

Algorithms for Online Matching, Assortment, and Pricing with Tight Weight-dependent Competitive Ratios

Will Ma

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, willma@mit.edu

David Simchi-Levi

Institute for Data, Systems, and Society, Department of Civil and Environmental Engineering, and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, dslevi@mit.edu

Motivated by the dynamic assortment offerings and item pricings occurring in e-commerce, we study a general problem of allocating finite inventories to heterogeneous customers arriving sequentially. We analyze this problem under the framework of competitive analysis, where the sequence of customers is unknown and does not necessarily follow any pattern. Previous work in this area, studying online matching, advertising, and assortment problems, has focused on the case where each item can only be sold at a single price, resulting in algorithms which achieve the best-possible competitive ratio of $1-1/e$.

In this paper, we extend all of these results to allow for items having multiple feasible prices. Our algorithms achieve the best-possible weight-dependent competitive ratios, which depend on the sets of feasible prices given in advance. Our algorithms are also simple and intuitive; they are based on constructing a class of universal “value functions” which integrate the selection of items and prices offered.

Finally, we test our algorithms on the publicly-available hotel data set of Bodea et al. (2009), where there are multiple items (hotel rooms) each with multiple prices (fares at which the room could be sold). We find that applying our algorithms, as a “hybrid” with algorithms which attempt to forecast and learn the future transactions, results in the best performance.

1. Introduction

In this paper we study a general online resource allocation problem, motivated by dynamic assortment and pricing in revenue management. Consider an airline website selling parallel flights, i.e.

different flights which depart from the same origin to the same destination around the same time. Each flight corresponds to an *item* which could be sold, and its seat capacity corresponds to the unreplenishable *starting inventory* of that item. Each flight has multiple fare classes (e.g. Economy, Basic Economy) which correspond to *prices* at which that item could be sold. We will refer to this collection of initial information (items, inventories, prices) as the *setup*.

Over the booking horizon, heterogeneous customers sequentially arrive to the airline's website. We assume that the airline can reliably estimate each customer's *choice probabilities* from historical data. That is, upon a customer's arrival, for any combination of items and prices that could be shown, the stochastic distribution of how the customer would choose among those items/prices is given. The customer is assumed to choose at most one item and one price, as the flights are parallel. The customer could also choose to make no purchase. Given the choice probabilities, the airline selects an *assortment* of items and corresponding prices to show the customer, where items with zero remaining inventory cannot be shown. The customer's decision is realized immediately afterward, and if she makes a purchase, then the airline earns the corresponding price as revenue, and depletes one unit of inventory of the corresponding item. The airline wants to maximize its cumulative revenue earned before the booking horizon is over (or all the flights are full).

We study this problem under the framework of competitive analysis, where the sequence of customers to arrive over the booking horizon is unknown and does not necessarily follow any pattern. Instead, the airline seeks to have a good relative performance on all possible sequences. It offers assortments and pricings using a (possibly randomized) *online algorithm*, which can make decisions based on only the setup and the arrival sequence/purchase realizations seen so far. For $c \leq 1$, the online algorithm is said to be *c-competitive*, or *achieve a competitive ratio of c*, if

$$\inf_{\text{arrival sequences } \mathcal{A}} \frac{\mathbb{E}[\text{ALG}(\mathcal{A})]}{\text{OPT}(\mathcal{A})} \geq c, \tag{1}$$

where $\mathbb{E}[\text{ALG}(\mathcal{A})]$ denotes the algorithm's expected revenue on arrival sequence \mathcal{A} , and $\text{OPT}(\mathcal{A})$ denotes the value of an optimum which knows the entirety of \mathcal{A} in advance. In this paper, we will allow the competitive ratio guarantee to be *setup-dependent*; that is, the value of c in (1) can be a

function of the items, their starting inventories, and prices. We are interested in online algorithms which achieve the best-possible competitive ratios for various families of setups.

1.1. Overview of Result, and Relation to Previous Results

For setups where each item has a single fare class, the competitive ratio of the above problem has been analyzed extensively under many streams of literature, which we review below.

1. **Online Assortment:** If each item has a single price, then the above problem formulation is exactly the online assortment problem of Golrezaei et al. (2014). The authors use an algorithm which judiciously “balances” between offering different items, based on their remaining inventory levels. They show, among other results, that their algorithm is $(1 - 1/e)$ -competitive in the *asymptotic regime*. That is, their value of c in (1) depends on the smallest starting inventory amount in the setup, and approaches $1 - 1/e$ as all starting inventories approach ∞ . Without large starting inventories, this problem has also been studied in the special case where all offered assortments must have size 1, in which case it becomes the *online matching with stochastic rewards* problem (Mehta and Panigrahi 2012, Mehta et al. 2014).
2. **Online Vertex-weighted Matching:** Consider the special case of the problem where the outcome of any assortment offering is deterministic, and given upon the customer’s arrival. In this case, we know the maximum (possibly 0) a customer is willing to pay for each item, and our decision can be reduced to selecting an item to offer to the customer at her maximum-willingness-to-pay (we can also offer no item). We will refer to this problem as the *deterministic case*; it can be viewed as an online weighted matching problem.

If each item is restricted to have a single price, then we get the *online vertex-weighted matching* problem of Aggarwal et al. (2011). The authors develop an algorithm which randomly “ranks” the items and matches higher-ranked items first, and show that it is $(1 - 1/e)$ -competitive. Their result generalizes the classical result of Karp et al. (1990) for the unweighted online bipartite matching problem (where all items have the *same* price).

3. **Adwords:** The Adwords problem of Mehta et al. (2007) is central to online advertising and features budget-constrained bidders instead of inventory-constrained items. It also uses the idea of “balancing”, between the bidders’ budgets in this case, to achieve $(1 - 1/e)$ -competitiveness in an asymptotic *small bids* regime. Although its budget constraints are not directly captured by our model, we show that in this asymptotic regime, the Adwords problem corresponds to a version of our problem where each item has a *single* price (despite each bidder having multiple bid values—we explain the reduction in Section 6).

Our main contribution, motivated by the parallel flights problem, is extending all of the preceding results to setups where items could have multiple prices, that are known in advance. Note that such a setup also arises naturally from models where the customers have been classified into “types”, and there is a “match quality” score between each item and each type—in that case, the price set of an item consists of the item’s possible match scores.

However, allowing for multiple prices per item runs into a known impossibility result: even in the deterministic case, which corresponds to the aforementioned online weighted matching problem, it is not possible to provide a non-zero competitive ratio guarantee c which holds for every multi-price setup (the way $c = 1 - 1/e$ was a constant guarantee for every single-price setup). This is because the moment we commit to a match, unboundedly larger edge weights can arrive afterward—see Mehta (2013, Ch. 7). Therefore, previous work in online weighted matching has assumed that matches can be freely disposed if larger weights arrive later (Feldman et al. 2009), or that arrivals appear in a random order (Kesselheim et al. 2013).

In our paper, we instead assume that the price sets (i.e., the possible edge weights) are known in advance, and derive *weight-dependent* competitive ratio guarantees, where our value of c in (1) will depend on the setup; specifically, the price sets $\mathcal{P}_1, \dots, \mathcal{P}_n$ of the n items. Our algorithms also make use of the knowledge of the price sets. Our competitive ratio results are based on establishing a universal mapping F from price sets \mathcal{P} to ratios in $[0, 1 - 1/e]$, such that:

1. Our MULTI-PRICE BALANCE algorithm, which extends the existing “inventory balancing” algorithms, is $\min_i F(\mathcal{P}_i)$ -competitive in the asymptotic regime;

2. Our MULTI-PRICE RANKING algorithm, which extends the existing “randomized ranking” algorithm, is $\min_i F(\mathcal{P}_i)$ -competitive in the deterministic case;
3. Any (deterministic or randomized) algorithm can be at most $\min_i F(\mathcal{P}_i)$ -competitive for the family of setups with price sets chosen from $\mathcal{P}_1, \dots, \mathcal{P}_n$, even if we restrict the setups to have asymptotic starting inventories and/or deterministic arrival sequences.

For any singleton price set \mathcal{P} , $F(\mathcal{P}) = 1 - 1/e$, and hence if $|\mathcal{P}_1| = \dots = |\mathcal{P}_n| = 1$, then our results recover existing results: the $(1 - 1/e)$ -competitiveness of inventory balancing in the asymptotic regime, the $(1 - 1/e)$ -competitiveness of randomized ranking in the deterministic case, and a single counterexample which shows that both of these algorithms are tight. $F(\mathcal{P})$ approaches 0 if \mathcal{P} contains both a large number of prices and large ratios between its prices, so our statement 3 also recovers the known impossibility result.

1.2. A Bid Price Algorithm when Items have Multiple Prices

We illustrate the necessity for our new algorithms by comparing MULTI-PRICE BALANCE to the existing inventory balancing algorithm of Golrezaei et al. (2014); similar arguments can be made in comparisons to the existing online vertex-weighted matching and Adwords algorithms.


Suppose there are parallel flights, whose seats have the same two fare classes: a lower price of $r^{(L)} = 150$, and a higher price of $r^{(H)} = 450$. At any point in time, for each flight i , let w_i denote the fraction its starting inventory which has been sold. The algorithm of Golrezaei et al. (2014) would associate each fare class $j \in \{L, H\}$ of each flight i with a “pseudorevenue” equal to

$$r^{(j)} \cdot \Psi(w_i), \tag{2}$$

where Ψ is a decreasing function that *penalizes* the revenues associated with flights i which are almost full. The algorithm then offers, to each customer, the assortment which maximizes the expected pseudorevenue of the (flight, fare)-combination that the customer would choose.

In (2), although $\Psi(w_i)$ will disincentivize the offering of a flight i whose w_i is large, the algorithm has no way of setting a “booking limit”—preventing sales at the lower price $r^{(L)}$ while still allowing

Figure 1 The Φ function for an item with feasible price set $\{150, 450\}$.



sales at price $r^{(H)}$. Given a stream of customers who are only interested in the lower price, the algorithm would sell all the seats at price $r^{(L)}$, without realizing the opportunity cost of $r^{(H)}$ it gave up. Since this could happen to every flight i , the algorithm’s competitiveness is at most $\frac{r^{(L)}}{r^{(H)}}$.

To improve upon this, our algorithm must implement some notion of “booking limits”. As a result, we define the pseudorevenue associated with fare class j of flight i to be

$$r^{(j)} - \Phi(w_i), \quad (3)$$

where Φ is an increasing function that sets a *cost* to selling flights i which are almost full. MULTIPRICE BALANCE uses the same idea of maximizing expected pseudorevenue, with the modification that it rejects the customer outright if the maximum expected pseudorevenue is non-positive.

The exact form of function Φ for this example is shown in Figure 1. Note that for a flight i , if the fraction sold w_i is between α and 1, then the pseudorevenue (3) will be negative for the lower price of 150 and positive for the higher price 450, producing a desirable “booking limit” at $w_i = \alpha$. Other than that, Φ still produces a continuously-increasing cost as w_i increases from 0 to 1, allowing us to trade off between offering the different flights i based on their values of w_i .

In (3), $\Phi(w_i)$ can be interpreted as a *bid price*, or the value placed on one unit of item i ’s inventory. Optimizing based on bid prices is a classical idea in revenue management (see Talluri and Van Ryzin (2006), Liu and Van Ryzin (2008)), where typically the bid prices are computed using a large LP which encompasses both the inventories and the forecasted distribution of future customers. However, since we make no assumptions about future customers, our bid prices are based on only the remaining inventories, like the balance algorithms from competitive analysis.

1.3. Our Competitive Ratio Guarantees

In general, for any price set \mathcal{P} we define a construction $\Phi_{\mathcal{P}}$ which we call a *value function*. In a setup, if $\mathcal{P}_1, \dots, \mathcal{P}_n$ denotes the price sets of the items, then MULTI-PRICE BALANCE defines the pseudorevenues of each item i using value function $\Phi_{\mathcal{P}_i}$ in expression (3). Our MULTI-PRICE RANKING algorithm uses the same value functions, but applies them to a “random seed” instead.

Note that for each item i , the construction of $\Phi_{\mathcal{P}_i}$ from \mathcal{P}_i is universal in that it does not depend on other parameters in the setup, e.g. the price sets of the other items. Our mapping F from price sets to ratios was also universal. The fact that separately determining the value function $\Phi_{\mathcal{P}_i}$ for each item i leads to the best-possible competitive ratio of $\min_i F(\mathcal{P}_i)$ is, in our opinion, very surprising—see also the discussions in Devanur and Jain (2012), Devanur et al. (2013). For any \mathcal{P} , our exact derivation of $\Phi_{\mathcal{P}}$ and $F(\mathcal{P})$ comes from the solution of a differential equation, which arises from a primal-dual analysis based on Buchbinder et al. (2007).


When $|\mathcal{P}| = 1$, with $\mathcal{P} = \{r\}$, value function $\Phi_{\mathcal{P}}(w) = r(1 - \Psi(w))$, and hence our notion of pseudorevenue in (3) coincides with the existing notion in (2). If each item has a single price, then our algorithms will coincide with the existing ones.

We now give a flavor of our new results with $|\mathcal{P}| = 2$. Let $\mathcal{P} = \{r, \xi r\}$, where $\xi > 1$ is the ratio from high price to low price. The value function $\Phi_{\mathcal{P}}$ depends on ξ (see Figure 1 for an example with $\xi = \frac{450}{150} = 3$). For any ξ , $F(\mathcal{P})$ equals

$$1 - \frac{\sqrt{1 + 4\xi(\xi - 1)/e} - 1}{2(\xi - 1)} =: F(\xi), \quad (4)$$

and the booking limit implied by $\Phi_{\mathcal{P}}$ (i.e. the corresponding value of α in Figure 1) equals $\ln(\frac{1}{1 - F(\xi)})$. We note that this is different from the booking limit of $\frac{\xi}{2\xi - 1}$ derived by Ball and Queyranne (2009), which is optimal for selling a *single* item whose price set is $\{r, \xi r\}$. For any value of ξ , the booking limit implied by our function $\Phi_{\{r, \xi r\}}$ is greater than $\frac{\xi}{2\xi - 1}$, which means that our algorithm is willing to sell a greater fraction of units at the lower price. The intuitive explanation of this is that with multiple items, there is less upside to reserving inventory for higher prices, because the reserved units still have to compete with other items to be sold.

Figure 2 Competitive ratios guarantees when each item has two prices. The guarantees increase from bottom to top (as the maximum ratio of an item’s high to low price, ξ_{\max} , ranges from ∞ to 1), and from left to right (as the minimum starting inventory, k_{\min} , ranges from 1 to ∞).



^b The smallest guarantee of $\frac{1}{4}$ in this diagram is also implied by the results of Chen et al. (2016).

If every item has two prices and ξ_{\max} denotes the maximum ratio of an item’s high price to low price, then $\min_i F(\mathcal{P}_i)$ equals $F(\xi_{\max})$ because F is decreasing. Also letting k_{\min} denote the minimum starting inventory among the items, we plot, in Figure 2, our competitive ratio guarantees for MULTI-PRICE BALANCE and MULTI-PRICE RANKING as both ξ_{\max} and k_{\min} range over $[1, \infty]$. This guarantee equals $F(\xi_{\max})$ in the asymptotic regime or deterministic case, which is tight.

The lower bound on the competitive ratio guarantee when each item has at most two prices occurs as $\xi_{\max} \rightarrow \infty$, in which case $F(\xi_{\max}) = 1 - \frac{1}{\sqrt{e}} \approx 0.393$. This is greater than the naive bound of $\frac{1}{2}(1 - \frac{1}{e})$, which would arise from randomly choosing between 2 prices and then using a $(1 - \frac{1}{e})$ -competitive algorithm on the chosen prices. Thus, using a function like Φ to integrate the selection of prices with the allocation across items is necessary for achieving the optimal competitive ratio.

Our guarantees may not be tight in the non-asymptotic, non-deterministic setting, which is an important open problem even in the single-price case (Devanur et al. 2013). Nonetheless, as k_{\min} increases, our bounds sharply approach the tight guarantee from the asymptotic regime. In the

single-price case, our bound is a factor of $(1 + k_{\min})(1 - e^{-1/k_{\min}})$ from the tight guarantee of $1 - 1/e$, improving the previous-best-known dependence on k_{\min} from Golrezaei et al. (2014).

1.4. Simulations on Hotel Data Set of Bodea et al. (2009)

We first summarize the general benefits of applying the algorithms from competitive analysis. In contrast to traditional algorithms, which optimize based on a forecast of future demand, or attempt to learn the demand, competitive algorithms guarantee some performance ratio under the worst case, and operate without any demand information. Most immediately, they are useful for products with highly unpredictable demand (Ball and Queyranne 2009, Lan et al. 2008), or for initializing new products with no historical sales data (Van Ryzin and McGill 2000). Second, by eschewing stochastic processes for generating demand, competitive algorithms are usually simple and flexible, leading to clean insights about the problem (Borodin and El-Yaniv 2005). Third, past research has reported on cases where competitive algorithms perform well in practice (Feldman et al. 2010), or on average in numerical experiments (Golrezaei et al. 2014, Chen et al. 2016).

In Section 7, we run simulations on the publicly-accessible hotel data set of Bodea et al. (2009). We use the product availability information to estimate customer choice models, and use the sequence of transactions as the sequence of arrivals. This leads to an online assortment problem like in Golrezaei et al. (2014), but with multiple prices (advance-purchase rate, rack rate, etc.) for each item (King room, Two-double room, etc.). We compare the performance of our MULTI-PRICE BALANCE algorithm to various benchmarks and forecasting algorithms.

The main conclusion from our simulations is that the best performance is achieved by *hybrid* algorithms (see Golrezaei et al. (2014)). These are forecasting-based algorithms which continuously reference our forecast-independent value functions Φ_1, \dots, Φ_n , and adjust their decisions accordingly. Although this only changes a small fraction ($\approx 5\%$) of decisions, these tend to be the decisions where the forecast is being most overconfident. Therefore, not only does this boost average performance, it drastically reduces the variance in performance caused when the forecast is wrong.

1.5. Other Related Work

We briefly mention some papers which study online resource allocation problems under other arrival models or performance metrics.

When a stochastic process generating the arrivals is given as input, the resulting optimization problem is generally still computationally intractable. Nonetheless, many effective heuristics have been proposed under various models of online resource allocation (Zhang and Cooper 2005, Jasin and Kumar 2012, Ciocan and Farias 2012, Chen and Farias 2013). These heuristics can earn $\frac{1}{2}$ of the LP optimum in general settings (Chan and Farias 2009, Wang et al. 2015). Manshadi et al. (2012) derive an improved performance ratio when the stochastic process is IID.

Competitive/approximation ratios both analyze the *fraction* of optimum achieved by an algorithm, but online resource allocation problems are also often analyzed under the *regret* metric, which measures the *difference* from optimum. This work often focuses on learning some unknown underlying stochastic model (Badanidiyuru et al. 2013, Ferreira et al. 2016). On the other hand, queueing-theoretic analyses have also been performed given a known stochastic model (Reiman and Wang 2008). Unlike in competitive analysis, all of these papers tend to focus on asymptotic performance as the number of customers grows to infinity. Finally, a recent metric which has been introduced is *regret ratio* (Zhang et al. 2016). For a comprehensive review of different metrics under different models of demand (for a single item), we refer to Araman and Caldentey (2011).

1.6. Organization of Paper

Throughout Sections 2–5 of this paper, we analyze a simplified model where each customer is *offered a single item at a single price* (but her purchase decision is still stochastic). This avoids the complexities of assortment optimization while still capturing our main techniques. In Section 6, we discuss the generalizations to the assortment and Adwords settings. In Section 7, we display the results of our simulations on the hotel data set.

2. Problem Definition, Algorithm Sketch, and Theorem Statements

A firm is selling $n \in \mathbb{N}$ different items. Each item $i \in [n]^1$ starts with a fixed inventory of $k_i \in \mathbb{N}$ units, and could be offered at any price in its price set \mathcal{P}_i . Throughout most of this paper, we assume that each \mathcal{P}_i consists of $m_i \in \mathbb{N}$ discrete prices satisfying $0 < r_i^{(1)} < \dots < r_i^{(m_i)}$. We will refer to $r_i^{(j)}$ as “price j of item i ”, and define $r_i^{(0)} := 0$. We extend to the case where \mathcal{P}_i is a continuum of prices in Appendix E.1.

There are $T \in \mathbb{N}$ customers arriving sequentially. Upon the arrival of customer $t \in [T]$, the firm is given $p_{t,i}^{(j)}$, the probability² that customer t would buy item i at price j , for all $i \in [n]$ and $j \in [m_i]$.³ The firm chooses up to one of the items i with remaining inventory and offers it to customer t , at any price $j \in [m_i]$. The customer accepts the offer with probability $p_{t,i}^{(j)}$, in which case the firm earns revenue $r_i^{(j)}$, and the inventory of item i is decremented by 1.

We divide the elements defined above into:

1. The *Setup* \mathcal{S} , consisting of parameters known at the start: $(n, (k_i, m_i, r_i^{(1)}, \dots, r_i^{(m_i)})_{i \in [n]})$; and
2. The *Arrival sequence* \mathcal{A} , consisting of parameters revealed over time: $(T, (p_{t,i}^{(j)})_{t \in [T], i \in [n], j \in [m_i]})$.

An *online algorithm* must decide, on any setup \mathcal{S} , what to offer to each customer t . This decision can be based on only the setup \mathcal{S} , the past arrivals/purchase realizations, and the purchase probabilities $p_{t,i}^{(j)}$ of the present customer t ; the online algorithm does not know the purchase probabilities associated with future customers. For an online algorithm, let $\text{ALG}(\mathcal{S}, \mathcal{A})$ denote the revenue earned on a run on setup \mathcal{S} with arrival sequence \mathcal{A} , which is a random variable with respect to the customers’ purchase decisions as well, as any randomness in the algorithm’s decisions.

Meanwhile, we can write the following LP for setup \mathcal{S} with arrival sequence \mathcal{A} :

$$\max \sum_{t=1}^T \sum_{i=1}^n \sum_{j=1}^{m_i} p_{t,i}^{(j)} r_i^{(j)} x_{t,i}^{(j)} \tag{5a}$$

¹ For a general positive integer b , let $[b]$ denote the set $\{1, \dots, b\}$.

² If \mathcal{P}_i is a continuum of prices, then we need to assume that the purchase probabilities can be input compactly. There are many parametric models for doing so, e.g. linear demand, where the purchase probability is $a - bP$ for prices P lying in an interval $[r^{\min}, r^{\max}]$.

³ These probabilities can be 0 for items the customer is not interested in, or prices that are too high.

$$\sum_{t=1}^T \sum_{j=1}^{m_i} p_{t,i}^{(j)} x_{t,i}^{(j)} \leq k_i \quad i \in [n] \quad (5b)$$

$$\sum_{i=1}^n \sum_{j=1}^{m_i} x_{t,i}^{(j)} \leq 1 \quad t \in [T] \quad (5c)$$

$$x_{t,i}^{(j)} \geq 0 \quad t \in [T], i \in [n], j \in [m_i] \quad (5d)$$

LP (5) encapsulates the execution of any algorithm, which could make full use of the arrival sequence \mathcal{A} at the start, on setup \mathcal{S} — $x_{t,i}^{(j)}$ represents the unconditional probability of the algorithm offering item i at price j to customer t ; (5b) enforces that starting inventories are respected; (5c) enforces that at most one combination of item and price is offered to each customer; and objective function (5a) represents the expected revenue earned by the algorithm. Let $\text{OPT}(\mathcal{S}, \mathcal{A})$ denote its optimal objective value. Note that although $\text{OPT}(\mathcal{S}, \mathcal{A})$ knows the arrival sequence in advance, it does not know the outcomes of the customers' potential purchase decisions.

For a fixed online algorithm and any setup \mathcal{S} , the online algorithm is said to *achieve a competitive ratio of c on \mathcal{S}* , if

$$\frac{\mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})]}{\text{OPT}(\mathcal{S}, \mathcal{A})} \geq c \quad \text{for all arrival sequences } \mathcal{A}. \quad (6)$$

In this paper, we will allow the competitive ratio guarantee c to depend on parameters in the setup \mathcal{S} , and derive results that hold for any \mathcal{S} .

Definition (6) provides a guarantee on $\mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})]$ relative to any algorithm which could have been possible, due to the following standard result.

LEMMA 1. *$\text{OPT}(\mathcal{S}, \mathcal{A})$ is an upper bound on the expected revenue of any algorithm, which could make full use of the arrival information at the start, on setup \mathcal{S} with arrival sequence \mathcal{A} .*

The proof of Lemma 1 is deferred to Appendix A. The definition of $\text{OPT}(\mathcal{S}, \mathcal{A})$ based on the LP is standard in problems with both stochastic purchase realizations and arbitrary customer arrivals—we refer to Mehta and Panigrahi (2012), Golrezaei et al. (2014) for its justification.

In the *deterministic* case of our problem, every $p_{t,i}^{(j)}$ is 0 or 1. The problem can be simplified by letting $j_{t,i} = \max\{j \in [m_i] : p_{t,i}^{(j)} = 1\}$, with $j_{t,i} = 0$ if the set is empty, for all $t \in [T]$ and $i \in [n]$. We

say that item i is *assigned* to customer t to indicate that i is offered to customer t at price $j_{t,i}$, which results in a sale; there is no reason to offer any other price. Customer t can also be *rejected*, e.g. if $j_{t,i}$ is low for every i . In the deterministic case, the LP (5) is integral, so $\text{OPT}(\mathcal{S}, \mathcal{A})$ is *equal* to the revenue of the best algorithm knowing the arrival sequence at the start.

2.1. Construction of Value Function for a Price Set

In this section, we specify a value function $\Phi_{\mathcal{P}}$ and a number $F(\mathcal{P})$, for any price set \mathcal{P} consisting of m discrete prices with $0 < r^{(1)} < \dots < r^{(m)}$. The derivation of $\Phi_{\mathcal{P}}$ and $F(\mathcal{P})$, as well as the case where \mathcal{P} is a continuum of prices, are deferred to Appendix E.

Consider an item with price set \mathcal{P} . Following the description from Section 1.2, we will interpret $\Phi_{\mathcal{P}}$ to be a function of $w \in [0, 1]$, which is the fraction of the item's starting inventory which has been sold. $\Phi_{\mathcal{P}}(w)$ specifies the value that should be placed on one unit of the item's inventory, when its fraction sold is w .

First we define “booking limits” $\alpha^{(1)}, \dots, \alpha^{(m)}$, which are the fractions of starting inventory “reserved” for the respective fares $r^{(1)}, \dots, r^{(m)}$, via the following proposition.

PROPOSITION 1. *Let $r^{(1)}, \dots, r^{(m)}$ be any numbers satisfying $0 < r^{(1)} < \dots < r^{(m)}$. Then there is a unique set of positive values $\alpha^{(1)}, \dots, \alpha^{(m)}$ which sum to 1 and satisfy*

$$1 - e^{-\alpha^{(1)}} = \frac{1}{1 - r^{(1)}/r^{(2)}} \cdot (1 - e^{-\alpha^{(2)}}) = \dots = \frac{1}{1 - r^{(m-1)}/r^{(m)}} \cdot (1 - e^{-\alpha^{(m)}}). \quad (7)$$

There is also a different, unique set of positive values $\sigma^{(1)}, \dots, \sigma^{(m)}$ which sum to 1 and satisfy

$$\sigma^{(1)} = \frac{1}{1 - r^{(1)}/r^{(2)}} \cdot \sigma^{(2)} = \dots = \frac{1}{1 - r^{(m-1)}/r^{(m)}} \cdot \sigma^{(m)}. \quad (8)$$

The proof of Proposition 1 is elementary and deferred to Appendix A. While finding the exact solution to (7) requires finding the roots of a degree- m polynomial, a numerical solution can easily be found via bisection search.

Proposition 1 contrasts $\alpha^{(1)}, \dots, \alpha^{(m)}$ in (7) with the booking limits $\sigma^{(1)}, \dots, \sigma^{(m)}$ in (8) originally derived by Ball and Queyranne (2009), which are optimal for selling a *single* item with price set $\{r^{(1)}, \dots, r^{(m)}\}$. With $\alpha^{(1)}, \dots, \alpha^{(m)}$, we can now complete the definition of $\Phi_{\mathcal{P}}$.

DEFINITION 1. Define the following, based on the values of $\alpha^{(1)}, \dots, \alpha^{(m)}$ from Proposition 1:

- $L^{(j)}$: the sum $\sum_{j'=1}^j \alpha^{(j')}$, defined for all $j = 0, \dots, m$ (note that $L^{(0)} = 0$ and $L^{(m)} = 1$);
- $\ell(\cdot)$: a function on $[0, 1]$, where $\ell(w)$ is the unique $j \in [m]$ for which $w \in [L^{(j-1)}, L^{(j)})$ (note that $\ell(L^{(j)}) = j + 1$ for $j = 0, \dots, m - 1$; we define $\ell(L^{(m)})$ to be m).

The *value function* $\Phi_{\mathcal{P}}$ for price set \mathcal{P} is then defined over $w \in [0, 1]$ by:

$$\Phi_{\mathcal{P}}(w) = r^{(\ell(w)-1)} + (r^{(\ell(w))} - r^{(\ell(w)-1)}) \frac{\exp(w - L^{(\ell(w)-1)}) - 1}{\exp(\alpha^{(\ell(w))}) - 1}. \quad (9)$$

An example of $\Phi_{\mathcal{P}}$ for $\mathcal{P} = \{150, 450\}$ was plotted in the Introduction, in Figure 1. In general, $\Phi_{\mathcal{P}}$ is continuously increasing and piecewise-convex over m *segments* of lengths $\alpha^{(1)}, \dots, \alpha^{(m)}$, separated by *segment borders* $L^{(0)}, \dots, L^{(m)}$. For each j , Φ reaches the value of $r^{(j)}$ at $L^{(j)}$.

DEFINITION 2. For price set $\mathcal{P} = \{r^{(1)}, \dots, r^{(m)}\}$, let $F(\mathcal{P}) = 1 - e^{-\alpha^{(1)}}$ and $G(\mathcal{P}) = \sigma^{(1)}$, where $\alpha^{(1)}$ and $\sigma^{(1)}$ are the values from Proposition 1.

Our competitive ratio guarantees will be based on the functions F and G . It can be checked that F maps a price set \mathcal{P} to $[1 - e^{-1/m}, 1 - e^{-1}]$ and G maps the price set to $[1/m, 1]$, where m is the number of prices in \mathcal{P} . When $m = 1$, our value function is $\Phi_{\mathcal{P}}(w) = r^{(1)} \cdot \frac{e^w - 1}{e - 1}$, which can be related back to the existing multiplicative “penalty functions” from the single-price case.

2.2. Sketch of our MULTI-PRICE BALANCE and MULTI-PRICE RANKING Algorithms

Having defined the value function $\Phi_{\mathcal{P}}$ for an arbitrary price set \mathcal{P} , we now sketch our algorithms.

We start with MULTI-PRICE RANKING, which is simpler. It assumes that $k_i = 1$ for all i , which does not lose generality since an item that starts with multiple units of inventory can be transformed into multiple disparate items. At the start, the algorithm fixes for each item i a random seed W_i , drawn independently and uniformly from $[0, 1]$. It then treats $\Phi_{\mathcal{P}_i}(W_i)$ as the bid price for the single unit of item i : it offers to each customer t the available item i and price j maximizing the expected pseudorevenue, $p_{t,i}^{(j)}(r_i^{(j)} - \Phi_{\mathcal{P}_i}(W_i))$.

MULTI-PRICE RANKING hedges against the ambiguity in customer arrivals using randomness, which is standard in competitive analysis. The random seed W_i determines the random minimum

price at which the algorithm is willing to sell item i , as well as a random priority for selling i when the algorithm is choosing between multiple items.

We now sketch MULTI-PRICE BALANCE, which updates the bid price of each item i based on the fraction w_i of its k_i units which has been sold. However, the algorithm does not directly use $\Phi_{\mathcal{P}_i}(w_i)$ as the bid price of item i , because w_i would always be a multiple of $\frac{1}{k_i}$, while the booking limits and segment borders which $\Phi_{\mathcal{P}_i}$ is based on may not be multiples of $\frac{1}{k_i}$. Instead, the algorithm first uses a *randomized procedure* for rounding the booking limits in $\Phi_{\mathcal{P}_i}$ to multiples of $\frac{1}{k_i}$.

Specifically, at the start, the algorithm fixes for each item i *random* segment borders $\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)}$, which are multiples of $\frac{1}{k_i}$ satisfying $0 = \tilde{L}_i^{(0)} \leq \dots \leq \tilde{L}_i^{(m_i)} = 1$. We note that having $\tilde{L}^{(i)} = \tilde{L}^{(i-1)}$ is possible (and guaranteed to happen if $m_i > k_i$), in which case the i 'th segment has length zero. In either case, the realizations of $\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)}$ imply a random value function $\tilde{\Phi}_i$ for item i , which is a *perturbation* of $\Phi_{\mathcal{P}_i}$. Function $\tilde{\Phi}_i$ is defined on $\{0, \frac{1}{k_i}, \dots, 1\}$, since the fraction sold w_i is always a multiple of $\frac{1}{k_i}$, and also satisfies $0 = \tilde{\Phi}_i(0) \leq \tilde{\Phi}_i(\frac{1}{k_i}) \leq \dots \leq \tilde{\Phi}_i(1)$. At any point in time, MULTI-PRICE BALANCE treats $\tilde{\Phi}_i(w_i)$ as the bid price for item i : it offers to each customer t the item i and price j maximizing

$$p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(w_i)). \quad (10)$$

In (10), the definition of pseudorevenue at price j is $\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(w_i)$ instead of $r_i^{(j)} - \tilde{\Phi}_i(w_i)$. This is because we want the expected pseudorevenue to be 0 when $w_i = \tilde{L}_i^{(j)}$. In general, the realized $\tilde{\Phi}_i$ will be close to $\Phi_{\mathcal{P}_i}$, so that $\tilde{\Phi}_i(\tilde{L}_i^{(j)}) \approx \Phi_{\mathcal{P}_i}(L_i^{(j)}) = r_i^{(j)}$. In the asymptotic regime with $k_i \rightarrow \infty$, $\tilde{\Phi}_i$ is deterministically initialized to $\Phi_{\mathcal{P}_i}$. However, for small k_i , optimizing a *randomized procedure* for initializing $\tilde{\Phi}_i$ (based on $r_i^{(1)}, \dots, r_i^{(m_i)}$ as well as k_i) instead of having the deterministic $\Phi_{\mathcal{P}_i}$ (which is based on only $r_i^{(1)}, \dots, r_i^{(m_i)}$) allows us to achieve a greater competitive ratio.

2.3. Statements of Our Results

Our competitive ratio results are based on the universal functions F and G from Definition 2, which assign a number to every price set \mathcal{P} .

THEOREM 1. *For any setup, with price sets denoted by $\mathcal{P}_1, \dots, \mathcal{P}_n$ and starting inventories denoted by k_1, \dots, k_n , MULTI-PRICE BALANCE achieves a competitive ratio of $\min_i \tilde{F}_i$, where for each item i , \tilde{F}_i is lower-bounded by all of: (i) $\frac{F(\mathcal{P}_i)}{(1+k_i)(e^{1/k_i}-1)}$; (ii) $\frac{G(\mathcal{P}_i)}{2}$; and (iii) $\frac{1-1/\epsilon}{(1+k_i)(1-e^{-1/k_i})}$ if $|\mathcal{P}_i| = 1$.*

COROLLARY 1. *MULTI-PRICE BALANCE achieves a competitive ratio approaching $\min_i F(\mathcal{P}_i)$ as each starting inventory k_i approaches ∞ .*

COROLLARY 2. *Suppose that each item has at most m discrete prices and at least k units of starting inventory. Then the competitive ratio achieved by MULTI-PRICE BALANCE is lower-bounded by $\frac{1-e^{-1/m}}{(1+k)(e^{1/k}-1)}$, which approaches $1 - e^{-1/m}$ as k approaches ∞ .*

Theorem 1 is our general result, where for each i , \tilde{F}_i is determined by the randomized procedure used to initialize the value function $\tilde{\Phi}_i$ for item i .

Lower bound (i) on \tilde{F}_i is attained by a randomized procedure which perturbs the “ideal” value function $\Phi_{\mathcal{P}_i}$ to define $\tilde{\Phi}_i$. This perturbation loses a factor of $(1+k_i)(e^{1/k_i}-1)$ in the denominator, which decreases to 1 as $k_i \rightarrow \infty$, resulting in Corollary 1. Corollary 2 is a further simplification of the bound presented, using the fact that $F(\mathcal{P}) \geq 1 - e^{-1/|\mathcal{P}|}$. Meanwhile, lower bound (ii) is attained by solving an optimization problem for the best randomized procedure to define $\tilde{\Phi}_i$ when $k_i = 1$; this procedure is not based on perturbing $\Phi_{\mathcal{P}_i}$ and the bound is based on $G(\mathcal{P}_i)$ instead of $F(\mathcal{P}_i)$. Finally, lower bound (iii) is an improvement of (i) in the single-price case, where we have gained a factor of e^{1/k_i} in the denominator. It simplifies and improves the dependence on k_i from Golrezaei et al. (2014).

MULTI-PRICE BALANCE is formalized and Theorem 1 is proven in Section 3. We explain the ideas behind our primal-dual analysis, why we need *random* value functions, and how to overcome the ensuing analytical challenges.

THEOREM 2. *For any setup in the deterministic case, MULTI-PRICE RANKING achieves a competitive ratio of $\min_i F(\mathcal{P}_i)$.*

MULTI-PRICE RANKING is formalized and Theorem 2 is proven in Section 4. Our analysis builds upon the framework of Devanur et al. (2013) and extends it to handle multiple prices.

THEOREM 3. *Let k be any positive integer. Let \mathcal{P} be any price set consisting of m prices with $0 < r^{(1)} < \dots < r^{(m)}$. Then there exists a setup \mathcal{S} where each item has starting inventory k and price set \mathcal{P} , along with a distribution over arrival sequences \mathcal{A} falling in the deterministic case, for which no online algorithm can have expected revenue greater than $F(\mathcal{P}) \cdot \mathbb{E}_{\mathcal{A}}[\text{OPT}(\mathcal{S}, \mathcal{A})]$.*

Theorem 3 is proven in Section 5. Since the starting inventory k can be made arbitrarily large and the arrival sequences fall in the deterministic case, Theorem 3 implies that the competitive ratio guarantees in Corollary 1 and Theorem 2 are tight, via Yao’s minimax principle (Yao 1977).

Our counterexample is based on those from Karp et al. (1990), Mehta et al. (2007), Golrezaei et al. (2014), where a large number of customers arrive according to a random permutation chosen uniformly from all possible permutations. In our case however, the customers are further split into m “phases”, where the customers in phase j are willing to pay $r^{(j)}$ for any of the items they are interested in. The phases lengths are optimized by an adversary to minimize the competitive ratio.

Interestingly, on the existing counterexamples, the random permutation implies that all (reasonable) algorithms are indifferent and have the same performance. By contrast, on our counterexample with the adversarially-optimized phase lengths, there is a *unique* optimal algorithm given the distribution over arrival sequences \mathcal{A} . When $k \rightarrow \infty$, this unique algorithms turns out to be our MULTI-PRICE BALANCE and MULTI-PRICE RANKING algorithms, which coalesce to the same algorithm in the asymptotic regime. This coalescence phenomenon has been noted in the single-price case as well by Aggarwal et al. (2011).

PROPOSITION 2. *For $m \geq 2$ prices satisfying $0 < r^{(1)} < \dots < r^{(m)}$, from which $\alpha^{(1)}$ and $\sigma^{(1)}$ are defined according to Proposition 1, the following inequalities hold:*

$$\left(1 - \frac{1}{e}\right) \cdot \sigma^{(1)} < 1 - e^{-\sigma^{(1)}} < 1 - e^{-\alpha^{(1)}}; \quad (11)$$

$$\frac{1}{1 + \ln \frac{r^{(m)}}{r^{(1)}}} < \sigma^{(1)}; \quad (12)$$

$$1 - e^{-\alpha} < 1 - e^{-\alpha^{(1)}}, \text{ where } \alpha \text{ is the unique solution to } 1 - e^{-\alpha} = \frac{1 - \alpha}{\ln \frac{r^{(m)}}{r^{(1)}}}. \quad (13)$$

Finally, Proposition 2, which is proven in Appendix A, puts our tight competitive ratio of $1 - e^{-\alpha^{(1)}}$ into perspective. $\sigma^{(1)}$ is the existing tight competitive ratio for a single item, while $1 - \frac{1}{e}$ is the existing tight competitive ratio for multiple items with one price each. (11) shows that our competitive ratio for multiple items with multiple prices is not a naive combination of the existing competitive ratios; our algorithms also cannot be obtained by naively combining existing algorithms.

With a single item whose price can take any value in the continuum $[r^{(1)}, r^{(m)}]$, the tight competitive ratio is $\frac{1}{1 + \ln(r^{(m)}/r^{(1)})}$ (Ball and Queyranne 2009). (12) says that if the prices are restricted to a discrete subset of $[r^{(1)}, r^{(m)}]$, then the competitive ratio of $\sigma^{(1)}$ can only be larger.

We have a corresponding relationship in the multi-price setting. $1 - e^{-\alpha}$, with α as defined⁴ in (13), is our competitive ratio when there are multiple items whose price sets are $[r^{(1)}, r^{(m)}]$. (13) says that if the prices are restricted to a discrete subset of $[r^{(1)}, r^{(m)}]$, then the competitive ratio of $1 - e^{-\alpha^{(1)}}$ can only be larger.

3. MULTI-PRICE BALANCE and the Proof of Theorem 1

MULTI-PRICE BALANCE, as sketched in Subsection 2.2, is formalized in Algorithm 1. For now, we consider a generic randomized procedure for initializing $\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)}$ and $\tilde{\Phi}_i$ in Step 1, where the realized initializations always satisfy the following monotonicity conditions:

$$\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)} \in \{0, \frac{1}{k_i}, \dots, 1\}, \quad 0 = \tilde{L}_i^{(0)} \leq \dots \leq \tilde{L}_i^{(m_i)} = 1; \quad (14)$$

$$\tilde{\Phi}_i(0), \tilde{\Phi}_i(\frac{1}{k_i}), \dots, \tilde{\Phi}_i(1) \in \mathbb{R}, \quad 0 = \tilde{\Phi}_i(0) \leq \tilde{\Phi}_i(\frac{1}{k_i}) \leq \dots \leq \tilde{\Phi}_i(1). \quad (15)$$

Since $\tilde{\Phi}_i$ is non-decreasing, the expression $\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_i}{k_i})$ in (16) is non-positive once the number sold N_i reaches k_i . Therefore, Algorithm 1 never tries to offer an item i which has stocked out.

⁴ α can be solved to equal $1 - W(Re^{R-1})/R$, where W is the inverse of the function $f(x) = xe^x$, and $R = \ln(r^{(m)}/r^{(1)})$ —see Appendix E.1.

Algorithm 1 MULTI-PRICE BALANCE

1: Initialize $\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)}, \tilde{\Phi}_i$ randomly and independently for each $i \in [n]$
2: $N_i \leftarrow 0$ for all $i \in [n]$ (N_i tracks the total number of copies of item i sold, at any price)
3: **for** $t = 1, 2, \dots$ **do**
4: Compute

$$\max_{i \in [n], j \in [m_i]} p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_i}{k_i})) \quad (16)$$

5: **if** the value of (16) is strictly positive **then**
6: Offer any item i_t^* and price j_t^* maximizing (16) to customer t
7: **if** customer t accepts (occurring with probability $p_{t,i_t^*}^{(j_t^*)}$) **then**
8: $Z_t \leftarrow \tilde{\Phi}_{i_t^*}(\tilde{L}_{i_t^*}^{(j_t^*)}) - \tilde{\Phi}_{i_t^*}(N_{i_t^*}/k_{i_t^*})$ (this is the pseudorevenue earned)
9: $N_{i_t^*} \leftarrow N_{i_t^*} + 1$
10: **end if**
11: **end if**
12: **end for**

THEOREM 4. *Suppose in Line 1 of Algorithm 1, for each $i \in [n]$, the segment borders $\tilde{L}_i^{(1)}, \dots, \tilde{L}_i^{(m_i)}$ and value function $\tilde{\Phi}_i$ are randomly initialized in a way such that*

$$k_i(\tilde{\Phi}_i(\frac{N+1}{k_i}) - \tilde{\Phi}_i(\frac{N}{k_i})) + \tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N}{k_i}) \leq \frac{r_i^{(j)}}{c}, \quad j \in [m_i], N \in \{0, \dots, \tilde{L}_i^{(j)} k_i - 1\}; \quad (17)$$

$$\mathbb{E}[\tilde{\Phi}_i(\tilde{L}_i^{(j)})] \geq r_i^{(j)}, \quad j \in [m_i]. \quad (18)$$

Then Algorithm 1 achieves a competitive ratio of c .

Theorem 4 identifies conditions which, when satisfied by the randomized procedure for each i , yields a competitive ratio of c . Note that (17) needs to hold for every potential initialization of $\tilde{\Phi}_i$, while (18) only needs to hold in expectation over the initializations. We prove Theorem 4 in Appendix B, but outline its proof here and provide some intuition.

First, we take the dual of the LP (5):

$$\min \sum_{i=1}^n k_i y_i + \sum_{t=1}^T z_t \quad (19a)$$

$$p_{t,i}^{(j)} y_i + z_t \geq p_{t,i}^{(j)} r_i^{(j)} \quad t \in [T], i \in [n], j \in [m_i] \quad (19b)$$

$$y_i, z_t \geq 0 \quad i \in [n], t \in [T] \quad (19c)$$

By weak duality, $\text{OPT}(\mathcal{S}, \mathcal{A})$ is upper-bounded by the objective value of any feasible dual solution.

During the (random) execution of Algorithm 1, it maintains a dual variable $y_i = \tilde{\Phi}_i(\frac{N_i}{k_i})$ for each i . At each time t , only if a sale is realized, does the algorithm set z_t to a non-zero value Z_t (Line 8) and increment the y_i -variables by incrementing N_{i^*} (Line 9). We prove three claims:

1. During each time $t \in [T]$, the gain in the dual objective is at most some multiple $\frac{1}{c}$ of the revenue earned by the algorithm;
2. During each time $t \in [T]$, the conditional expectation of Z_t over the random purchase decision of customer t , combined with the current value of y_i , make the LHS of (19b) at least $p_{t,i}^{(j)} \cdot \tilde{\Phi}_i(\tilde{L}_i^{(j)})$, for all $i \in [n]$ and $j \in [m_i]$;
3. The expectation of $\tilde{\Phi}_i(\tilde{L}_i^{(j)})$, over the random segment borders and value function initially chosen by the algorithm, is at least $r_i^{(j)}$, for all $i \in [n]$ and $j \in [m_i]$.

Claim 1 follows from condition (17), while Claim 3 follows from condition (18). Claims 2 and 3 can be combined to show that the dual variables y_i and z_t maintained by the algorithm are feasible, after taking an expectation over all sample paths.

We explain the intuition behind our idea of a *random* value function, and the resulting analysis. Even for a single item, with a small starting inventory and a large ratio r from its highest to lowest price, in order to achieve a constant competitive ratio which does not scale with r , one must use *random* booking limits (Ball and Queyranne 2009). With multiple items, our equivalent is to have the configuration of segment borders $\tilde{L}_i^{(0)}, \dots, \tilde{L}_i^{(m_i)}$ be random, and define an arbitrary value function $\tilde{\Phi}_i$ corresponding to each one. In order to “average” over these configurations in the analysis, we relax dual feasibility to only hold in expectation. The idea of feasibility in expectation has been previously seen, but in different contexts: in Devanur et al. (2013), over a random seed, and in Golrezaei et al. (2014), over a random purchase decision (similar to our Claim 2).

3.1. Optimizing the Randomized Procedures

Theorem 4 reduces the problem of deriving a competitive algorithm to that of finding a randomized procedure for initializing $\tilde{\Phi}_1, \dots, \tilde{\Phi}_n$ satisfying (17)–(18). We can consider this problem separately for each i , based on $r_i^{(1)}, \dots, r_i^{(m_i)}$ and k_i , and omit the subscript i .

A randomized procedure consists of a distribution over the all of the configurations satisfying (14), and for each configuration, values for $\tilde{\Phi}(\frac{1}{k}), \tilde{\Phi}(\frac{2}{k}), \dots, \tilde{\Phi}(1)$ satisfying (15). We would like to find a randomized procedure which satisfies (17)–(18) with a maximal value of c . While this optimization problem is intractable in general, we can use the intuition behind the definitions of $L^{(0)}, \dots, L^{(m)}$ and $\Phi_{\mathcal{P}}$ from Subsection 2.1 to specify a near-optimal randomized procedure.

DEFINITION 3. Define the following randomized procedure for initializing $\tilde{\Phi}$:

1. Draw a random seed W uniformly from $[0, 1]$;
2. For each j , set $\tilde{L}^{(j)} = \frac{\lfloor L^{(j)}k \rfloor + 1}{k}$ if $W < L^{(j)}k - \lfloor L^{(j)}k \rfloor$, and $\tilde{L}^{(j)} = \frac{\lfloor L^{(j)}k \rfloor}{k}$ otherwise;
3. For $q \in \{0, \frac{1}{k}, \dots, 1\}$, let $\tilde{\ell}(q)$ be the unique $j \in [m]$ such that $\tilde{L}^{(j-1)} \leq q < \tilde{L}^{(j)}$ (note that $\tilde{\ell}(\tilde{L}^{(j)}) = j + 1$ for $j = 0, \dots, m - 1$; we define $\tilde{\ell}(\tilde{L}^{(m)})$ to be m).

The value function $\tilde{\Phi}$ is then defined over $q \in \{0, \frac{1}{k}, \dots, 1\}$ by

$$\tilde{\Phi}(q) = \sum_{j=1}^{\tilde{\ell}(q)-1} (r^{(j)} - r^{(j-1)}) \frac{\exp(\tilde{L}^{(j)} - \tilde{L}^{(j-1)}) - 1}{\exp(\alpha^{(j)}) - 1} + (r^{(\tilde{\ell}(q))} - r^{(\tilde{\ell}(q)-1)}) \frac{\exp(q - \tilde{L}^{(\tilde{\ell}(q)-1)}) - 1}{\exp(\alpha^{(\tilde{\ell}(q))}) - 1}. \quad (20)$$

$\tilde{\Phi}$ increases over the m (possibly empty) “segments” of its domain $\{0, \frac{1}{k}, \dots, 1\}$, which are “bordered” by $\tilde{L}^{(0)}, \dots, \tilde{L}^{(m)}$. (20) is similar to definition (9) for $\Phi_{\mathcal{P}}$, except the sum in (20) does not telescope, since $\tilde{L}^{(j)} - \tilde{L}^{(j-1)}$ equals $\alpha^{(j)}$ only in expectation.

Note that in Step 2 above, the random segment borders $\tilde{L}^{(0)}, \dots, \tilde{L}^{(m)}$ are rounded *comonotonically* (in a perfectly positively correlated fashion) using a single seed. This ensures that the borders are increasing as required in (14), as well as the following properties.

PROPOSITION 3. *The random values of $\tilde{L}^{(0)}, \dots, \tilde{L}^{(m)}$ from Definition 3 satisfy:*

$$\mathbb{E}[\tilde{L}^{(j)}] = L^{(j)}, \quad j = 0, \dots, m; \quad (21)$$

$$|(\tilde{L}^{(j)} - \tilde{L}^{(j')}) - (L^{(j)} - L^{(j')})| \leq \frac{1}{k}, \quad 1 \leq j' < j \leq m. \quad (22)$$

(22) is the key property derived from comonotonicity: although the rounding could move each $\tilde{L}^{(j)}$ by up to $1/k$ in either direction, the distance *between* two different $\tilde{L}^{(j)}, \tilde{L}^{(j')}$ never changes by more than $1/k$ from $L^{(j)} - L^{(j')}$. Proposition 3 is then used to prove our main result about the randomized procedure from Definition 3.

Algorithm 2 MULTI-PRICE RANKING in the Deterministic Case

```

1: Initialize  $W_i$  uniformly at random from  $[0, 1]$ , independently for each  $i \in [n]$ 
2:  $\text{available}_i \leftarrow \mathbf{true}$  for all  $i \in [n]$ 
3: for  $t = 1, 2, \dots$  do
4:   Compute
      
$$\max_{i \in [n], j \in [m_i]: \text{available}_i = \mathbf{true}} (r_i^{(j_{t,i})} - \Phi_{\mathcal{P}_i}(W_i)) \tag{23}$$

5:   if the value of (23) is strictly positive then
6:     Offer any item  $i_t^*$  maximizing (23) to customer  $t$ , at price  $j_{t,i_t^*}$ 
7:      $\text{available}_{i_t^*} \leftarrow \mathbf{false}$ 
8:   end if
9: end for

```

THEOREM 5. *The randomized procedure for initializing $\tilde{\Phi}$ from Definition 3 satisfies (17)–(18) with $c = \frac{1-e^{-\alpha(1)}}{(1+k)(e^{1/k}-1)}$. Furthermore, if $m = 1$, then the value of c can be improved to $\frac{1-e^{-\alpha(1)}}{(1+k)(1-e^{-1/k})}$.*

Theorem 5 is proven in Appendix B. It, in conjunction with Theorem 4, establishes bounds (i) and (iii) from our main result for MULTI-PRICE BALANCE, Theorem 1. In Appendix B, we state the complete proof of Theorem 1, including bound (ii), which involves explicitly formulating the optimization problem over randomized procedures and solving it when $k = 1$.

4. MULTI-PRICE RANKING and the Proof of Theorem 2

In Subsection 2.2, we sketched MULTI-PRICE RANKING for our general problem. In Algorithm 2, we formalize it specifically for the deterministic case, which is the case analyzed in Theorem 2. Note that we have assumed, without loss of generality, that $k_i = 1$ for each item i .

Our analysis extends the framework of Devanur et al. (2013) to incorporate multiple prices. It uses the dual LP defined in (19), where every $p_{t,i}^{(j)}$ is 0 or 1.

If Algorithm 2 assigns item i to customer t (charging price $j_{t,i}$), then we set dual variables $Z_t = r_i^{(j_{t,i})} - \Phi_{\mathcal{P}_i}(W_i)$ and $Y_i = \Phi'_{\mathcal{P}_i}(W_i)$, where $\Phi_{\mathcal{P}_i}$ is the fixed function defined in Subsection 2.1 (we ignore the measure-zero set where $\Phi'_{\mathcal{P}_i}$ is undefined). All dual variables not set during a time period are defined to be zero. The following lemmas are proven in Appendix C:

LEMMA 2. *If Algorithm 2 assigns item i to customer t , then $(1 - e^{-\alpha_i^{(1)}})(Y_i + Z_t) \leq r_i^{(j_{t,i})}$ w.p.1.*

LEMMA 3. *Setting $y_i = \mathbb{E}[Y_i]$, $z_t = \mathbb{E}[Z_t]$ for all i, t forms a feasible solution to the dual LP (19).*

The proof of Theorem 2 is then easy given these lemmas:

Proof of Theorem 2. Lemma 3 implies $\text{OPT}(\mathcal{S}, \mathcal{A}) \leq \sum_{i=1}^n \mathbb{E}[Y_i] + \sum_{t=1}^T \mathbb{E}[Z_t]$, via weak duality.

However, by Lemma 2, the revenue earned by Algorithm 2 is at least $\min_{i \in [n]} \{1 - e^{-\alpha_i^{(1)}}\} \cdot (\sum_{i=1}^n Y_i + \sum_{t=1}^T Z_t)$, with probability 1. Thus, $\mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})] \geq (1 - \exp(-\min_{i \in [n]} \alpha_i^{(1)})) \cdot \text{OPT}(\mathcal{S}, \mathcal{A})$. \square

5. Randomized Counterexample and the Proof of Theorem 3

We first formalize the setup and randomized arrival sequence described in Subsection 2.3.

There are $n \in \mathbb{N}$ items, indexed by i , which all have $m_i = m$, $r_i^{(j)} = r^{(j)}$ for all j , and $k_i = k$ for some $k \in \mathbb{N}$. We think of n as going to ∞ , while k is arbitrary. Throughout this example, we often express quantities as portions τ of n . We abuse notation and write τn to refer to an integer, even if τ is irrational, since the error from rounding τn to the nearest integer is negligible as $n \rightarrow \infty$.

The arrival sequence is randomized following the classical construction of Karp et al. (1990). There are $T = nk$ customers, split into n “groups” of k identical customers each. Uniformly draw a random permutation $\pi = (\pi_1, \dots, \pi_n)$ of $(1, \dots, n)$ from the $n!$ possibilities. For $i \in [n]$, all k customers in group i would deterministically buy any item in $\{\pi_i, \dots, \pi_n\}$. Our construction differs from existing ones in that the n groups of customers are further split into m “phases”. Let β_1, \dots, β_m be positive numbers summing to 1, corresponding to the fraction of groups in each phase, whose values we specify later. For all $j \in [m]$, the customers in groups $(\beta_1 + \dots + \beta_{j-1})n + 1, \dots, (\beta_1 + \dots + \beta_j)n$ are willing to pay $r^{(j)}$ for any of the items in their interest set.

DEFINITION 4. Define the following shorthand notation for all $j = 1, \dots, m + 1$:

- $A_j := \sum_{\ell=j}^m \alpha^{(\ell)}$ (note that $A_1 = 1$ and $A_{m+1} = 0$);
- $B_j := \sum_{\ell=j}^m \beta_\ell$ (note that $B_1 = 1$ and $B_{m+1} = 0$).

PROPOSITION 4. Given $m \in \mathbb{N}$, $0 < r^{(1)} < \dots < r^{(m)}$, and $\alpha^{(1)}, \dots, \alpha^{(m)}$ as defined in Proposition 1, there exists a unique solution to the following system of equations in variables B_2, \dots, B_m :

$$B_m r^{(m)} e^{-\alpha^{(m)}} = \dots = B_2 r^{(2)} e^{-\alpha^{(2)}} = r^{(1)} e^{-\alpha^{(1)}}, \quad (24)$$

with $0 < B_m < \dots < B_2 < B_1 = 1$.

We define B_2, \dots, B_m according to Proposition 4. This implies definitions for β_1, \dots, β_m , which are strictly positive and sum to 1.

Now, regardless of the permutation π , the optimal algorithm allocates the k copies of item π_i to the customers in group i , for each $i \in [n]$, successfully serving all $T = nk$ customers and earning revenue $\sum_{j=1}^m r^{(j)}(\beta_j n)k$. This is also the optimal objective value of the LP (5). Therefore, regardless of the realized arrival sequence \mathcal{A} , $\text{OPT}(\mathcal{S}, \mathcal{A}) = \sum_{j=1}^m r^{(j)}(\beta_j n)k$ which we can rewrite as

$$\sum_{j=1}^m (r^{(j)} - r^{(j-1)})B_j nk. \quad (25)$$

5.1. Upper Bound on Performance of Online Algorithms

LEMMA 4. *The expected revenue of an online algorithm with this randomized \mathcal{A} is upper-bounded by the maximum value of*

$$\sum_{j=1}^m r^{(j)} B_j n (1 - e^{-\lambda_j}) k \quad (26)$$

subject to $0 \leq \lambda_j \leq \ln \frac{B_j}{B_{j+1}}$ for $j \in [m-1]$, $0 \leq \lambda_m$, and $\sum_{j=1}^m \lambda_j \leq 1$.

Lemma 4 drastically simplifies the analysis of the online algorithm, because it restricts to algorithms which are *indifferent* to the realized permutation π , allowing for a deterministic analysis. However, our analysis differs from existing ones (e.g. (Golrezaei et al. 2014, Lem. 6)) in that despite the item symmetry, the online algorithm has a decision—how many customers in each phase to serve, as opposed to reserving inventory for customers in future phases.

This is controlled by the λ -variables, where λ_j denotes the expected fraction of item π_n 's inventory sold to phase- j customers. The expected number of groups served during phase j is then at most $B_j n (1 - e^{-\lambda_j})$, resulting in the upper bound (26). Constraint $\lambda_j \leq \ln \frac{B_j}{B_{j+1}}$ comes from the fact that $B_j n (1 - e^{-\lambda_j})$ must not exceed the total number of groups in phase j , $\beta_j n$.

LEMMA 5. *Let $j \in [m]$ and $\tau \in [0, 1]$. The maximum value of*

$$\sum_{\ell=j}^m r^{(\ell)} B_\ell n (1 - e^{-\lambda_\ell}) k \quad (27)$$

subject to $\lambda_\ell \geq 0$ for all $\ell = j, \dots, m$ as well as $\sum_{\ell=j}^m \lambda_\ell \leq \tau$ is

$$nk \sum_{\ell=j}^m r^{(\ell)} B_\ell \left(1 - \exp \left(-\alpha^{(\ell)} + \frac{A_j - \tau}{m - j + 1} \right) \right). \quad (28)$$

Lemma 5 establishes the optimal objective value of the optimization problem from Lemma 4. The upper bound of $\ln \frac{B_j}{B_{j+1}}$ on λ_j for $j \in [m-1]$ turns out to not be binding. With both lemmas, the proof of Theorem 3 is easy.

Proof of Theorem 3. The value of (28) with $j = 1$ and $\tau = 1$ is

$$nk \sum_{\ell=1}^m r^{(\ell)} B_{\ell} (1 - e^{-\alpha^{(\ell)}}) = (1 - e^{-\alpha^{(1)}}) \sum_{\ell=1}^m (r^{(\ell)} - r^{(\ell-1)}) B_{\ell} nk, \quad (29)$$

where we have used (7) to derive the equality. Combining Lemmas 4–5, we get that the RHS of (29) is an upper bound on $\mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})]$, for any online algorithm. Meanwhile, $\text{OPT}(\mathcal{S}, \mathcal{A})$ is equal to (25) regardless of \mathcal{A} , which is exactly the RHS of (29) divided by $(1 - e^{-\alpha^{(1)}})$. We have established that $\mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})] \leq (1 - e^{-\alpha^{(1)}}) \mathbb{E}[\text{OPT}_{\mathcal{A}}(\mathcal{S}, \mathcal{A})]$, completing the proof of the theorem. \square

REMARK 1. Suppose that $k \rightarrow \infty$. It can be seen that our algorithm (either MULTI-PRICE BALANCE or MULTI-PRICE RANKING, which behave identically when $k \rightarrow \infty$ —see Aggarwal et al. (2011)), with booking limits $\alpha^{(1)}, \dots, \alpha^{(m)}$, is the unique optimal algorithm given this distribution over arrival sequences. Indeed, the proof of Lemma 4 shows that given $\lambda_1, \dots, \lambda_m$, the dominant strategy for the online algorithm is to deplete the inventories of items evenly (which is possible since $k \rightarrow \infty$), in which case upper bound (26) is attained. The proof of Lemma 5 shows that the unique optimal values for $\lambda_1, \dots, \lambda_m$ are $\alpha^{(1)}, \dots, \alpha^{(m)}$. It only remains to show that $\lambda_j = \alpha^{(j)}$ is feasible, namely $\alpha^{(j)} \leq \ln \frac{B_j}{B_{j+1}}$ for $j < m$. Applying (24), this is equivalent to showing $e^{-\alpha^{(j)}} \geq \frac{r^{(j)} e^{-\alpha^{(j)}}}{r^{(j+1)} e^{-\alpha^{(j+1)}}}$, or $e^{-\alpha^{(j+1)}} \geq \frac{r^{(j)}}{r^{(j+1)}}$, which follows from (7) since $1 - e^{-\alpha^{(1)}} \leq 1$.

6. Extending our Techniques

We explain how our techniques can be extended to allow for fractional inventory consumption like in the Adwords problem (Mehta et al. 2007), or offering multiple items like in the online assortment problem (Golrezaei et al. 2014).

Consider the following modification of our problem from Section 2: when customer t is offered item i at price j , she *deterministically* pays $p_{t,i}^{(j)} r_i^{(j)}$ and consumes a fractional amount $p_{t,i}^{(j)} \leq 1$ of item i 's inventory, instead of paying $r_i^{(j)}$ and consuming 1 unit with probability $p_{t,i}^{(j)}$. We assume

that $\min_i k_i \rightarrow \infty$. This generalizes the Adwords problem under the small bids assumption, by allowing each budget i to be depleted at m_i different rates $r_i^{(1)}, \dots, r_i^{(m_i)}$.

For this problem, we use MULTI-PRICE BALANCE, except since we are taking $\min_i k_i \rightarrow \infty$, we can deterministically set each $\tilde{\Phi}_i = \Phi_{\mathcal{P}_i}$. The three claims used to establish Theorem 4 are simpler: Claim 2 now holds deterministically instead of requiring a conditional expectation over Z_t , while Claim 3 also holds deterministically since $\tilde{\Phi}_i$ is always $\Phi_{\mathcal{P}_i}$. In Theorem 5, condition (17) is now only satisfied under an additional error term ε , since N is no longer a discrete integer. Nonetheless, the rounding error ε approaches 0 as $k_i \rightarrow \infty$, so the optimal competitive ratio is still achieved.

For online assortment, we use the term *product* to refer to an (item, price)-combination (i, j) . Consider the following modification of our problem from Section 2: upon the arrival of customer t , for any subset (assortment) S of products and $(i, j) \in S$, we are given $p_{t,i}^{(j)}(S)$, the probability that customer t would pick product (i, j) when offered the choice from S . After being given these probabilities, we must offer an assortment S to customer t . This generalizes the original online assortment problem of Golrezaei et al. (2014), by allowing each item be offered at different prices. We note that the assortment offered S can be constrained to lie in an arbitrary downward-closed family \mathcal{F} of subsets of $\{(i, j) : i \in [n], j \in [m_i]\}$; for example, we could disallow assortments where an item is simultaneously offered at multiple prices. The execution of an algorithm can be encapsulated by the following modification of the LP (5):

$$\max \sum_{t=1}^T \sum_{S \in \mathcal{F}} x_t(S) \sum_{(i,j) \in S} r_i^{(j)} p_{t,i}^{(j)}(S) \quad (30a)$$

$$\sum_{t=1}^T \sum_{S \in \mathcal{F}} x_t(S) \sum_{j:(i,j) \in S} p_{t,i}^{(j)}(S) \leq k_i \quad i \in [n] \quad (30b)$$

$$\sum_{S \in \mathcal{F}} x_t(S) = 1 \quad t \in [T] \quad (30c)$$

$$x_t(S) \geq 0 \quad t \in [T], S \in \mathcal{F} \quad (30d)$$

MULTI-PRICE BALANCE can be directly applied to this problem, with the change that it offers the *assortment* $S \in \mathcal{F}$ maximizing expected pseudorevenue, $\sum_{(i,j) \in S} p_{t,i}^{(j)}(S) (\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(w_i))$, to

each customer t . We assume the existence of an oracle for solving this single-shot assortment optimization problem, which admits an efficient algorithm under many commonly-used choice models (see Cheung and Simchi-Levi (2016) for a summary). In the analysis, dual constraints (19b) now require $z_t \geq \sum_{(i,j) \in S} p_{t,i}^{(j)}(S)(r_i^{(j)} - y_i)$ for all t and S , which is still implied by the conditions of Theorem 4 as long as the choice probabilities for customers satisfy a mild *substitutability* assumption (see Golrezaei et al. (2014) for details).

7. Simulations on Hotel Data Set of Bodea et al. (2009)

We test our algorithms on the publicly-accessible hotel data set collected by Bodea et al. (2009). Based on the data, we consider a multi-price online assortment problem, as defined in Section 6.

7.1. Experimental Setup

We consider Hotel 1 from the data set, which has more transactions than the other four hotels. For each transaction, we use *booking* to refer to the date the transaction occurred, and *occupancy* to refer to the dates the customer will stay in the hotel. We consider occupancies spanning the 5-week period from Sunday, March 11th, 2007 to Sunday, April 15th, 2007. Although the data contains occupancies for a couple of weeks outside this range, such transactions are sparse.

We merge the different rooms into 4 categories: King rooms, Queen rooms, Suites, and Two-double rooms. Rooms under the same category draw from the same inventory. We merge the different fare classes into two: discounted advance-purchase fares and regular rack rates. We use *product* to refer to any of the 8 combinations formed by the 4 room categories and 2 fares.

We estimate a Multinomial Logit (MNL) choice model on these 8 products, for each of 8 customer types. The customer types are based on the booking channel, party size, and VIP status (if any) associated with a transaction. These types capture preference heterogeneity (for example, party sizes greater than 1 tend to prefer Suites and Two-double rooms). The details of our choice estimation are deferred to Appendix F.

We should point out that more sophisticated segmentation and estimation techniques have been employed on this data set (van Ryzin and Vulcano 2014, Newman et al. 2014). Nonetheless, MNL

has been reported to perform relatively well (van Ryzin and Vulcano 2014, sec. 5.2). The MNL choice model is convenient for our purposes because under it, both the assortment optimization problem, as well as the choice-based LP (30) with exponentially many variables, can be solved efficiently (Talluri and Van Ryzin 2004, Liu and Van Ryzin 2008, Cheung and Simchi-Levi 2016).

We treat each occupancy date as a separate instance of the problem, for which we define a sequence of arrivals, with one arrival for each transaction which occupies that date. The choice probabilities for each arrival are determined by the customer type associated with the corresponding transaction.⁵ The number of days in advance of occupancy that each arrival occurred is also recorded, but this information is only relevant for algorithms which attempt to forecast the remaining number of arrivals based on the remaining length of time.

Before we proceed, we discuss the limitations of our analysis and the data set:

1. In the data set, 55% of the transactions occupy multiple, *consecutive* days. However, we treat such a transaction as a separate arrival in the instances for each of those occupancy dates. While this is a simplifying assumption, the focus of our paper is on the basic allocation problem without complementarity effects across consecutive days, and our goal in using the data set is to extract an arrival pattern over time.
2. It is not possible to deduce from the data the fixed capacity for each category of room. To compensate, we consider a wide range of starting capacities in our tests.
3. Estimating the number of customers who do not make a purchase is a standard challenge in choice modeling, which is exacerbated in this data set by the fact that the arrivals are rather non-stationary. We test various assumptions on the weight of the no-purchase option in the MNL model for each customer type. In general, we assume that this weight is large, which causes the revenue-maximizing assortments to be large, allowing for tension between offering large assortments which maximize immediate revenue, and offering small assortments which regulate inventory consumption (details in Appendix F).

⁵ The choice realized in that transaction was used for choice model estimation, but is not used in defining the arrival.

Table 1 Details on Room Categories and Fares

Room Category	Low Fare	High Fare	Fraction of Rooms
King	\$307	\$361	52%
Queen	\$304	\$361	15%
Suite	\$384	\$496	13%
Two Double	\$306	\$342	20%

7.2. Instance Definition

A test instance corresponds to a specific occupancy date, which has a finite inventory of each room category. Each customer interested in that occupancy date arrives in sequence, after which her characteristics (channel, party size, VIP status) are revealed. The problem is to show a *personalized* assortment of (room, fare)-options to each customer. The instances we test are described below.

- Arrival sequence: 35 possibilities, one for each day in the 5-week occupancy period. We multiply the arrivals by 10 (i.e. instead of a type-1 customer followed by a type-2 customer, we have 10 type-1 customers followed by 10 type-2 customers), being interested in the high-inventory regime. After multiplication, the average number of arrivals per day is 1340, peaking on Sundays and Mondays, although the number and breakdown of customers varies by day.
- Number of products: 8 (room, fare)-combinations, identical for all instances.
- Prices of products: displayed in Table 1, identical for all instances. These prices were determined by taking the average price of that (room, fare)-combination over all transactions.
- Starting inventories: 3 possibilities, where we set the starting inventories to yield a desired *loading factor*. The loading factor is defined by the average (over all 35 days) number of customers per unit of starting inventory, and we use the same loading factors (1.4, 1.6, 1.8) as Golrezaei et al. (2014). For a fixed loading factor, all 35 instances have the same starting inventories. The fraction of total starting inventory corresponding to each room type is based on the relative frequency with which that type is booked over all transactions (see Table 1).

We test additional synthetic instances, where we increase the high fares and consider a greater range of loading factors, in Subsection 7.5.

7.3. Algorithms Compared

We compare the performances of 10 algorithms on each instance.

First we describe the forecast-independent algorithms we test.

1. **Myopic**: offer each customer the assortment maximizing immediate expected revenue, from the items that have not stocked out.
2. **Conservative**: only offer items at their maximum prices.⁶
3. **GNR**: offer to each customer t the assortment S maximizing

$$\sum_{(i,j) \in S} p_{t,i}^{(j)}(S) \cdot r_i^{(j)} \Psi(w_i),$$

where w_i is the fraction of item i sold and Ψ is the inventory balancing function from Golrezaei et al. (2014). This would represent the algorithm of Golrezaei et al. (2014) applied to the multi-price setting (see Section 1.2).

4. **Our Algorithm**: offer to each customer t the assortment S maximizing

$$\sum_{(i,j) \in S} p_{t,i}^{(j)}(S) \cdot (r_i^{(j)} - \Phi_{\mathcal{P}_i}(w_i)), \quad (31)$$

where w_i is the fraction of item i sold. This is essentially MULTI-PRICE BALANCE, except we have used the fixed value function $\Phi_{\mathcal{P}_i}$ instead of the random value function $\tilde{\Phi}_i$ to define the bid price of each item i , which is a simplifying approximation for the high-inventory regime.

The Myopic and Conservative algorithms represent two extremes, where the former extracts the maximum in expectation from every customer and is optimal as the loading factor approaches 0, while the latter extracts the maximum from every unit of inventory and is optimal as the loading factor approaches ∞ . In-between these extremes, our algorithm attempts to balance revenue-per-customer and revenue-per-item as it selects items and prices to put in the assortment.

Next we describe the forecasting-based algorithms we test. These algorithms all estimate the number of each type of customer yet to arrive, and then incorporate this information into the

⁶ This algorithm selects between the items (at their high prices) using the algorithm of Golrezaei et al. (2014).

LP (30) to set bid prices. They differ in how they perform the forecasting, and how frequently they update the bid prices by re-solving the LP. Further details about these algorithms, as well as discussion of alternative algorithms, are deferred to Appendix F.1.

5. **One-shot LP**: solve the LP only once, at the start, using the average number of customers of each type to appear on a given day.
6. **LP Resolving**: re-solve the LP every 100 arrivals, using updated forecasts and inventory counts. During each re-solve, the estimated number of remaining customers is updated, taking into account the length of time remaining until occupancy, and the number of customers that have arrived. The estimated type breakdown is fixed, based on the aggregate distribution.
7. **LP Learning**: same as LP Resolving, except the estimated type breakdown is also updated, based on the empirical distribution observed thus far.
8. **LP Clairvoyant**: same as LP Resolving, but given the true number of customers of each type remaining.

Finally, we describe the hybrid algorithms we test. These algorithms combine a forecasting algorithm with “Our Algorithm” as described above, based on a parameter $\gamma > 1$. For each customer t , the hybrid algorithm considers the expected pseudorevenue (as defined in (31)) of the assortment S^{fcst} suggested by the forecasting algorithm. If this is at least $\frac{1}{\gamma}$ of the maximum value of (31) over all assortments S , then the hybrid algorithm offers S^{fcst} . Otherwise, the hybrid algorithm offers the assortment suggested by our algorithm, which maximizes (31) by definition.

9. **Resolve-1.5**: hybrid algorithm based on LP Resolving and parameter $\gamma = 1.5$.
10. **Learn-1.5**: hybrid algorithm based on LP Learning and parameter $\gamma = 1.5$.

We selected $\gamma = 1.5$ above by taking the better of the two values 1.5, 2.0 tested in Golrezaei et al. (2014). We did not search over $\gamma > 1$ for the best γ , as the reported performance of such a hybrid algorithm would be greatly inflated, since γ would be chosen after seeing the performance.

7.4. Results

On every instance, we express the performance of each algorithm as a percentage of the LP upper bound. That is, we take the expected revenue of the algorithm (approximated over 10 runs), and

Table 2 The percentages of optimum achieved by different algorithms. The 3 highest percentages in each row are **bolded**. The 3 lowest standard deviations in each row are *italicized*.

Loading Factor		Forecast-independent				Forecast-dependent				Hybrid	
		Myopic	Conservative	GNR	OurAlg	One-shot	Resolve	Learn	Clairvoyant	Resolve-1.5	Learn-1.5
1.4	Mean	0.974	0.940	0.973	0.976	0.973	0.962	0.958	0.991	0.977	0.977
	Stdev	<i>0.023</i>	<i>0.034</i>	<i>0.020</i>	<i>0.013</i>	<i>0.016</i>	<i>0.039</i>	<i>0.041</i>	<i>0.008</i>	<i>0.018</i>	<i>0.020</i>
1.6	Mean	0.965	0.960	0.964	0.971	0.964	0.961	0.963	0.990	0.977	0.978
	Stdev	<i>0.025</i>	<i>0.036</i>	<i>0.020</i>	<i>0.014</i>	<i>0.021</i>	<i>0.031</i>	<i>0.030</i>	<i>0.008</i>	<i>0.008</i>	<i>0.010</i>
1.8	Mean	0.957	0.972	0.960	0.968	0.808	0.962	0.968	0.990	0.977	0.977
	Stdev	<i>0.020</i>	<i>0.036</i>	<i>0.017</i>	<i>0.012</i>	<i>0.100</i>	<i>0.029</i>	<i>0.023</i>	<i>0.009</i>	<i>0.008</i>	<i>0.007</i>


divide it by the optimal objective value of the LP (30) with the true arrival sequence. In Table 2, we report the mean and standard deviation of each algorithm’s percentages over the 35 arrival sequences, for each loading factor.

In general, our MULTI-PRICE BALANCE algorithm is the most profitable and consistent among the forecast-independent algorithms, while the forecast-dependent algorithms have much greater fluctuation in their performance for different occupancy days, depending on how accurate their forecasts were for that day. LP Learning is slightly better than the others, but is most prone to overfitting in its forecasts. We note that although the forecast-independent algorithms do not make use of information about the remaining time horizon (which can be used to estimate the remaining number of customers), they perform comparably well to the forecasting algorithms. Nonetheless, by combining the forecasting algorithms with MULTI-PRICE BALANCE, the hybrid algorithms are able to correct for forecast overconfidence and achieve the best performance overall (aside from the Clairvoyant algorithm, which has a perfect forecast of the future). We find that although the hybrid algorithm only changes a small fraction ($\approx 5\%$) of the forecasting algorithm’s decisions, this drastically improves the profitability and consistency.

7.5. Results under Greater Fare Differentiation

The instances tested in Subsection 7.4 were “easy” in that there was not so much difference between selling rooms at their low or high fares. In this subsection, we synthetically modify the higher fare for each room category to be *twice* its lower fare. We also increase the utility of the no-purchase

Figure 3 Algorithm performances in the setting with greater fare differentiation. The lines corresponding to the two hybrid algorithms, which perform the best overall, have been **bolded**.



option in the MNL model for each customer type (see Appendix F), to maintain the tension between low fares which maximize expected revenue, and high fares which limit inventory consumption.

Furthermore, we test a complete range of loading factors, including both the extreme where the Myopic algorithm is optimal, and the extreme where the Conservative algorithm is optimal. In Figure 3, we plot the average percentages of optimum attained by each algorithm over the 35 arrival sequences, for each loading factor.

The conclusion again is that our two hybrid algorithms, which use forecasts but continuously reference our forecast-independent value functions, are the most profitable and consistent. Note that our MULTI-PRICE BALANCE algorithm comes third, and performs significantly better than the

inventory-balancing algorithm of Golrezaei et al. (2014), which is similar to the Myopic algorithm as it does not take the two different prices of the same item into account.

However, it is important to note that overall, our methodology is only relevant in scenarios where the loading factor is in-between the extremes where there is a non-trivial tradeoff. If the loading factor is very low, and the hotel does not even get close to full on any day, then it would be best to always use the Myopic algorithm; similarly, if the loading factor is very high, and the hotel gets full every day, then it would be best to always use the Conservative algorithm. Nonetheless, we argue that most hotels do lie in-between the extremes, where it is sometimes full and sometimes empty depending on sudden local events. Otherwise, the hotel either over-built or under-built in an higher-order decision.

8. Conclusion

Competitive analysis is a well-established methodology in sequential decision-making problems, providing a baseline decision in the absence of a reliable forecast of the future. Previously, online resource allocation algorithms based on competitive analysis have assumed that each resource can only be converted to reward at a fixed rate. We extend these results and derive algorithms which jointly consider the tradeoffs between different resources and different reward rates. This broadly expands the applicability of competitive analysis in areas such as online matching, online advertising, personalized e-commerce, and appointment scheduling.

Acknowledgments

The authors would like to thank Rong Jin of Alibaba for pointing out a detailed technical error in an earlier version of the appendix. The authors would also like to thank Ozan Candogan and James Orlin for asking questions which led to the simpler bound presented in Corollary 2. A preliminary version of this article appeared in the *20th ACM Conference on Economics and Computation (EC)*, whose anonymous reviewers helped clarify the manuscript in several places.

References

- Aggarwal G, Goel G, Karande C, Mehta A (2011) Online vertex-weighted bipartite matching and single-bid budgeted allocations. *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, 1253–1264 (Society for Industrial and Applied Mathematics).
- Araman VF, Caldentey R (2011) Revenue management with incomplete demand information. *Wiley Encyclopedia of Operations Research and Management Science* .
- Badanidiyuru A, Kleinberg R, Slivkins A (2013) Bandits with knapsacks. *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, 207–216 (IEEE).
- Ball MO, Queyranne M (2009) Toward robust revenue management: Competitive analysis of online booking. *Operations Research* 57(4):950–963.
- Bodea T, Ferguson M, Garrow L (2009) Data set—choice-based revenue management: Data from a major hotel chain. *Manufacturing & Service Operations Management* 11(2):356–361.
- Borodin A, El-Yaniv R (2005) *Online computation and competitive analysis* (cambridge university press).
- Buchbinder N, Jain K, Naor JS (2007) Online primal-dual algorithms for maximizing ad-auctions revenue. *European Symposium on Algorithms*, 253–264 (Springer).
- Chan CW, Farias VF (2009) Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Mathematics of Operations Research* 34(2):333–350.
- Chen X, Ma W, Simchi-Levi D, Xin L (2016) Dynamic recommendation at checkout under inventory constraint. *manuscript on SSRN* .
- Chen Y, Farias VF (2013) Simple policies for dynamic pricing with imperfect forecasts. *Operations Research* 61(3):612–624.
- Cheung WC, Simchi-Levi D (2016) Efficiency and performance guarantees for choice-based network revenue management problems with flexible products. *available on SSRN* .
- Ciocan DF, Farias V (2012) Model predictive control for dynamic resource allocation. *Mathematics of Operations Research* 37(3):501–525.
- Devanur NR, Jain K (2012) Online matching with concave returns. *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 137–144 (ACM).

- Devanur NR, Jain K, Kleinberg RD (2013) Randomized primal-dual analysis of ranking for online bipartite matching. *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, 101–107 (SIAM).
- Feldman J, Henzinger M, Korula N, Mirrokni V, Stein C (2010) Online stochastic packing applied to display ad allocation. *Algorithms-ESA 2010* 182–194.
- Feldman J, Korula N, Mirrokni V, Muthukrishnan S, Pál M (2009) Online ad assignment with free disposal. *International Workshop on Internet and Network Economics*, 374–385 (Springer).
- Ferreira KJ, Simchi-Levi D, Wang H (2016) Online network revenue management using thompson sampling. *manuscript on SSRN* .
- Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-time optimization of personalized assortments. *Management Science* 60(6):1532–1551.
- Jasin S, Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* 37(2):313–345.
- Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, 352–358 (ACM).
- Kesselheim T, Radke K, Tönnis A, Vöcking B (2013) An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. *European Symposium on Algorithms*, 589–600 (Springer).
- Lan Y, Gao H, Ball MO, Karaesmen I (2008) Revenue management with limited demand information. *Management Science* 54(9):1594–1609.
- Liu Q, Van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2):288–310.
- Manshadi VH, Gharan SO, Saberi A (2012) Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37(4):559–573.
- Mehta A (2013) Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8(4):265–368.

-
- Mehta A, Panigrahi D (2012) Online matching with stochastic rewards. *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, 728–737 (IEEE).
- Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5):22.
- Mehta A, Waggoner B, Zadimoghaddam M (2014) Online stochastic matching with unequal probabilities. *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1388–1404 (SIAM).
- Newman JP, Ferguson ME, Garrow LA, Jacobs TL (2014) Estimation of choice-based models using sales data from a single firm. *Manufacturing & Service Operations Management* 16(2):184–197.
- Reiman MI, Wang Q (2008) An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research* 33(2):257–282.
- Talluri K, Van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1):15–33.
- Talluri KT, Van Ryzin GJ (2006) *The theory and practice of revenue management*, volume 68 (Springer Science & Business Media).
- Van Ryzin G, McGill J (2000) Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science* 46(6):760–775.
- van Ryzin G, Vulcano G (2014) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* 61(2):281–300.
- Wang X, Truong V, Bank D (2015) Online advance admission scheduling for services, with customer preferences. *Working paper*.
- Yao ACC (1977) Probabilistic computations: Toward a unified measure of complexity. *Foundations of Computer Science, 1977., 18th Annual Symposium on*, 222–227 (IEEE).
- Zhang D, Cooper WL (2005) Revenue management for parallel flights with customer-choice behavior. *Operations Research* 53(3):415–431.
- Zhang H, Shi C, Qin C, Hua C (2016) Stochastic regret minimization for revenue management problems with nonstationary demands. *Naval Research Logistics (NRL)* 63(6):433–448.

Appendix A: Deferred Proofs from Section 2

Proof of Lemma 1. Fix any adaptive algorithm (which knows the arrival sequence, but not the realizations of the customers' purchase decisions, at the start) and consider its execution on setup \mathcal{S} with arrival sequence \mathcal{A} . Let $X_{t,i}^{(j)}$ be the indicator random variable (0 or 1) for the algorithm offering item i at price j to customer t , and $P_{t,i}^{(j)}$ be the indicator random variable for customer t accepting when item i is offered to her at price j . On a given run, the constraints $\sum_{t=1}^T \sum_{j=1}^{m_i} P_{t,i}^{(j)} X_{t,i}^{(j)} \leq k_i$ and $\sum_{i=1}^n \sum_{j=1}^{m_i} X_{t,i}^{(j)} \leq 1$ are satisfied. Therefore, they are still satisfied after taking an expectation over all runs, and furthermore we can use independence to show that $\mathbb{E}[P_{t,i}^{(j)} X_{t,i}^{(j)}] = \mathbb{E}[P_{t,i}^{(j)}] \cdot \mathbb{E}[X_{t,i}^{(j)}] = p_{t,i}^{(j)} x_{t,i}^{(j)}$. Therefore, the algorithm must satisfy constraints (5b) and (5c) of the LP. Since its revenue on a given run is $\sum_{t=1}^T \sum_{i=1}^n \sum_{j=1}^{m_i} P_{t,i}^{(j)} r_i^{(j)} X_{t,i}^{(j)}$, taking an expectation over it yields (5a), completing the proof. \square

Proof of Proposition 1. The statement for $\sigma^{(j)}, \dots, \sigma^{(j)}$ is immediate from the fact that the explicit value of $\sigma^{(j)}$ is $(1 - \frac{r^{(j-1)}}{r^{(j)}})(1 + \sum_{j'=2}^m (1 - \frac{r^{(j'-1)}}{r^{(j')}}))^{-1}$, for all $j \in [m]$. To prove the statement for $\alpha^{(1)}, \dots, \alpha^{(m)}$, we show that the solution to the system of n equations formed by (7) and $\alpha^{(1)} + \dots + \alpha^{(m)} = 1$ is unique and strictly positive.

Let $\gamma^{(j)} = e^{-\alpha^{(j)}}$ for all j . Then the constraint $\alpha^{(1)} + \dots + \alpha^{(m)} = 1$ can be rewritten as $\prod_{j=1}^m \gamma^{(j)} = \frac{1}{e}$. Furthermore, we derive from (7) that for all $j > 1$, $\gamma^{(j)} = (1 - \frac{r^{(j-1)}}{r^{(j)}})\gamma^{(1)} + \frac{r^{(j-1)}}{r^{(j)}}$. Therefore,

$$\gamma^{(1)} \cdot \prod_{j=2}^m \left(\left(1 - \frac{r^{(j-1)}}{r^{(j)}}\right)\gamma^{(1)} + \frac{r^{(j-1)}}{r^{(j)}} \right) = \frac{1}{e}. \quad (32)$$

Consider the LHS of (32) as a function of $\gamma^{(1)}$ on $[\frac{1}{e}, 1]$. This is a continuous, strictly increasing function which is at most $\frac{1}{e}$ when $\gamma^{(1)} = \frac{1}{e}$ and 1 when $\gamma^{(1)} = 1$. Therefore, there is a unique solution with $\gamma^{(1)} \in [\frac{1}{e}, 1)$, and the resulting value of $\alpha^{(1)}$ is positive. For $j > 1$, since $\gamma^{(j)}$ can also be written as $\gamma^{(1)} + \frac{r^{(j-1)}}{r^{(j)}}(1 - \gamma^{(1)})$, it can be seen that $\gamma^{(j)} \in [\frac{1}{e}, 1)$, hence the unique value for $\alpha^{(j)}$ is positive as well. \square

Proof of Proposition 2. For the first inequality in (11), observe that $f(x) = \frac{x}{1-e^{-x}}$ is a strictly increasing function on $[0, 1]$. Since $\sigma^{(1)} \in (0, 1)$, $\frac{\sigma^{(1)}}{1-e^{-\sigma^{(1)}}} < \frac{1}{1-\frac{1}{e}}$, which is the desired result.

For the second inequality in (11), we show $\alpha^{(1)} > \sigma^{(1)}$, by showing that for all $j = 2, \dots, m$, $\alpha^{(j)}$ is a smaller multiple of $\alpha^{(1)}$ than $\sigma^{(j)}$ is of $\sigma^{(1)}$. This suffices because both the fractions $\alpha^{(1)}, \dots, \alpha^{(m)}$ and $\sigma^{(1)}, \dots, \sigma^{(m)}$ must sum to 1. For a given j , we must establish that $\frac{\alpha^{(j)}}{\alpha^{(1)}} < \frac{\sigma^{(j)}}{\sigma^{(1)}}$. By definition, $\frac{\sigma^{(j)}}{\sigma^{(1)}} = 1 - \frac{r^{(j-1)}}{r^{(j)}} = \frac{1-e^{-\alpha^{(j)}}}{1-e^{-\alpha^{(1)}}}$. Therefore, it suffices to show that $\frac{\alpha^{(j)}}{\alpha^{(1)}} < \frac{1-e^{-\alpha^{(j)}}}{1-e^{-\alpha^{(1)}}}$, or $\frac{\alpha^{(j)}}{1-e^{-\alpha^{(j)}}} < \frac{\alpha^{(1)}}{1-e^{-\alpha^{(1)}}}$. This follows from the fact that the function $f(x) = \frac{x}{1-e^{-x}}$ is strictly increasing.

To prove (12), note that $\sigma^{(1)} = (1 + \sum_{j=2}^m (1 - \frac{r^{(j-1)}}{r^{(j)}}))^{-1}$, while $1 + \ln \frac{r^{(m)}}{r^{(1)}} = 1 + \sum_{j=2}^m \ln \frac{r^{(j)}}{r^{(j-1)}}$. Therefore, it suffices to show that for any $j = 2, \dots, m$, $\ln \frac{r^{(j)}}{r^{(j-1)}} > 1 - \frac{r^{(j-1)}}{r^{(j)}}$. Letting $x = \ln \frac{r^{(j-1)}}{r^{(j)}} < 0$, the desired inequality becomes $-x > 1 - e^x$, which is immediate.

For (13), we would like to prove that $\alpha < \alpha^{(1)}$. Note that $\alpha^{(1)}$ is the unique solution to

$$\alpha^{(1)} + \sum_{j=2}^m \left[-\ln \left(1 - (1 - e^{-\alpha^{(1)}}) \left(1 - \frac{r^{(j-1)}}{r^{(j)}} \right) \right) \right] = 1, \quad (33)$$

while α is the unique solution to

$$\alpha + \sum_{j=2}^m (1 - e^{-\alpha}) \ln \frac{r^{(j)}}{r^{(j-1)}} = 1. \quad (34)$$

The LHS of (33), as a function of $\alpha^{(1)}$, is increasing over $(0, 1)$; the same can be said about the LHS of (34) as a function of α . Therefore, it suffices to show that if $\alpha^{(1)} = \alpha = x$, then the LHS of (33) is strictly less than the LHS of (34), for all $x \in (0, 1)$.

Let $F = 1 - e^{-x}$ and consider any $j > 1$. Let $s = \frac{r^{(j-1)}}{r^{(j)}} \in (0, 1)$. It suffices to show that $-\ln(1 - F(1 - s)) < F \cdot \ln \frac{1}{s}$, which can be rearranged as $\frac{1-s^F}{1-s} > F$. For the final inequality, note that $f(s) = s^F$ is a strictly concave function on $(0, 1)$, since $F \in (0, 1)$. Therefore, $\frac{1-s^F}{1-s} > F$, because the LHS is the slope of the secant line through (s, s^F) and $(1, 1)$, while the RHS is the slope of the tangent line through $(1, 1)$. \square

Appendix B: Supplement to Section 3

The first subsection contains the deferred proofs from Section 3. In the second subsection, we explain how to optimize the randomized procedure for generating a single value function. In the third subsection, we put together the proof of Theorem 1.

The following inequality will be useful throughout the paper. For all $j = 2, \dots, m$, (7) says that $1 - e^{-\alpha^{(j)}} \leq 1 - \frac{r^{(j-1)}}{r^{(j)}}$, where we have used the fact that $1 - e^{-\alpha^{(1)}} \leq 1$. Therefore, for all $j = 2, \dots, m$, we can derive that

$$\frac{r^{(j-1)}}{r^{(j)}} \leq e^{-\alpha^{(j)}}. \quad (35)$$

B.1. Deferred Proofs

Proof of Theorem 4. Define $N_{t,i}$ to be the algorithm's value for N_i at the end of time t ($N_{0,i}$ is understood to be 0), for all $t \in [T]$ and $i \in [n]$. For all $t \in [T]$, define $R_t = r_{i_t^*}^{(j_t^*)}$ and $Z_t = \tilde{\Phi}_{i_t^*}(\tilde{L}_{i_t^*}^{(j_t^*)}) - \tilde{\Phi}_{i_t^*}(N_{i_t^*}/k_{i_t^*})$ if a sale was made during time t ; define $R_t = Z_t = 0$ otherwise.

Consider the solution to the dual LP (19) formed by setting $y_i = \mathbb{E}[\tilde{\Phi}_i(\frac{N_{T,i}}{k_i})]$ for all $i \in [n]$, and $z_t = \mathbb{E}[Z_t]$ for all $t \in [T]$. We claim that this solution is feasible. The non-negativity constraint (19c) can be verified directly from the definitions.

Now, consider constraint (19b) for a fixed $t \in [T], i \in [n], j \in [m_i]$. Given the initializations of $\tilde{L}_i^{(1)}, \dots, \tilde{L}_i^{(m_i)}, \tilde{\Phi}_i$ and the value of $N_{t-1,i}$, the algorithm will always make a decision during time t

which earns pseudorevenue whose conditional expectation is at least $p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i}))$, by definition (16). Formally,

$$\mathbb{E}[Z_t | \tilde{L}_i^{(1)}, \dots, \tilde{L}_i^{(m_i)}, \tilde{\Phi}_i, N_{t-1,i}] \geq p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})),$$

for all values of $\tilde{L}_i^{(1)}, \dots, \tilde{L}_i^{(m_i)}, \tilde{\Phi}_i, N_{t-1,i}$. By the tower property of conditional expectation, $z_t = \mathbb{E}[Z_t] \geq \mathbb{E}[p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i}))]$. Meanwhile, y_i has been set to $\mathbb{E}[\tilde{\Phi}_i(\frac{N_{T,i}}{k_i})]$. Since $N_{T,i} \geq N_{t-1,i}$ and $\tilde{\Phi}_i$ is increasing, $y_i \geq \mathbb{E}[\tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})]$. Therefore, the LHS of (19b), $p_{t,i}^{(j)}y_i + z_t$, is at least $\mathbb{E}[p_{t,i}^{(j)}(\tilde{\Phi}_i(\tilde{L}_i^{(j)}))]$. By (18), this is at least $r_i^{(j)}$, completing the proof of feasibility.

Applying weak duality, we obtain

$$\begin{aligned} \text{OPT}(\mathcal{S}, \mathcal{A}) &\leq \sum_{i=1}^n k_i \mathbb{E}[\tilde{\Phi}_i(\frac{N_{T,i}}{k_i})] + \sum_{t=1}^T \mathbb{E}[Z_t] \\ &= \sum_{i=1}^n k_i \mathbb{E}\left[\sum_{t=1}^T (\tilde{\Phi}_i(\frac{N_{t,i}}{k_i}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i}))\right] + \sum_{t=1}^T \mathbb{E}[Z_t] \\ &= \sum_{t=1}^T \mathbb{E}\left[\sum_{i=1}^n k_i (\tilde{\Phi}_i(\frac{N_{t,i}}{k_i}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})) + Z_t\right]. \end{aligned} \quad (36)$$

We now analyze the term inside the expectation,

$$\sum_{i=1}^n k_i (\tilde{\Phi}_i(\frac{N_{t,i}}{k_i}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})) + Z_t, \quad (37)$$

for every $t \in [T]$. We would like to argue that it is at most $\frac{R_t}{c}$, on every sample path.

There are two cases. If an item $i = i_t^*$ was sold at price $j = j_t^*$ during time t , then (37) equals

$$k_i (\tilde{\Phi}_i(\frac{N_{t-1,i}+1}{k_i}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})) + \tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i}). \quad (38)$$

Indeed, $N_{t,i} = N_{t-1,i} + 1$, $N_{t,i} = N_{t-1,i}$ for all $i \neq i_t^*$, and $Z_t = \tilde{\Phi}_i(\tilde{L}_i^{(j)}) - \tilde{\Phi}_i(\frac{N_{t-1,i}}{k_i})$ by definition. Furthermore, since Z_t is positive, $N_{t-1,i}$ must be less than $\tilde{L}_i^{(j)}k$. Therefore, we can invoke (17) to get that (38) is at most $r_i^{(j)}/c$, which is equal to $\frac{R_t}{c}$ by definition. In the other case, if no item was sold during time t , then (38) is 0, while $R_t = 0$ too, so (38) is still at most $\frac{R_t}{c}$.

Substituting back into (36), we conclude that $\text{OPT}(\mathcal{S}, \mathcal{A}) \leq \sum_{t=1}^T \mathbb{E}[\frac{R_t}{c}]$, which is equal to $\frac{1}{c} \mathbb{E}[\text{ALG}(\mathcal{S}, \mathcal{A})]$ by definition. This completes the proof of Algorithm 1 having a competitive ratio at least c . \square

Proof of Proposition 3. For (21), note that $\mathbb{E}[\tilde{L}^{(j)}] = \frac{\lfloor L^{(j)}k \rfloor + 1}{k} (L^{(j)}k - \lfloor L^{(j)}k \rfloor) + \frac{\lfloor L^{(j)}k \rfloor}{k} (1 - (L^{(j)}k - \lfloor L^{(j)}k \rfloor)) = \frac{1}{k} (L^{(j)}k - \lfloor L^{(j)}k \rfloor) + \frac{\lfloor L^{(j)}k \rfloor}{k} = L^{(j)}$.

For (22), note that $|(\tilde{L}^{(j)} - \tilde{L}^{(j')}) - (L^{(j)} - L^{(j')})| = |(\tilde{L}^{(j)} - L^{(j)}) - (\tilde{L}^{(j')} - L^{(j')})|$. We will prove that $(\tilde{L}^{(j)} - L^{(j)}) - (\tilde{L}^{(j')} - L^{(j')}) \leq \frac{1}{k}$; the inequality that $(\tilde{L}^{(j)} - L^{(j)}) - (\tilde{L}^{(j')} - L^{(j')}) \geq -\frac{1}{k}$ follows

by symmetry. The maximum value of $k\tilde{L}^{(j)}$ is $\lfloor kL^{(j)} \rfloor + 1$ while the minimum value of $k\tilde{L}^{(j')}$ is $\lfloor kL^{(j')} \rfloor$, hence the result is immediate unless $(\lfloor kL^{(j)} \rfloor + 1) - kL^{(j)} + kL^{(j')} - \lfloor kL^{(j')} \rfloor > 1$, i.e. $kL^{(j')} - \lfloor kL^{(j')} \rfloor > kL^{(j)} - \lfloor kL^{(j)} \rfloor$. However, in this case, if $k\tilde{L}^{(j)} = \lfloor kL^{(j)} \rfloor + 1$, then $W < kL^{(j)} - \lfloor kL^{(j)} \rfloor < kL^{(j')} - \lfloor kL^{(j')} \rfloor$ and hence $\tilde{L}^{(j')}$ is rounded up as well. Similarly, if $\tilde{L}^{(j')}$ is rounded down, then $\tilde{L}^{(j)}$ must be rounded down as well. If $\tilde{L}^{(j)}$ and $\tilde{L}^{(j')}$ are rounded in the same direction, then (iii) holds. \square

Proof of Theorem 5. First we prove (18), the claim that $\mathbb{E}[\tilde{\Phi}(\tilde{L}^{(j)})] \geq r^{(j)}$, inductively. Clearly $\mathbb{E}[\Phi(\tilde{L}^{(0)})] \geq r^{(0)} = 0$. Now consider $j \in [m]$ and suppose we have established (18) for the $j-1$ case. We can compare expression (20) with $q = \tilde{L}^{(j)}$ and $q = \tilde{L}^{(j-1)}$ to obtain $\tilde{\Phi}(\tilde{L}^{(j)}) = \tilde{\Phi}(\tilde{L}^{(j)}) + (r^{(j)} - r^{(j-1)}) \frac{\exp(\tilde{L}^{(j)} - \tilde{L}^{(j-1)}) - 1}{\exp(\alpha^{(j)}) - 1}$. Therefore,

$$\begin{aligned} \mathbb{E}[\tilde{\Phi}(\tilde{L}^{(j)})] &\geq r^{(j-1)} + (r^{(j)} - r^{(j-1)}) \frac{\mathbb{E}[\exp(\tilde{L}^{(j)} - \tilde{L}^{(j-1)})] - 1}{\exp(\alpha^{(j)}) - 1} \\ &\geq r^{(j-1)} + (r^{(j)} - r^{(j-1)}) \frac{\exp(\mathbb{E}[\tilde{L}^{(j)} - \tilde{L}^{(j-1)}]) - 1}{\exp(\alpha^{(j)}) - 1} \\ &= r^{(j-1)} + (r^{(j)} - r^{(j-1)}) \frac{\exp(\alpha^{(j)}) - 1}{\exp(\alpha^{(j)}) - 1} \end{aligned}$$

where the first inequality uses the induction hypothesis, and the second inequality uses Jensen's inequality (the exponential function \exp is convex). The equality follows from (21) and the definition that $\alpha^{(j)} = L^{(j)} - L^{(j-1)}$, completing the induction.

Now we prove (17) for an arbitrary $j \in [m]$ and $N \in \{0, \dots, \tilde{L}^{(j)}k - 1\}$. Let $q = \frac{N}{k}$ and $\ell = \tilde{\ell}(q)$. Note that $1 \leq \ell \leq j$, and $\tilde{L}^{(\ell-1)} \leq q < \tilde{L}^{(\ell)}$. Substituting $q = \frac{N}{k}$ into the LHS of (17), we get $k(\tilde{\Phi}(q + \frac{1}{k}) - \tilde{\Phi}(q)) + \tilde{\Phi}(\tilde{L}^{(j)}) - \tilde{\Phi}(q)$. Adding and subtracting $\tilde{\Phi}(\tilde{L}^{(\ell)})$ and rearranging, we get

$$k\left(\tilde{\Phi}\left(q + \frac{1}{k}\right) - \tilde{\Phi}(q)\right) + \tilde{\Phi}(\tilde{L}^{(\ell)}) - \tilde{\Phi}(q) + \tilde{\Phi}(\tilde{L}^{(j)}) - \tilde{\Phi}(\tilde{L}^{(\ell)}). \quad (39)$$

The following upper bound can be derived for expression (39):

$$\begin{aligned} &k\left(\tilde{\Phi}\left(q + \frac{1}{k}\right) - \tilde{\Phi}(q)\right) + \tilde{\Phi}(\tilde{L}^{(\ell)}) - \tilde{\Phi}(q) + \tilde{\Phi}(\tilde{L}^{(j)}) - \tilde{\Phi}(\tilde{L}^{(\ell)}) \\ &= (r^{(\ell)} - r^{(\ell-1)}) \frac{e^{q+1/k - \tilde{L}^{(\ell-1)}} (k - (k+1)e^{-1/k}) + e^{\tilde{L}^{(\ell)} - \tilde{L}^{(\ell-1)}}}{e^{\alpha^{(\ell)}} - 1} + \sum_{\ell'=\ell+1}^j (r^{(\ell')} - r^{(\ell'-1)}) \frac{e^{\tilde{L}^{(\ell')} - \tilde{L}^{(\ell'-1)}} - 1}{e^{\alpha^{(\ell')}} - 1} \\ &\leq (r^{(\ell)} - r^{(\ell-1)}) \frac{e^{\tilde{L}^{(\ell)} - \tilde{L}^{(\ell-1)}} (k - (k+1)e^{-1/k}) + e^{\tilde{L}^{(\ell)} - \tilde{L}^{(\ell-1)}}}{e^{\alpha^{(\ell)}} - 1} + \sum_{\ell'=\ell+1}^j (r^{(\ell')} - r^{(\ell'-1)}) \frac{e^{\tilde{L}^{(\ell')} - \tilde{L}^{(\ell'-1)}} - 1}{e^{\alpha^{(\ell')}} - 1} \\ &= (r^{(\ell)} - r^{(\ell-1)}) \frac{e^{\tilde{L}^{(\ell)} - \tilde{L}^{(\ell-1)} - \alpha^{(\ell)}} (1+k)(1 - e^{-1/k})}{1 - e^{-\alpha^{(\ell)}}} + \sum_{\ell'=\ell+1}^j (r^{(\ell')} - r^{(\ell'-1)}) \frac{e^{\tilde{L}^{(\ell')} - \tilde{L}^{(\ell'-1)} - \alpha^{(\ell')}} - e^{-\alpha^{(\ell')}}}{1 - e^{-\alpha^{(\ell')}}}. \quad (40) \end{aligned}$$

The inequality holds because $k - (1+k)e^{-1/k} > 0$ for all $k \in \mathbb{N}$, and q is at most $\tilde{L}^{(\ell)} - 1/k$.

It suffices to show that expression (40) is bounded from above by

$$r^{(j)} \frac{(1+k)(e^{1/k} - 1)}{1 - e^{-\alpha^{(1)}}}. \quad (41)$$

when $\ell' = \ell + 2$, $\ell'' = \ell$ if we arrived at case two during iteration $\ell + 1$ and $\ell'' = \ell + 1$ otherwise,...), we conclude that expression (40) is bounded from above by

$$(r^{(j)} - r^{(j-1)}) \frac{e^{\tilde{L}^{(j)} - \tilde{L}^{(\ell''-1)} - L^{(j)} + L^{(\ell''-1)}} (1+k)(1 - e^{-1/k})}{1 - e^{-\alpha^{(j)}}}$$

for some $\ell'' \in \{\ell, \dots, j\}$. The fact that $1 - e^{-\alpha^{(1)}} = \frac{r^{(j)}}{r^{(j)} - r^{(j-1)}}(1 - e^{-\alpha^{(j)}})$, due to (7), and the fact that $(\tilde{L}^{(j)} - \tilde{L}^{(\ell''-1)}) - (L^{(j)} - L^{(\ell''-1)}) \leq 1/k$, due to (22), complete the proof of expression (40) being at most expression (41), and thus the proof of Theorem 5 for general m .

Finally, when $m = 1$, $\alpha^{(1)} = 1$. In the above proof, since j and ℓ are always 1, (40) can be replaced by $r^{(1)} \cdot \frac{(1)(1+k)(1-e^{-1/k})}{1-e^{-1}}$, where we have used the fact that $L^{(1)} = k$ always. This is immediately at most $\frac{r^{(j)}}{c}$, for the improved value of $c = \frac{1-e^{-\alpha^{(1)}}}{(1+k)(1-e^{-1/k})}$, completing the proof of Theorem 5 in its entirety. \square

B.2. Optimizing the Randomized Procedure

We can explicitly formulate the optimization problem over randomized procedures for a single item with starting inventory k and m prices $r^{(1)}, \dots, r^{(m)}$. Using the ‘‘balls in bins’’ counting argument, the number of configurations satisfying (14) is $D := \binom{k+m-1}{m-1}$.

We refer to these configurations in an arbitrary order using the index $d \in [D]$, where we let ρ_d denote the probability of choosing configuration d , $f_d(\cdot)$ denote the value function for d , and $L_d^{(j)}$ denote the value of $\tilde{L}^{(j)}$ under configuration d for all $j = 0, \dots, m$. The optimization problem of satisfying (17)–(18) with a maximal value of c can be formulated as follows:

$$\tilde{F} := \sup c \tag{44a}$$

$$k(f_d(\frac{N+1}{k}) - f_d(\frac{N}{k})) + f_d(L_d^{(j)}) - f_d(\frac{N}{k}) \leq \frac{r^{(j)}}{c} \quad d \in [D], j \in [m], 0 \leq N \leq kL_d^{(j)} - 1 \tag{44b}$$

$$f_d(1) \geq \dots \geq f_d(\frac{1}{k}) \geq f_d(0) = 0 \quad d \in [D] \tag{44c}$$

$$\sum_{d=1}^D \rho_d f_d(L_d^{(j)}) \geq r^{(j)} \quad j \in [m] \tag{44d}$$

$$\sum_{d=1}^D \rho_d = 1 \tag{44e}$$

$$f_d(0), f_d(\frac{1}{k}), \dots, f_d(1) \in \mathbb{R}; \rho_d \geq 0 \quad d \in [D] \tag{44f}$$

Constraint (44b) corresponds to (17), constraint (44d) corresponds to (18), while constraint (44c) enforces the definition of a value function in (15). We let \tilde{F} denote the optimal objective value of (44). Unfortunately, it is difficult to solve (44) exactly, since the number of configurations D is exponential in the number of prices m , and constraint (44d) is non-linear.

Nonetheless, (44) is useful at determining the best competitive ratio which could be established *using our analysis*. We know that the randomized procedure from Definition 3 (based on Φ) is an optimal solution to (44) as $k \rightarrow \infty$, since it achieves the optimal competitive ratio possible.

We can also solve (44) exactly when $k = 1$, in which case $D = m$, where we will let $d \in [D]$ denote the configuration with $\tilde{L}^{(0)} = \dots = \tilde{L}^{(d-1)} = 0$ and $\tilde{L}^{(d)} = \dots = \tilde{L}^{(m)} = 1$. (44b) reduces to $2f_d(1) \leq \frac{r^{(j)}}{c}$, and needs to hold for $d \in [D]$, $j \geq d$ (for $j < d$, $kL_d^{(j)} - 1 = -1$). However, clearly only the constraint with $j = d$ is binding. As a result, (44b) corresponds to m constraints. (44d) corresponds to m constraints of the form $\sum_{d=1}^j \rho_d f_d(1) \geq r^{(j)}$, for $j \in [m]$.

Not counting $f_d(0)$, which must be set to 0, there are $2m + 1$ variables: $\{f_d(1), \rho_d : d \in [D]\}$ and c . Consider the system of equations obtained in these $2m + 1$ variables by setting (44b), (44d), and (44e) to equality. It can be checked that the unique solution is

$$f_d(1) = \frac{r^{(d)}}{\sigma^{(1)}}, \forall d \in [D]; \rho_d = \sigma^{(d)}, \forall d \in [D]; c = \frac{\sigma^{(1)}}{2} \quad (45)$$

with $\sigma^{(1)}, \dots, \sigma^{(m)}$ defined from $r^{(1)}, \dots, r^{(m)}$ according to (8). Furthermore, this solution is both feasible, satisfying the non-negativity constraints in (44c) and (44f), and optimal. Therefore, the value of \tilde{F} is $\frac{\sigma^{(1)}}{2}$.

B.3. Proof of Theorem 1

Now we put together the proof of Theorem 1. For all items $i \in [n]$, \tilde{F}_i is defined to be the optimal objective value of (44), with $k = k_i$, $m = m_i$, and $r^{(1)} = r_i^{(1)}, \dots, r^{(m)} = r_i^{(m_i)}$. Consider Algorithm 1, where for all i , the randomized procedure used to initialize $\tilde{\Phi}_i$ is an optimal solution to (44) achieving the objective value of \tilde{F}_i . For all i , (17)–(18) is satisfied as long as $c \leq \tilde{F}_i$. Therefore, the maximum value of c satisfying the conditions of Theorem 4 is $\min_i \tilde{F}_i$. By Theorem 4, this algorithm achieves a competitive ratio of $\min_i \tilde{F}_i$.

To establish bounds (i)–(iii) from Theorem 1, for all i , we need to find a feasible randomized procedure with an objective value in (44) equal to the bound. For bounds (i) and (iii), this is established directly by the randomized procedure from Definition 3, via the statement of Theorem 5. For bound (ii), we need to split the k_i units of item i into k_i disparate items. For each single-unit item, its value function in Algorithm 1 is initialized according to the randomized procedure described by (45). This yields a value of $\frac{\sigma_i^{(1)}}{2}$, completing the proof of Theorem 1.

Appendix C: Deferred Proofs from Section 4

Proof of Lemma 2. Since the algorithm was willing to sell item i at price j , it must be the case that $W_i < L_i^{(j)}$. Let ℓ denote $\ell_i(W_i)$, which is at most j . We ignore measure-zero events and assume that $W_i \neq L_i^{(\ell-1)}$. We can rearrange Z_t as

$$\begin{aligned} & r_i^{(j)} - r_i^{(\ell)} + r_i^{(\ell)} - \left(r_i^{(\ell-1)} + (r_i^{(\ell)} - r_i^{(\ell-1)}) \frac{\exp(W_i - L_i^{(\ell-1)}) - 1}{\exp(\alpha_i^{(\ell)}) - 1} \right) \\ &= r_i^{(j)} - r_i^{(\ell)} + (r_i^{(\ell)} - r_i^{(\ell-1)}) \frac{\exp(\alpha_i^{(\ell)}) - \exp(W_i - L_i^{(\ell-1)})}{\exp(\alpha_i^{(\ell)}) - 1}. \end{aligned}$$

Adding $Y_i = \Phi'_{\mathcal{P}_i}(W_i) = (r_i^{(\ell)} - r_i^{(\ell-1)}) \frac{\exp(W_i - L_i^{(\ell-1)})}{\exp(\alpha_i^{(\ell)}) - 1}$ to this expression, we get $r_i^{(j)} - r_i^{(\ell)} + \frac{r_i^{(\ell)} - r_i^{(\ell-1)}}{1 - \exp(-\alpha_i^{(\ell)})}$, which can be re-written as $r_i^{(j)} - r_i^{(\ell)} + \frac{r_i^{(\ell)}}{1 - \exp(-\alpha_i^{(1)})}$ due to (7). The result follows immediately. \square

Proof of Lemma 3. It suffices to show that constraint (19b) holds for all $t \in [T]$ and $i \in [n]$. Since $p_{t,i}^{(j)} \in \{0, 1\}$ and the constraint clearly holds when $p_{t,i}^{(j)} = 0$, it suffices to show that $\mathbb{E}[Y_i + Z_t] \geq r_i^{(j_{t,i})}$, where $j_{t,i} \neq 0$. We will let $j = j_{t,i}$ for brevity.

Fix the realization of $W_{i'}$ for all $i' \neq i$, and consider the run of the algorithm on a modified setup with item i removed. Having fixed the values of $W_{i'}$, such a run is deterministic. Let Z^{crit} denote the pseudorevenue earned on this run during time t , possibly 0. $\Phi_{\mathcal{P}_i}$ maps $[0, L_i^{(j)}]$ to $[0, r_i^{(j)}]$ bijectively, so we can set W^{crit} to be the value in $[0, L_i^{(j)}]$ for which $\Phi_{\mathcal{P}_i}(W^{\text{crit}}) = \max\{r_i^{(j)} - Z^{\text{crit}}, 0\}$.

We now consider the run of the algorithm on the full setup with item i , which is dependent on the realization of W_i . The following two claims from Devanur et al. (2013) generalize to our multi-price setting.

1. Dominance: if $W_i \in [0, W^{\text{crit}})$, then in the run with item i , item i gets matched.

Proof: Since $W^{\text{crit}} > W_i$ and $W_i \geq 0$, $W^{\text{crit}} > 0$. Therefore, $\Phi_{\mathcal{P}_i}(W^{\text{crit}}) > 0$. Thus $\Phi_{\mathcal{P}_i}(W^{\text{crit}}) = r_i^{(j)} - Z^{\text{crit}}$ (as opposed to $\Phi_{\mathcal{P}_i}(W^{\text{crit}}) = 0$), and moreover since $W_i < W^{\text{crit}}$ and $\Phi_{\mathcal{P}_i}$ is strictly increasing, $\Phi_{\mathcal{P}_i}(W_i) < r_i^{(j)} - Z^{\text{crit}}$. This implies $r_i^{(j)} - \Phi_{\mathcal{P}_i}(W_i) > \max\{Z^{\text{crit}}, 0\}$, since $Z^{\text{crit}} \geq 0$. Thus on the run with item i , either i is already matched before time t , or it is matched to customer t .

2. Monotonicity: $Z_t \geq Z^{\text{crit}}$ (regardless of the realization of W_i).

Proof: fix the realization of W_i . We compare two deterministic runs of the algorithm: one with item i , and one without. We can inductively establish over $t = 0, \dots, T$ that at the end of time t , the set of unmatched items in the run with i is a superset of that in the run without i . Therefore, in the run with i , since the algorithm is maximizing pseudorevenue over a superset of items, its pseudorevenue Z_t can be no less than Z^{crit} .

Now, conditioned on the realizations of $W_{i'}$ for $i' \neq i$, which determines the values of Z^{crit} and W^{crit} , we have $Z_t \geq Z^{\text{crit}}$ (by Monotonicity) and in turn $Z^{\text{crit}} \geq r_i^{(j)} - \Phi_{\mathcal{P}_i}(W^{\text{crit}})$ (by the definition of W^{crit}). Meanwhile, as long as i gets matched, Y_i gets set to $\Phi'_{\mathcal{P}_i}(W_i)$, so by Dominance, $\mathbb{E}[Y_i | \{W_{i'} : i' \neq i\}] \geq \int_0^{W^{\text{crit}}} \Phi'_{\mathcal{P}_i}(w) dw = \Phi_{\mathcal{P}_i}(W^{\text{crit}}) - \Phi_{\mathcal{P}_i}(0) = \Phi_{\mathcal{P}_i}(W^{\text{crit}})$. Therefore, $\mathbb{E}[Y_i + Z_t | \{W_{i'} : i' \neq i\}] \geq r_i^{(j)}$. The proof follows from the tower property of conditional expectation. \square

Appendix D: Deferred Proofs from Section 5

Proof of Proposition 4. The unique solution to the system (24) is obtained inductively over $j = 2, \dots, m$ by setting $B_j = \frac{r^{(j-1)} e^{-\alpha(j-1)}}{r^{(j)} e^{-\alpha(j)}} B_{j-1}$. By (35), $\frac{r^{(j-1)}}{r^{(j)}} \leq e^{-\alpha(j)}$, hence $B_j \leq e^{-\alpha(j-1)} B_{j-1}$. But $\alpha^{(j-1)} > 0$ by Proposition 1, completing the proof that $B_j < B_{j-1}$ for $j = 2, \dots, m$. The fact that $0 < B_m$ is immediate. \square

Proof of Lemma 4. Consider the execution of an online algorithm with this randomized arrival sequence. For all $i \in [n]$ and group of customers $t \in [n]$, let $Q_{t,i}$ denote the number of group- t customers to which item π_i is sold, which is a random variable with respect to the random permutation π as well as any randomness in the algorithm. Let $q_{t,i} = \mathbb{E}[Q_{t,i}]$.

Clearly if $i < t$, then $Q_{t,i} = 0$, because group- t customers have no interest in item π_i . Otherwise, for any $i, i' \geq t$, we argue that $q_{t,i} = q_{t,i'}$. This is because while group t is arriving, the online

algorithm cannot distinguish between items π_i and $\pi_{i'}$, hence any items it allocates are equally likely to be item π_i and item $\pi_{i'}$. Therefore, we let q_t denote the value of $q_{t,i}$ for $i \geq t$.

Now, consider item π_n . Since it only has k units of inventory, we know that $\sum_{t=1}^n Q_{t,n} \leq k$ on every sample path. Using the linearity of expectation, we get that

$$\sum_{t=1}^n q_t \leq k. \quad (46)$$

Furthermore, for a $t \in [n]$, on every sample path, $\sum_{i=t}^n Q_{t,i} \leq k$, since there are only k customers in group t . Therefore, $(n+1-t)q_t \leq k$, or

$$q_t \leq \frac{k}{n+1-t}. \quad (47)$$

For this proof, let $M_j = \sum_{j'=1}^j \beta_{j'}$, for all $j = 0, \dots, m$. For all $j \in [m]$, let $\lambda_j = \frac{1}{k} \sum_{t=M_{j-1}n+1}^{M_j n} q_t$. Substituting into (46), we get the constraint that $\sum_{j=1}^m \lambda_j \leq 1$. For any $j \in [m-1]$, summing inequality (47) for $t = M_{j-1}n+1, \dots, M_j n$ yields $\lambda_j \leq \ln \frac{B_j}{B_{j+1}}$, since $n \rightarrow \infty$, and $B_j = 1 - M_{j-1}$, $B_{j+1} = 1 - M_j$ by definition. It is also clear from definition that $\lambda_j \geq 0$ for all $j \in [m]$.

Finally, the total expected revenue is

$$\sum_{j=1}^m r^{(j)} \sum_{t=M_{j-1}n+1}^{M_j n} q_t (n+1-t), \quad (48)$$

since for each group t there are $n+1-t$ items for each of which q_t copies are sold in expectation.

Consider any $j \in [m]$. Since $\sum_{t=M_{j-1}n+1}^{M_j n} q_t = \lambda_j k$ by definition, $\sum_{t=M_{j-1}n+1}^{M_j n} q_t (n+1-t)$ is maximized by setting q_t to its upper bound in (47) for $t = M_{j-1}n+1, M_{j-1}n+2, \dots$ until the capacity of $\lambda_j k$ is reached. Since $n \rightarrow \infty$, we can simply compute the value of t for which

$$\frac{k}{n - M_{j-1}n} + \dots + \frac{k}{n-t} = \lambda_j k, \quad (49)$$

with $t \in [M_{j-1}n, M_j n]$. Letting $t = (M_{j-1} + y\beta_j)n$ with $y \in [0, 1]$, and using the definition of B_j ,

(49) becomes $\ln \frac{B_j}{B_{j-y\beta_j}} = \lambda_j$, or $y\beta_j = B_j(1 - e^{-\lambda_j})$. Therefore,

$$\begin{aligned} \sum_{t=M_{j-1}n+1}^{M_j n} q_t (n+1-t) &\leq \sum_{t=M_{j-1}n+1}^{(M_{j-1} + B_j(1 - e^{-\lambda_j}))n} \frac{k}{n+1-t} \cdot (n+1-t) \\ &= B_j(1 - e^{-\lambda_j})nk \end{aligned}$$

Substituting into (48), we get that the expected revenue of the online algorithm is at most (26), where $\sum_{j=1}^m \lambda_j \leq 1$, $\lambda_j \leq \ln \frac{B_j}{B_{j+1}}$ for $j \in [m-1]$, and $\lambda_j \geq 0$ for $j \in [m]$, completing the proof. \square

Proof of Lemma 5. We use backward induction over $j = m, \dots, 1$. When $j = m$, (28) becomes $nk r^{(m)} B_m (1 - \exp(-\tau))$, since $A_m = \alpha^{(m)}$ by definition. Meanwhile, (27) is maximized by setting $\lambda_m = \tau$, resulting in the same expression and establishing the base case.

Now suppose $j < m$ and that we have already established the lemma in the $j+1$ case. If we set $\lambda_j = \lambda$, for some $\lambda \in [0, \tau]$, then the maximum value of (27) subject to $\lambda_{j+1}, \dots, \lambda_m \geq 0$ and $\lambda_{j+1} + \dots + \lambda_m \leq \tau - \lambda$ is, by the inductive hypothesis,

$$r^{(j)} B_j (1 - \exp(-\lambda)) nk + nk \sum_{\ell=j+1}^m r^{(\ell)} B_\ell \left(1 - \exp \left(-\alpha^{(\ell)} + \frac{A_{j+1} - (\tau - \lambda)}{m - (j+1) + 1} \right) \right). \quad (50)$$

Consider this expression as a function of λ . The derivative is

$$r^{(j)} B_j \exp(-\lambda) nk + nk \sum_{\ell=j+1}^m r^{(\ell)} B_\ell \cdot \frac{-1}{m-j} \cdot \exp \left(-\alpha^{(\ell)} + \frac{A_{j+1} - (\tau - \lambda)}{m-j} \right) \quad (51)$$

and the second derivative is clearly negative, so the function is concave. Therefore, it is maximized by setting the derivative to 0. By definition (24), $r^{(\ell)} B_\ell e^{-\alpha^{(\ell)}}$ is identical for all $\ell = j+1, \dots, m$, and equal to $r^{(j)} B_j e^{-\alpha^{(j)}}$. Thus setting (51) to 0 implies:

$$\begin{aligned} \exp(\alpha^{(j)} - \lambda) &= \frac{1}{m-j} \sum_{\ell=j+1}^m \exp \left(\frac{A_{j+1} - (\tau - \lambda)}{m-j} \right) \\ \alpha^{(j)} - \lambda &= \frac{A_{j+1} - (\tau - \lambda)}{m-j}. \end{aligned}$$

Rearranging and using the definition that $A_{j+1} = A_j - \alpha^{(j)}$, we get $\lambda = \alpha^{(j)} - \frac{A_j - \tau}{m-j+1}$. Substituting this value of λ into (50), the expression $\frac{A_{j+1} - (\tau - \lambda)}{m - (j+1) + 1}$ is equal to $\frac{A_j - \tau}{m-j+1}$, hence (50) is equal to (28), completing the induction and the proof of the lemma. \square

Appendix E: Deriving the Multi-price Value Function $\Phi_{\mathcal{P}}$

In this section we explain how we optimized the value function $\Phi_{\mathcal{P}}$ for a given price set \mathcal{P} , leading to the system of equations in (7), and the functional form in (9). In Appendix E.1, we use the same method to derive the optimal value function when the price of an item can take any value in the continuum $[r^{\min}, r^{\max}]$.

Consider constraints (17)–(18) in Theorem 5 for a single item with $k \rightarrow \infty$. Let $w = \frac{N}{k}$, and we deterministically set $\tilde{\Phi}$ to some Φ . The goal is to solve for the Φ which maximizes the value of F .

Observe that

$$\lim_{k \rightarrow \infty} k(\Phi(\frac{N+1}{k}) - \Phi(\frac{N}{k})) = \lim_{k \rightarrow \infty} \frac{\Phi(w + 1/k) - \Phi(w)}{1/k},$$

which is equal to the derivative of Φ as w , by definition (Φ will end up not being differentiable on a discrete set of measure 0, which can be ignored). Therefore, (17) is equivalent to

$$\Phi'(w) - \Phi(w) \leq r^{(j)}(\frac{1}{F} - 1), \quad (52)$$

and needs to hold for all $j \in [m]$, $w \in [0, L^{(j)}]$. For a fixed $w \in (L^{(j-1)}, L^{(j)})$, (52) needs to hold for all $j' = j, \dots, m$, but is clearly binding when $j' = j$. Therefore, it suffices to fix a $j \in [m]$ and consider (52) when $w \in (L^{(j-1)}, L^{(j)})$.

We should point out that this simplification via the “binding” argument is not possible for a finite k and random $\tilde{\Phi}$, because then (52) becomes $\tilde{\Phi}'(w) - \tilde{\Phi}(w) \leq \frac{r^{(j)}}{F} - \tilde{\Phi}(L^{(j)})$, and the RHS in fact may not be increasing in j . This is why we resort to first solving for Φ when $k \rightarrow \infty$ and then defining $\tilde{\Phi}$ as a random perturbation of Φ .

If we set (52) to equality for some $j \in [m]$ and all $w \in (L^{(j-1)}, L^{(j)})$, and solve the differential equation, we get that $\Phi(w)$ must be of the form $Ce^w - r^{(j)}(\frac{1}{F} - 1)$ on $(L^{(j-1)}, L^{(j)})$. Setting $\Phi(L^{(j-1)}) = r^{(j-1)}$ and $\Phi(L^{(j)}) = r^{(j)}$, we obtain

$$\begin{aligned} C &= \frac{r^{(j)} - r^{(j-1)}}{e^{L^{(j)}} - e^{L^{(j-1)}}}; \\ F &= \frac{1}{1 - \frac{r^{(j-1)}}{r^{(j)}}} \cdot (1 - e^{-\alpha^{(j)}}). \end{aligned} \quad (53)$$

The RHS of (53) is the largest value of F which allows (52) to hold on segment j . It is dependent on $\alpha^{(j)}$, which is equal to $L^{(j)} - L^{(j-1)}$, the length of segment j . For (52) to hold on all segments $j \in [m]$, F must be set to $\min_j \frac{1}{1 - r^{(j-1)}/r^{(j)}} \cdot (1 - e^{-\alpha^{(j)}})$.

Therefore, we would like to choose segment lengths $\alpha^{(1)}, \dots, \alpha^{(m)}$ summing to 1 to maximize the minimum $\frac{1}{1 - r^{(j-1)}/r^{(j)}} \cdot (1 - e^{-\alpha^{(j)}})$, which is accomplished by setting $\frac{1}{1 - r^{(j-1)}/r^{(j)}} \cdot (1 - e^{-\alpha^{(j)}})$ equal

for all $j \in [m]$. This yields the system of equations (7), and Proposition 1. The resulting value of F is equal to $1 - e^{-\alpha^{(1)}}$, since $r^{(0)} = 0$. The resulting value of C , when substituted into the equation for $\Phi(w)$ on each segment $(L^{(j-1)}, L^{(j)})$, yields (9).

The derivation of Φ we just completed, starting from condition (52), comes from our analysis of MULTI-PRICE BALANCE. We note that the exact same inequality (52) can also be derived from our analysis of MULTI-PRICE RANKING, which shows that the same value function should be used for both algorithms.

E.1. Continuum of Feasible Prices

Let the feasible price set for the item be $[r^{\min}, r^{\max}]$, where $0 < r^{\min} < r^{\max}$. Using the same “binding” argument, it suffices to maximize the value of F for which the following can hold:

$$\Phi'(w) - \Phi(w) \leq r^{\min} \left(\frac{1}{F} - 1 \right), \quad w \in (0, \alpha); \quad (54)$$

$$\Phi'(w) - \frac{\Phi(w)}{F} \leq 0, \quad w \in (\alpha, 1). \quad (55)$$

Φ must also satisfy $\Phi(0) = 0$, $\Phi(\alpha) = r^{\min}$, $\Phi(1) = r^{\max}$, while $\alpha \in (0, 1)$ is an arbitrary “booking limit” for the lowest price of r^{\min} .

We know from before that under the optimal solution to (54), the value of F can be at most $1 - e^{-\alpha}$. Solving the differential equation where (55) is set to equality, $\Phi(w)$ must take the form $Ce^{w/F}$ on $(\alpha, 1)$. Substituting $\Phi(\alpha) = r^{\min}$ and $\Phi(1) = r^{\max}$ yields

$$C = (r^{\min})^{\frac{1}{1-\alpha}} (r^{\max})^{-\frac{\alpha}{1-\alpha}};$$

$$F = \frac{1 - \alpha}{\ln \frac{r^{\max}}{r^{\min}}}.$$

Therefore, the value of F is also bounded from above by $\frac{1-\alpha}{\ln(r^{\max}/r^{\min})}$. F is maximized by setting $\frac{1-\alpha}{\ln(r^{\max}/r^{\min})}$ equal to the other upper bound of $1 - e^{-\alpha}$; the value at which equality is achieved is then the competitive ratio.

Letting $R = \ln(r^{\max}/r^{\min})$, the solution to $\frac{1-\alpha}{R} = 1 - e^{-\alpha}$ can be written as $W(Re^{R-1}) - R + 1$, where W is the Lambert-W function, the inverse function to $f(x) = xe^x$ for $x \in \mathbb{R}_{\geq 0}$. Indeed, when $\alpha = W(Re^{R-1})$, the following can be derived:

$$\frac{1 - \alpha}{R} = 1 - e^{-\alpha}$$

$$Re^{-\alpha} = \alpha + R - 1$$

$$Re^{R-1} = (\alpha + R - 1)e^{\alpha+R-1}$$

$$W(Re^{R-1}) = \alpha + R - 1$$

Substituting $\alpha = W(\ln(r^{\max}/r^{\min})e^{\ln(r^{\max}/r^{\min})-1}) - \ln(r^{\max}/r^{\min}) + 1$ into the formula for C , and using the fact that $\Phi(w) = Ce^{w/F}$, we get

$$\Phi(w) = (r^{\min})^{\frac{1-w}{1-\alpha}} (r^{\max})^{\frac{w-\alpha}{1-\alpha}}, \quad w \in [\alpha, 1].$$

Meanwhile, the earlier derivation implies that

$$\Phi(w) = r^{\min} \cdot \frac{e^w - 1}{e^\alpha - 1}, \quad w \in [0, \alpha].$$

It can be checked that indeed $\Phi(0) = 0$, $\Phi(\alpha) = r^{\min}$ (Φ is continuous at $w = \alpha$), and $\Phi(1) = r^{\max}$. Furthermore, unlike the case of discrete prices, it can be checked that Φ is also differentiable at $w = \alpha$ (on $[\alpha, 1]$, use the form that $\Phi(w) = Ce^{w/F}$, hence $\Phi'(\alpha) = \frac{\Phi(\alpha)}{F}$).

Appendix F: Supplement to Numerical Experiments

We provide additional details about our choice estimation. We define 8 customer types, one for each combination of the 3 following binary features.

1. Group: whether the customer indicated a party size greater than 1.
2. CRO: whether the customer booked using the Central Reservation Office, as opposed to the hotel's website or a Global Distribution System (for details on these terms, see Bodea et al. (2009)).
3. VIP: whether the customer had any kind of VIP status.

We did not use features such as: whether the booking date is a weekend, whether the check-in date is a weekend, the length of stay, or the number of days in advance booked. Such features did not result in a more predictive model.

We estimate the mean MNL utilities for each of the 8 products separately for each customer type. The results are displayed in Table 3. The total share of each customer type (out of all the

Table 3 MNL choice models for the 8 customer types. The suffix “L” on a room type means lower fare, while the suffix “H” on a room type means higher fare.

Customer Type			MNL Mean Utilities									
Group?	CRO?	VIP?	Share	KingL	QueenL	SuiteL	2DoubleL	KingH	QueenH	SuiteH	2DoubleH	NoBuy
		✓	0.16	-0.36	-1.22	-2.56	-1.04	0	-0.23	-2.25	-1.8	0
	✓		0.03	-0.82	-1.98	-2.16	-2.09	0	-1.02	-1.45	-1.82	0
	✓	✓	0.28	-1.67	$-\infty$	-3.78	-2.71	0	-1.33	-1.8	-1.58	0
	✓	✓	0.09	-2.13	$-\infty$	-3.38	-3.76	0	-2.12	-1	-1.59	0
✓			0.19	-0.54	-0.97	-2.26	0	-0.91	-1.47	-2.78	-1.41	0
✓		✓	0.04	-0.09	-0.82	-0.95	-0.14	0	-1.35	-1.07	-0.51	0
✓	✓		0.18	-0.93	$-\infty$	-2.56	-0.76	0	-1.66	-1.41	-0.27	0
✓	✓	✓	0.03	-1.39	$-\infty$	-2.16	-1.8	0	-2.45	-0.61	-0.28	0

transactions) is also displayed. We should point out that it is possible for a customer to choose the higher fare for a room, even if the lower fare was also offered. This is because the higher fares are often packaged with additional offers, such as airline services, city attractions, in-room services, etc.


We have shifted the mean utilities so that for each customer type, the weights of both the no-purchase option, and the most-preferred purchase option, is equal to 0. (We synthetically set the weight of the no-purchase option because it is not possible to estimate from the data.) The large weights on the no-purchase options ensure that the revenue-maximizing assortments tend to include both the low and high fares.

In the setting with greater fare differentiation (Subsection 7.5), the high prices of the King, Queen, Suite, and Two-double rooms are adjusted to \$614, \$608, \$768, \$612, respectively (twice the lower fares). The mean utility of the no-purchase option is increased by 2 for every customer type, to ensure that the revenue-maximizing assortments still include both the low and high fares.

F.1. Details on the Forecasting Bid-price Algorithms

To forecast the remaining number of customers, we assume that we know the average number of customers interested in each occupancy date (1340), as well as the overall trend for how far in advance customers book, which is plotted in Figure 4. As an example of how to use these numbers, consider the occupancy date March 31st. At the start, we forecast there to be 1340 arrivals. However, suppose by March 6th, 500 customers have arrived. Since we know from Figure 4 that roughly 50% of the total population interested in March 31st will have already booked by March 6th (25 days in advance), we expect there to only be 500 customers remaining.

Figure 4 Distribution of arrivals over the days before check-in, formed by aggregating all transactions.



To forecast the breakdown of remaining customers by type, we assume that we know the aggregate distribution of customer type over all occupancy dates. For example, from Appendix F, we know that 28% of all customers are of Type 3. Then we would estimate $28\% \times 500 = 140$ of the 500 remaining customers to be of Type 3. Alternatively, one can try to learn the specific distribution of customers interested in March 31st. Suppose that only 100, or 20%, of the 500 bookings made before March 6th came from customers of Type 3. Then we would instead estimate $20\% \times 500 = 100$ of the 500 remaining customers to be of Type 3.

To use the forecasted information, algorithms incorporate it into the LP (30), and set the *bid price* of each item i equal to the shadow price of constraint i in (30b). These algorithms then offer each customer t the assortment S (from the available items) maximizing $\sum_{(i,j) \in S} p_{t,i}^{(j)}(S)(r_i^{(j)} - \lambda_i)$.

We clarify the exact way in which the forecasted information is incorporated into the LP. Let there be A customer types, indexed by $a = 1, \dots, A$. We use $p_{a,i}^{(j)}(S)$ to denote the probability of a customer of type a choosing product (i, j) from assortment S . Suppose that when we want to re-solve the LP (30), the forecasted number of remaining customers of type a is N_a , for all $a \in [A]$, and the remaining inventory of item i is K_i , for all $i \in [n]$. We can formulate the following LP, which is a modification of (30):

$$\max \sum_{a=1}^A \sum_S x_a(S) \sum_{(i,j) \in S} r_i^{(j)} p_{a,i}^{(j)}(S)$$

$$\begin{aligned}
\sum_{a=1}^A \sum_S x_a(S) \sum_{j:(i,j) \in S} p_{a,i}^{(j)}(S) &\leq k_i && i \in [n] \\
\sum_S x_a(S) &= N_a && a \in [A] \\
x_a(S) &\geq 0 && a \in [A], S \subseteq \{(i,j) : i \in [n], j \in [m_i]\}
\end{aligned}$$

We have set $T = \sum_{a=1}^A N_a$ and $|\{t : \text{type of customer } t \text{ is } a\}| = N_a$; note that the ordering of remaining customers is inconsequential for the LP.

Although this LP has an exponential number of variables, we can easily solve it using column generation (e.g., see Liu and Van Ryzin (2008)). Fix an optimal primal solution $(x_a^*(S) : a \in [A], S \subseteq \{(i,j) : i \in [n], j \in [m_i]\})$ and an optimal dual solution $(y_i^* : i \in [n]), (z_a^* : a \in [A])$. The bid-price algorithm sets the bid price of each item i equal to y_i^* .

We should point out that for every bid-price algorithm based on dual variables, there is a corresponding *random assignment* algorithm based on primal variables. Such an algorithm would, for each customer type a , offer each assortment S with probability proportional to $x_a^*(S)$. We have confirmed that these algorithms perform similarly in the simulations. We compare with the bid-price algorithms instead of the random assignment algorithms because they follow a form more similar to our MULTI-PRICE BALANCE algorithm.