \leq \int_{x_i}^{x_i+1} \psi(x) dx$$

Rewriting the left hand expression through a summation we precisely get the desired result,

$$\sum_{x=x_0}^{x_N-1} \psi(x) \leq \int_{x=x_0}^{x_N} \psi(y)dy.$$

Appendix B: Supplemental Figures and Tables from Case Studies

B.1. Airline Case Study

Prices of Ancillary Services & Distribution of Consumers Across Tier Levels

These are the ancillary products offered to consumers “un-bundled” from the ticket itself, currently offered at their full prices from Table 8. Figure 9 shows the distribution of consumers across the 7 cluster types.

Ancillary Good	Full Price
In-Flight Wi-Fi Access	\$25
Premium Entertainment	\$20
Priority Security	\$20
Priority Boarding	\$10
Priority Baggage Handling	\$25
Comfort Seating Upgrade	\$50
Excess Checked Baggage	\$45
VIP Lounge Access	\$50
In-Flight Gourmet Meals	\$15
2,000 Bonus Miles	\$100
4,000 Bonus Miles	\$200

Table 8 This table summarizes the full prices at which each of the ancillary products are typically offered.

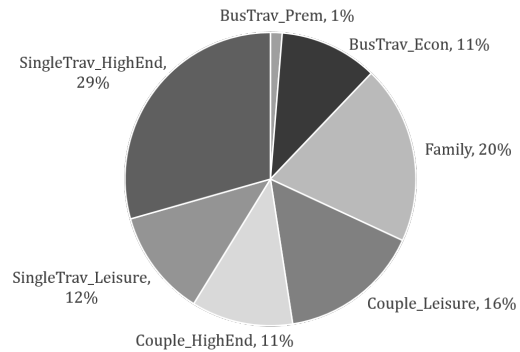


Table 9 This chart shows the distribution of consumers in the airline data set across various persona profile clusters.

Demand Estimation

We treated the products in the airline data set as independent because they did not have a direct effect on one another’s prices. Therefore, we fit a MNL logit single choice model for each (consumer type, product) pair. We considered 7 consumer types as explained on page 18: business travelers in premium class, business travelers in economy class, family travelers, last minute groups, couples traveling in economy, high end couples traveling in premium, single

leisure travelers, and single high end travelers. The ancillary products of interest were: wi-fi access, premium on-board entertainment, priority security, priority boarding, priority baggage handling, seat upgrades, checked/excess baggage, VIP lounge access, gourmet on-board meals, 2,000 bonus miles and 4,000 bonus miles (their nominal prices are displayed on page 41 of the paper). For a given pair, for example (family traveler, priority boarding), we would fit a MNL logit model to obtain the probability with which a traveler with feature vector \mathbf{x} of the type family would purchase priority boarding (product i), given by the equation:

$$\phi_i^{k,t}(\mathbf{p}_{S_{k,t}}) = \frac{e^{\beta_i^i \mathbf{x}}}{1 + e^{\beta_i^i \mathbf{x}}}, \quad \text{where,} \quad (18)$$

$$\beta_i^i \mathbf{x} = \beta_0^i + \beta_1^i \cdot \text{price}_i^t + \beta_2^i \cdot \text{miles}_{k,t} + \beta_3^i \cdot \text{tier}_{k,t} + \beta_4^i \cdot \text{mton}_{k,t} + \beta_5^i \cdot \text{neco}_{k,t} + \beta_6^i \cdot \text{npre}_{k,t} + \beta_7^i \cdot \text{ttod}_{k,t} \\ + \beta_8^i \cdot \text{fUSD}_{k,t} + \beta_9^i \cdot \text{dwek}_{k,t} + \beta_{10}^i \cdot \text{tjin}_{k,t} + \beta_{11}^i \cdot \text{mred}_{k,t} + \beta_{12}^i \cdot \text{npsg}_{k,t} + \beta_{13}^i \cdot \text{ctim}_{k,t}$$

Note that across all products i we consider the same structure of feature vector \mathbf{x} for each consumer $\{k,t\}$ and since $\phi(\cdot)$ is fit uniquely to each (consumer type, product) pair, we are able to generate a personalized estimate of willingness to buy $\phi_i^{k,t}(\mathbf{p}_{S_{k,t}})$ for each consumer $\{k,t\}$ considering each product i . The independent variables in Eq. (18) used to estimate the coefficients are defined as follows: miles = miles balance, tier = tier level, mton = miles to next tier level, neco = number of previous flights in economy in past 2 years, npre = number of flights in premium in past 2 years, ttod = time to departure (in weeks), fUSD = fare in USD, dwek = day of week of departure (categorical), tjin = time since joining mileage program (in weeks), mred = miles redeemed for rewards in past two years, npsg = number of passengers in booking, and ctim = connection time for itinerary booked. In this case due to independence, bundle buy probabilities were constructed by multiplying the above willingness-to-buy functions for the products in a particular bundle under consideration.

B.2. Retail Case Study

Cumulative Expenditure of Medium and High Frequency Loyalty Groups over Selling Horizon

Figures 3 and 4 show the time-series behavior of the middle and high frequency consumer loyalty groups with respect to their cumulative expenditure over the two year selling horizon. The y-axis is the average cumulative expenditure of each loyalty group depending on the week, given by the x-axis.

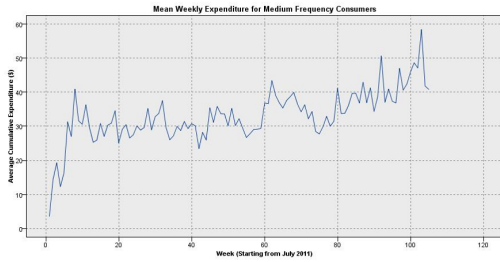


Figure 3 This plot shows the mean expenditure of medium loyalty repeat customers.

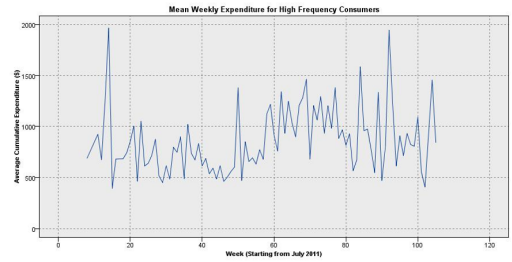


Figure 4 This plot shows the mean expenditure of high loyalty repeat customers.

Interdependence of Seasonal Home Decor Demand Groups

A detailed analysis of the demand groups across the top 500 SKUs in the seasonal home decor department demonstrated that many products were interconnected through historical purchases, as shown in Figures 5 and 6 below. The thickness of the lines between pairs of SKUs indicate how often these products were historically purchased together.

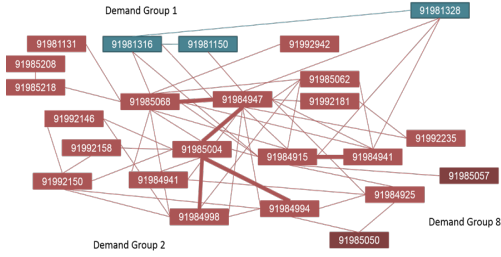


Figure 5 This is a visualization of the joint historical transactions between multiple demand groups.



Figure 6 This visualization focuses on the most frequent interactions between products in demand group 1.

Demand Estimation

The longer selling horizon in the retail case required consideration of time-dependent information not present in the airline data. We used the demand estimation approach from Harsha and Subramanian (2016) in order to capture correlation among products, lost sales and variability in market size and share. For each (demand group, loyalty group) pair, we aggregated the transaction data to the weekly level and considered as input mean values for net charged amount, sold quantities and various seasonality, holiday and promotional period flags in order to also estimate market size and share in addition to willingness-to-buy. We then fit a multi-choice MNL model for all customers in a specific loyalty group, when considering bundles S in a demand group \hat{S} , given by:

$$\phi_S^t(\mathbf{p}_{S,t}) = \frac{e^{\beta_t^S \cdot \mathbf{x}}}{\sum_{S \in \hat{S}} e^{\beta_t^S \cdot \mathbf{x}}}, \quad \text{where,} \tag{19}$$

$$\beta_t^S \cdot \mathbf{x} = \beta_0^S + \beta_1^S \cdot \text{price}_t^S + \beta_2^S \cdot \text{CumVis}_t + \beta_3^S \cdot \text{CumExp}_t + \beta_4^S \cdot \mathbb{1}_{\text{promo}_t} + \beta_5^S \cdot \mathbb{1}_{\text{clear}_t} + \beta_6^S \cdot \mathbb{1}_{\text{hldy}_t}$$

Note that the feature vector \mathbf{x} that we use is the same structure for all consumers and loyalty groups, but its magnitude and values change. Specifically, the willingness to pay function for consumer $\{k, t\}$ is $\phi_S^{k,t}(\mathbf{p}_{S,k,t})$, which is Eq. (19) populated with the specific characteristics of that consumer; for example, the cumulative number of visits and expenditure of consumer $\{k, t\}$ are used in conjunction with the estimated coefficients β_2^S and β_3^S . Furthermore, each bundle S is defined uniquely based on its composition (which is reflected in its price and various associated flags) and therefore has a different interaction with a given consumer $\{k, t\}$ than other bundles in \hat{S} . The variables corresponding to the estimated coefficients above are defined as follows: CumVis = number of cumulative consumer visits at time t , CumExp = amount of cumulative consumer expenditure at time t , promo = promotion period flag, clear = clearance period flag, and hldy = holiday flag. As each bundle is uniquely defined through its composition, bundle size is not explicitly a covariate. Our demand modeling approach accounts for lost sales as well as market size estimation within the demand fitting process. We do not assume a fixed market size and instead capture its variability at any given time period t through the following two expressions in Eq. (20):

$$\text{Weekly Sales}_{S,t} \approx \text{Market Size}_t \cdot (\phi_S^t(\mathbf{p}_S)), \quad \text{where,} \tag{20}$$

$$\ln(\text{Market Size}_t) = \gamma_0 + \gamma_1 \cdot t + \gamma_2 \cdot (T - t) + \gamma_3 \cdot \mathbb{1}_{\text{hldy}_t}$$

We found that the coefficients for personalized loyalty features such as cumulative expenditure and cumulative visits were all positive and significant, indicating that these metrics were effective in capturing individualized information indicative of purchase preferences. The out of sample WMAPE was approximately 0.40 when averaged across all loyalty and demand groups. Furthermore, the estimated coefficients across all of the models intuitively corresponded to realistic consumer choice behaviors. For example, the coefficients related to seasonality were all strongly positive, reinforcing the idea that peak periods and popularity drive consumers to have a higher propensity-to-buy. Conversely, steep clearance periods resulted in negative coefficients as this corresponds to scenarios in which the prime life span of the seasonal good has expired. This estimation approach provided us with a model for every (loyalty group, product bundle) pair across all demand groups, which we leveraged dynamically to make personalized bundle offers.

B.3. Analyzing the Effects of Marginal Costs and Competition

Marginal Costs

The original retail data set contained marginal cost information for each product offered at each point in time, which was incorporated into our analysis. To gain a greater understanding of how marginal costs impact profitability, we studied the largest demand group of seasonal holiday decor from the retail data consisting of a mix of individual kitchen products (kitchen towels) and associated complementary ones (potholders, etc). Within this group we considered three settings where marginal costs were: (1) left at their actual historical values, (2) increased by 50%, and (3) decreased by 50%. The results of this study are summarized in Table 10 below, in which the percentage values reflect the average discount amount offered to each loyalty group under marginal cost scenarios (1)-(3).

	Consumer Loyalty Group	Costs Down 50%	Actual Costs	Costs Up 50%
SIAA	Loyalty 1 (One-Time Shoppers)	4.1%	2.9%	1.2%
	Loyalty 2 (Repeat, < Mean Expenditure)	3.2%	2.4%	1.0%
	Loyalty 3 (Repeat, > Mean Expenditure)	2.0%	1.7%	0.8%
EMPA	Loyalty 1 (One-Time Shoppers)	6.2%	4.6%	2.5%
	Loyalty 2 (Repeat, < Mean Expenditure)	4.7%	3.5%	1.9%
	Loyalty 3 (Repeat, > Mean Expenditure)	3.3%	2.7%	1.5%

Table 10 This table summarizes the personalized discounts by loyalty group for a bundle containing a kitchen towel and a complementary product such as a potholder (bundle T&Misc).

These results mirror our original retail study in which the consumers in Loyalty Group 3 have the least price elasticity and therefore receive the least amount of relative change in discount (only 0.5-1.5% compared to the range of 1-2.5% in the lower loyalty groups) regardless of marginal costs. SIAA is overall less affected as EMPA is a more profit-driven algorithm that makes larger discount adjustments when marginal costs change. Finally, the differences across the columns demonstrate that increasing costs has a larger effect on discount levels when compared to decreasing costs.

Consumer Loyalty Group	Costs Down 50%	Costs Up 50%
Loyalty 1 (One-Time Shoppers)	93.7%	91.5%
Loyalty 2 (Repeat, < Mean Expenditure)	95.1%	90.9%
Loyalty 3 (Repeat, > Mean Expenditure)	96.5%	89.7%
Total (Across All Loyalty Groups)	94.0%	91.2%

Table 11 This table summarizes the percentage of profitability achieved by the SIAA compared to the Clairvoyant model when marginal costs are increased or decreased 50% for all products. The performance of SIAA in the retail study was 93.4%.

We also analyzed the aggregate changes in profitability overall by loyalty group, as shown in Table 11. The percentages reflect the fraction of optimal Clairvoyant profit attained by SIAA when the marginal costs of all products were uniformly increased (or decreased) by 50% from their actual historical values. The effect of increasing marginal costs (second column) is larger within each loyalty group than the benefit of decreasing them (first column), which is also true in aggregate across all groups (last row) Higher costs result in profit loss of 1-4% whereas lowering costs causes profit gains of 0.5-3%. Ultimately, increasing marginal costs have a larger negative effect on profitability when compared to the gain of similarly decreasing costs.

Competition

The retail case study was not particularly subject to competition as the products were offered solely by the retailer. However, certain products from the airline case study could be subject to competition: (i) VIP lounge access, which is a space often shared by multiple airlines, and (ii) on-board gourmet meals, which are often cannibalized by consumers choosing to eat in the airport prior to the flight. We investigated the effect of competition through a numerical experiment consisting of a sensitivity analysis in which we changed the additive constant in the denominator of the MNL buy probability functions across the range of values in the set {1, 5, 10, 50, 100, 250}. The product-specific

results are shown in Table 12 below, where only one product was subject to this competition setting at a time. The percentages in this table represent the proportion of prior profit achieved by EMPA before competition was introduced. If we consider the row-wise and column-wise differences in both directions, the decrease in profitability as a result of competition has a more pronounced marginal effect in more inventory-constrained cases. This average loss in profitability relative to prior performance (without competition) can reach up to 20-25% depending on the nominal price of the product facing competition.

	Denominator	Initial Inventory Level					
		Unlimited	100%	90%	80%	75%	50%
VIP Lounge Access (\$50)	1	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
	5	98.3%	97.7%	96.9%	96.0%	94.9%	93.4%
	10	95.6%	94.8%	93.6%	92.5%	91.3%	90.0%
	50	92.9%	91.6%	89.8%	89.2%	88.1%	85.8%
	100	89.2%	88.5%	86.0%	85.2%	83.2%	81.2%
	250	85.4%	83.8%	82.6%	79.7%	77.6%	75.5%
On-Board Gourmet Meals (\$15)	1	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
	5	99.2%	99.0%	98.8%	98.5%	97.7%	96.6%
	10	96.8%	96.6%	95.9%	95.3%	94.4%	93.7%
	50	94.8%	94.0%	93.1%	92.6%	91.9%	91.2%
	100	92.3%	91.7%	89.8%	89.2%	88.7%	88.2%
	250	88.9%	88.3%	87.7%	85.6%	84.2%	82.9%

Table 12 This table presents the average expected percentage of prior revenue attained by the EMPA with the introduction of competition into the products VIP lounge access and on-board gourmet meals in the airline case study.

We also observe a marked difference between the two products due to their prices: the differences in proportion of prior non-competitive profitability captured ranges from 3.5% to up to 7.5% (with lounge access recovering less profit than gourmet meals). As initial inventory levels become more constrained, the margins of under-performance increasingly grow, mirroring our initial non-competition results. In addition, the additive approach outperforms the multiplicative method and is overall less affected by the introduction of competition, as shown in Table 13 below. The percentages in this table are the expected average profit of each algorithm relative to the Clairvoyant model when it also faces competition in the same product (as opposed to the proportion of prior profit captured by the same algorithm without competition as in Table 12). In these results the MNL constants are changed in both algorithms as well as in the Clairvoyant for a specific fixed product. Table 13 analyzes only two inventory levels and records the difference in performance between SIAA and EMPA. The outperformance of SIAA over EMPA grows increasingly as the level of competition increases. We observe up to a 3% gap in performance between the two algorithms when we compare them to an optimal Clairvoyant strategy facing the same level of competition.

Initial Inventory	Method	Product	Denominator					
			1	5	10	50	100	250
Unlimited	SIAA	Lounge	96.9%	95.8%	95.1%	94.2%	93.3%	92.6%
		Meals	96.9%	96.2%	95.8%	95.2%	94.7%	94.5%
	EMPA	Lounge	97.4%	96.2%	95.4%	93.6%	91.5%	90.7%
		Meals	97.4%	96.6%	96.1%	95.2%	94.2%	93.6%
75%	SIAA	Lounge	88.6%	87.8%	86.7%	85.9%	84.8%	84.1%
		Meals	88.6%	88.2%	87.8%	87.2%	86.3%	86.0%
	EMPA	Lounge	87.3%	86.4%	85.2%	84.1%	82.7%	80.7%
		Meals	87.3%	86.9%	86.3%	85.5%	84.4%	83.1%

Table 13 This table shows the average expected percentage of Clairvoyant revenue attained by the both SIAA and EMPA with the introduction of competition into the setting for VIP lounge access and on-board gourmet meals in the airline case study.