



## Operations Research

Publication details, including instructions for authors and subscription information:  
<http://pubsonline.informs.org>

### An Approximation Algorithm for Network Revenue Management Under Nonstationary Arrivals

Yuhang Ma, Paat Rusmevichientong, Mika Sumida, Huseyin Topaloglu

#### To cite this article:

Yuhang Ma, Paat Rusmevichientong, Mika Sumida, Huseyin Topaloglu (2020) An Approximation Algorithm for Network Revenue Management Under Nonstationary Arrivals. *Operations Research* 68(3):834-855. <https://doi.org/10.1287/opre.2019.1931>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact [permissions@informs.org](mailto:permissions@informs.org).

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

## Crosscutting Areas

# An Approximation Algorithm for Network Revenue Management Under Nonstationary Arrivals

Yuhang Ma,<sup>a</sup> Paat Rusmevichientong,<sup>b</sup> Mika Sumida,<sup>a</sup> Huseyin Topaloglu<sup>a</sup>

<sup>a</sup>School of Operations Research and Information Engineering, Cornell Tech, New York, New York 10044; <sup>b</sup>Marshall School of Business, University of Southern California, Los Angeles, California 90089

Contact: [ym367@cornell.edu](mailto:ym367@cornell.edu) (YM); [rusmevic@marshall.usc.edu](mailto:rusmevic@marshall.usc.edu),  <https://orcid.org/0000-0001-9584-4203> (PR); [ms3268@cornell.edu](mailto:ms3268@cornell.edu) (MS); [topaloglu@orie.cornell.edu](mailto:topaloglu@orie.cornell.edu),  <https://orcid.org/0000-0002-3049-6719> (HT)

Received: April 8, 2018

Revised: March 5, 2019

Accepted: August 9, 2019

Published Online in Articles in Advance:  
April 28, 2020

Subject Classification: dynamic programming;  
marketing: pricing

Area of Review: Revenue Management and  
Market Analytics

<https://doi.org/10.1287/opre.2019.1931>

Copyright: © 2020 INFORMS

**Abstract.** We provide an approximation algorithm for network revenue management problems. In our approximation algorithm, we construct an approximate policy using value function approximations that are expressed as linear combinations of basis functions. We use a backward recursion to compute the coefficients of the basis functions in the linear combinations. If each product uses at most  $L$  resources, then the total expected revenue obtained by our approximate policy is at least  $1/(1+L)$  of the optimal total expected revenue. In many network revenue management settings, although the number of resources and products can become large, the number of resources used by a product remains bounded. In this case, our approximate policy provides a constant-factor performance guarantee. Our approximate policy can handle nonstationarities in the customer arrival process. To our knowledge, our approximate policy is the first approximation algorithm for network revenue management problems under nonstationary arrivals. Our approach can incorporate the customer choice behavior among the products, and allows the products to use multiple units of a resource, while still maintaining the performance guarantee. In our computational experiments, we demonstrate that our approximate policy performs quite well, providing total expected revenues that are substantially better than its theoretical performance guarantee.

**Funding:** P. Rusmevichientong and H. Topaloglu were supported by the National Science Foundation Division of Civil, Mechanical and Manufacturing Innovation [Grants 1824860 and 1825406]. H. Topaloglu was also supported in part by Schmidt Sciences.

**Supplemental Material:** The e-companion is available at <https://doi.org/10.1287/opre.2019.1931>.

**Keywords:** network revenue management • approximate dynamic programming • dynamic pricing

## 1. Introduction

In network revenue management problems, we manage the limited capacities for a collection of resources to satisfy the requests for different products that arrive randomly over time. Such problems find applications in a variety of settings, including airlines, hospitality, railways, and cloud computing. In airlines, for example, the resources are the flight legs and the products are the itineraries offered to customers that can consume capacities on multiple flight legs. In hospitality, on the other hand, the resources are the availabilities of hotel rooms on each day and the products are the multiple night stays offered to customers that can consume capacities on multiple days. The main tradeoff in network revenue management problems involves keeping a balance between accepting a product request that is currently in the system to generate some immediate revenue and reserving the resource capacities for a potentially more profitable product request that can arrive in the future. Nevertheless, the key difficulty in

finding the optimal course of action arises from the fact that serving a request for a product consumes the capacities of the different resources used by the product. Thus, computing the optimal course of action requires keeping track of the remaining capacities for all the resources simultaneously, creating the curse of dimensionality as the number of resources increases.

In this paper, we provide an approximation algorithm for network revenue management problems. Our problem setup follows the standard network revenue management literature. We have resources with limited capacities that can be used to serve the requests for products arriving randomly over a finite selling horizon. At each time period in the selling horizon, a customer arrives with a request for a particular product. If we accept the product request, then we generate a certain revenue and consume the capacities of the different resources used by the product. The arrivals of the product requests can have arbitrary nonstationarities, but they are independent between the time periods. The goal is to find

a policy to determine which product requests to accept to maximize the total expected revenue over the selling horizon. The dynamic programming formulation for this problem requires a high-dimensional state variable that keeps track of the remaining resource capacities. Thus, it is intractable to compute the optimal policy.

### 1.1. Contributions

Letting  $L$  be the maximum number of resources used by a product, we give an approximate policy that is guaranteed to obtain at least  $1/(1 + L)$  of the optimal total expected revenue. In many network revenue management settings, the number of resources and products can become large, but the number of resources used by a product remains bounded. In airlines, for example,  $L$  corresponds to the maximum number of flight legs included in an itinerary, which usually does not exceed two or three. When the number of resources used by a product is bounded, our approximate policy provides a constant-factor performance guarantee. To our knowledge, our approximate policy is the first approximation algorithm that can handle arbitrary nonstationarities in the arrivals of the product request. Also, we note that our performance guarantee is independent of the numbers of resources and products, and it does not involve any hidden constants that can potentially depend on other input data.

The idea behind our approximate policy is to use value function approximations that are expressed as linear combinations of basis functions. The coefficients in the linear combinations are computed through a backward recursion over the time periods in the selling horizon. The approach that we use to construct our approximate policy provides flexibility on two important dimensions. First, our value function approximations are a member of a relatively broad class. In our value function approximations, we have one basis function for each product. The basis function associated with each product takes the value zero when we do not have sufficient capacities to serve a request for the product. Therefore, we refer to our basis functions as availability-tracking basis functions. For any choice of availability-tracking basis functions, we can use our backward recursion over the time periods to compute coefficients for the basis functions in the linear combinations. In our backward recursion, we have a tuning parameter  $\theta$  whose specific allowable values are determined by the availability-tracking basis functions that we use. We prove that if we construct an approximate policy using the value functions computed through our backward recursion, then the approximate policy is guaranteed to obtain at least  $1/(1 + \theta L)$  of the optimal total expected revenue. This result holds for any choice of availability-tracking basis functions. The performance guarantee of  $1/(1 + \theta L)$

improves as  $\theta$  gets smaller. In our approach, the tuning parameter  $\theta$  must be at least one, and there exist availability-tracking basis functions that permit choosing the smallest possible value of one for the tuning parameter  $\theta$ , in which case, we obtain the performance guarantee of  $1/(1 + L)$  discussed in the previous paragraph.

Second, we start with a network revenue management setup where each customer arrives into the system to purchase a particular product and each product uses at most one unit of a resource. This setup allows us to convey the key ideas without notational clutter, but we can extend our approach to more general setups. In particular, we show that we can extend our approach to a case in which we offer a set of products to each customer, and the customer chooses among the offered products. Similarly, we show that we can extend our approach to a case in which a product may use more than one unit of a resource, which occurs, for example, in airlines when group reservations are allowed. If the customers choose among the offered products, then the performance guarantee of our approximate policy is still  $1/(1 + L)$ , whereas if a product uses at most  $M$  units of a particular resource, then the performance guarantee of our approximate policy is  $1/(1 + (2M - 1)L)$ . Lastly, our backward recursion is simple and does not require solving any involved optimization problems, but it can use the solution to a linear programming approximation to construct our value function approximations, while retaining the performance guarantee of  $1/(1 + L)$ .

### 1.2. Literature Review

One approach to construct policies in network revenue management problems is based on bid prices, where we associate a bid price for each resource, measuring the value of a unit of resource. In this case, we accept the request for a product if the revenue from the product exceeds the total value of the resources consumed by the product. Simpson (1989) and Williamson (1992) compute bid prices using a linear programming approximation that is constructed under the assumption that the numbers of product requests take on their expected values. They use the optimal values of the dual variables associated with certain capacity constraints to measure the value of a unit of resource. Talluri and van Ryzin (1998) show that such a bid price policy is asymptotically optimal, as the expected numbers of product requests and the capacities of the resources increase linearly with the same rate. Talluri and van Ryzin (1999) use samples of the product requests in the linear program to capture information about the distributions of the product requests. Bertsimas and Popescu (2003) measure the value of a unit of resource directly by perturbing the right side of a capacity constraint and resolving

the linear program. A number of papers characterize the loss in the optimal total expected revenue for the policies derived from linear programming approximations; see Cooper (2002), Maglaras and Meissner (2006), Jasin and Kumar (2012), and Jasin and Kumar (2013). These papers focus on an asymptotic regime, where the total expected demand and the capacities of the resources increase linearly with the same rate. The characterization of the losses are additive and it depends on the input data. In contrast, our performance guarantee holds without an asymptotic regime. Also, our performance guarantee of  $1/(1+L)$  is multiplicative and does not depend on the input data other than  $L$ .

The value of a unit of resource should depend on the time left in the selling horizon to utilize the resource, as well as the remaining capacity of the resource. Thus, bid prices should, in principle, be time and capacity dependent. There is work on computing such bid prices. Adelman (2007), Zhang and Adelman (2009), Kunnumkal and Topaloglu (2010a), and Kirshner and Nediak (2015) develop methods that yield time-dependent bid prices, whereas Cooper and Homem-de-Mello (2007), Topaloglu (2009), and Zhang (2011) develop methods that yield capacity-dependent bid prices. Tong and Topaloglu (2013), Vossen and Zhang (2015a,b), and Kunnumkal and Talluri (2016a) show that some of these approaches are equivalent, although their derivations use seemingly unrelated paths. There is also work on incorporating customer choice behavior into network revenue management problems, where customers choose among the offered products. Some of the work extends and analyzes the linear programming approximation to incorporate customer choice behavior; see Gallego et al. (2004), Liu and van Ryzin (2008), Kunnumkal and Topaloglu (2008), Bront et al. (2009), Mendez-Diaz et al. (2010), Meissner et al. (2012), Talluri (2014), and Strauss and Talluri (2017). There is also work on approximating the value functions under customer choice behavior; see Zhang and Cooper (2005, 2009), Zhang and Adelman (2009), Kunnumkal and Topaloglu (2010b), and Kunnumkal and Talluri (2016b). There are papers that use stochastic approximation to compute booking limits and bid prices; see van Ryzin and Vulcano (2008a, b), Topaloglu (2008), and Chaneton and Vulcano (2011). These papers do not give performance guarantees.

Online packing problems are closely related to network revenue management problems, as each packing constraint may capture the capacity of a resource and each arriving item may capture the request for a product. Working with the random-order arrival model for the product requests, Kesselheim et al. (2014) give a policy with a competitive ratio of  $1 - O(\sqrt{(\log L)/c_{\min}})$ , where  $c_{\min}$  is the smallest

resource capacity. The authors also give an upper bound of  $1 - \Omega(1 - 1/L^{1/(c_{\min}-1)})$  on the competitive ratio. Thus, their competitive ratio approaches one as the capacities of the resources become large. The strong competitive ratio in Kesselheim et al. (2014) comes at the cost of having to work with the random-order arrival model, which is restrictive for the network revenue management setting. In the random-order arrival model, the set of product requests that are to arrive over the selling horizon is fixed a priori, possibly by an adversary. Once the set of product requests that are to arrive is fixed, these product requests arrive according to a random permutation. A competitive ratio under the random-order arrival model implies the same competitive ratio under independent and identically distributed arrivals. However, the competitive ratio in Kesselheim et al. (2014) does not hold under nonstationary arrivals. Indeed, in Appendix E in the e-companion, we give a simple example to show that the policy proposed by Kesselheim et al. (2014) performs arbitrarily poorly under nonstationary arrivals. Online packing problems are often motivated by the adwords setting, where assuming stationary arrivals is reasonable, because the users doing a search today may not be too different from those doing a search tomorrow. However, assuming stationary arrivals is not reasonable in the network revenue management setting, because the customers booking airline tickets close to the departure time are clearly different from those booking early. We also tested the policy in Kesselheim et al. (2014) in our computational experiments under both stationary and nonstationary arrivals, though it does not have a competitive ratio in the latter case. The policy in Kesselheim et al. (2014) noticeably lags behind our approximate policy in both cases. Much of the other work on online packing problems also uses the random-order arrival model; see Devanur and Hayes (2009), Molinaro and Ravi (2014), Agrawal et al. (2014). Devanur et al. (2011) work with stationary or adversarial arrivals, but with restrictions on choices of the adversary. Tan and Srikant (2012) use nonstationary arrivals, but allow violations of capacity constraints.

The adwords problem is a special case of the online packing problem, where the revenue and resource consumption of a product request are both dictated by the bid that an advertiser places for a keyword. A number of papers use the primal-dual framework to design algorithms for the adwords problem; see Mehta et al. (2007), Buchbinder and Naor (2009), and Goel et al. (2010). In the primal-dual framework, a tradeoff function is used to prioritize the bidders for receiving keywords. For example, the tradeoff function may prioritize the bidders with the larger remaining budgets for receiving keywords. After making decisions over the selling horizon by following the

tradeoff function, one uses the decisions to construct a primal-dual solution pair for a linear program that provides an upper bound on the optimal total expected revenue, in which case, studying the optimality gap of this solution yields the competitive ratio.

Motivated by product assortment personalization, Golrezaei et al. (2014) study a problem involving multiple products with limited inventories. The firm offers a set of products to an arriving customer. The customer makes a choice among the offered products according to a choice model. The goal is to find a policy to determine which set of products to offer to each customer to maximize the total expected revenue over the selling horizon. If each product uses exactly one resource so that  $L = 1$ , then our network revenue management problem under the customer choice behavior corresponds to the problem studied by Golrezaei et al. (2014). The authors use the primal-dual framework to construct a policy with a 50% competitive ratio. The tradeoff function used by Golrezaei et al. (2014) adjusts the revenue of each product as a function of its remaining inventory, prioritizing the products with larger inventories. Their policy offers a set of products to maximize the immediate expected adjusted-revenue from each customer. Our availability-tracking basis functions are somewhat analogous to the tradeoff functions in the primal-dual framework in the sense that our basis functions also adjust the revenue of each product as a function of the remaining inventories of the resources, in which case, we accept a product request only if the immediate adjusted-revenue from the product is nonnegative. However, our basis functions, by themselves, are not sufficient to obtain a performance guarantee. To obtain our performance guarantee, we need to weigh our basis functions with coefficients and the computation of the coefficients requires designing a backward recursion over the time periods.

Recently, Rusmevichientong et al. (2020) consider dynamic assortment problems with random usage durations. In their problem setting, each customer uses a product for a random duration of time, before returning it. The authors use linear value function approximations to obtain a policy with a 50% performance guarantee. The slopes of the linear value function approximations are computed by using a backward recursion that resembles ours on the surface, but both papers tackle rather different challenges. In Rusmevichientong et al. (2020), there is no resource network. If a customer chooses a product, the customer consumes only the inventory of this product. In the absence of a resource network, linear value function approximations are sufficient to get a performance guarantee, but under a resource network, linear approximations were not sufficient

for us. When we use nonlinear basis functions, the opportunity cost of the resource capacities used by a product becomes dependent on the system's state. Furthermore, since different products may use the same resource, it is difficult to account for the opportunity cost of each resource separately. We get around these difficulties by constructing a state-independent upper bound on the opportunity cost and constructing a backward recursion over the time periods that accounts for the optimal total expected revenue from a particular product. Nevertheless, we cannot claim that our work is a generalization of Rusmevichientong et al. (2020) either, because incorporating random usage durations in the presence of a resource network brings its own complications.

There is work in dynamic optimization with implications on revenue management. Chan and Farias (2009) give approximation algorithms for a certain class of stochastic control problems. Their work implies that if each product uses only one resource, then a myopic policy provides a 50% performance guarantee even under customer choice behavior. Asadpour and Nazerzadeh (2016) give algorithms for maximizing random monotone submodular functions. Their nonadaptive algorithm implies that if there is a single resource, then a static algorithm that *a priori* chooses the product requests to accept provides a performance guarantee of  $(e - 1)/(2e)$ . Lastly, Wang et al. (2016) and Gallego et al. (2016) give algorithms with performance guarantees for dynamic resource allocation problems, but they also consider the case where each product uses one resource.

### 1.3. Organization

In Section 2, we give a dynamic programming formulation for the network revenue management problem. In Section 3, we construct our approximate policy, provide a performance guarantee, and show that this guarantee is tight. In Section 4, we give extensions of our approximate policy to the cases where the customers choose among the products and a product may consume more than one unit of a resource. We also discuss how to leverage a linear programming approximation when constructing our approximate policies. In Section 5, we provide computational experiments. In Section 6, we conclude.

## 2. Problem Formulation

The set of resources is  $\mathcal{L}$  and the set of products is  $\mathcal{J}$ . The capacity of resource  $i$  is  $C_i$ . If we accept a request for product  $j$ , then we consume one unit of capacity of each resource in the set  $A^j \subseteq \mathcal{L}$ . We use  $L$  to denote the maximum number of resources that can be used by a product, so  $L = \max_{j \in \mathcal{J}} |A^j|$ . Accepting a request for product  $j$  generates a revenue of  $r_j$ . We have  $T$  time periods in the selling horizon indexed by  $\mathcal{T} = \{1, \dots, T\}$ .

Each time period corresponds to a sufficiently small interval of time so that there is at most one product request at each time period. We get a request for product  $j$  at time period  $t$  with probability  $\lambda_j^t$ . With the remaining probability  $1 - \sum_{j \in \mathcal{J}} \lambda_j^t$ , there is no request for a product at time period  $t$ .

Our goal is to find a policy to determine which product requests to accept at each time period to maximize the total expected revenue over the selling horizon, while adhering to the capacity availabilities of the resources. To capture the system's state, we let  $x_i$  be the remaining capacity of resource  $i$  at the beginning of a generic time period. Therefore, we can use the vector  $\mathbf{x} = (x_i : i \in \mathcal{L}) \in \mathbb{Z}_+^{|\mathcal{L}|}$  to capture the state of the resources. The set of possible states is  $\mathcal{Q} = \{\mathbf{x} \in \mathbb{Z}_+^{|\mathcal{L}|} : x_i \leq C_i \forall i \in \mathcal{L}\}$ . We can accept a request for product  $j$  when we have at least one unit of capacity for each resource used by product  $j$ . Therefore, letting  $\mathbf{1}_{\{\cdot\}}$  be the indicator function, we can accept a request for product  $j$  if and only if  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} = 1$ .

We use  $V^t(\mathbf{x})$  to denote the maximum total expected revenue over time periods  $t, t+1, \dots, T$ , given that the system is in state  $\mathbf{x}$  at time period  $t$ . Letting  $e_i \in \mathbb{Z}_+^{|\mathcal{L}|}$  be the unit vector with a one in the  $i$ th component and defining  $[a]^+ = \max\{a, 0\}$ , we can find the optimal policy by computing the optimal value functions  $\{V^t(\mathbf{x}) : \mathbf{x} \in \mathcal{Q}, t \in \mathcal{T}\}$  through the dynamic program

$$\begin{aligned} V^t(\mathbf{x}) &= \sum_{j \in \mathcal{J}} \lambda_j^t \left( \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} \right) \\ &\quad \cdot \max \left\{ r_j + V^{t+1} \left( \mathbf{x} - \sum_{i \in A^j} e_i \right) V^{t+1}(\mathbf{x}) \right\} \\ &\quad + \left( 1 - \sum_{j \in \mathcal{J}} \lambda_j^t + \sum_{j \in \mathcal{J}} \lambda_j^t \left( 1 - \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} \right) \right) V^{t+1}(\mathbf{x}) \\ &= V^{t+1}(\mathbf{x}) + \sum_{j \in \mathcal{J}} \lambda_j^t \left( \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} \right) \\ &\quad \cdot \left[ r_j - V^{t+1}(\mathbf{x}) + V^{t+1} \left( \mathbf{x} - \sum_{i \in A^j} e_i \right) \right]^+, \end{aligned} \quad (1)$$

with the boundary condition that  $V^{T+1} = 0$ . In this dynamic program, if we have a request for product  $j$  at time period  $t$  and we have capacities on all resources used by product  $j$ , then we have a choice to accept or reject the request. If we accept, then we generate a revenue of  $r_j$ , and the state of the resources at the next time period is  $\mathbf{x} - \sum_{i \in A^j} e_i$ . If we reject, then we do not generate revenue and the state of the resources at the next time period remains at  $\mathbf{x}$ . Also, if there is no request at time period  $t$  or there is a request for some

product, but we do not have capacity on some resource used by the product, then we do not accept a request, in which case, the state of the resources at the next time period remains at  $\mathbf{x}$ . The second equality in Equation (1) follows by arranging the terms. Letting  $\mathbf{C} = (C_i : i \in \mathcal{L})$  be the vector of initial resource capacities, the optimal total expected revenue is  $V^1(\mathbf{C})$ . By Equation (1), given that the state of the resources at time period  $t$  is  $\mathbf{x}$ , if  $r_j \geq V^{t+1}(\mathbf{x}) - V^{t+1}(\mathbf{x} - \sum_{i \in A^j} e_i)$ , then it is optimal to accept a request for product  $j$  as long as we have capacity on all resources used by product  $j$ . Here, we can view  $V^{t+1}(\mathbf{x}) - V^{t+1}(\mathbf{x} - \sum_{i \in A^j} e_i)$  as the opportunity cost of the capacities used by product  $j$ .

The size of the state space is  $|\mathcal{Q}| = O(\prod_{i \in \mathcal{L}} C_i)$ , which increases exponentially with the number of resources, making the computation of the optimal value functions intractable.

### 3. Approximate Policy

We construct approximations to the optimal value functions using a linear combination of basis functions. We use the following outline. In Section 3.1, we describe our basis functions. In Section 3.2, we show how to compute the coefficient of each basis function in the value function approximations. In Section 3.3, we show the performance guarantee for our approximate policy. In Section 3.4, we show that this performance guarantee is tight.

#### 3.1. Requirements for Basis Functions

We approximate the optimal value function  $V^t$  using the value function approximation  $H^t$ . To construct the value function approximation  $H^t$ , we use a linear combination of basis functions. In particular, for each product  $j$ , we have a basis function  $\varphi_j : \mathcal{Q} \rightarrow [0, 1]$ . In this case, the value function approximation  $H^t$  is given by

$$H^t(\mathbf{x}) = \sum_{j \in \mathcal{J}} \gamma_j^t \varphi_j(\mathbf{x}), \quad (2)$$

where  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is a prespecified collection of basis functions and  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  are adjustable coefficients. To construct our approximate policy, we replace  $V^{t+1}$  on the right side of Equation (1) with  $H^{t+1}$ . Therefore, given that the state of the resources at time period  $t$  is  $\mathbf{x}$ , our approximate policy accepts a request for product  $j$  as long as  $r_j \geq H^{t+1}(\mathbf{x}) - H^{t+1}(\mathbf{x} - \sum_{i \in A^j} e_i)$  and we have capacity on all resources used by product  $j$ . We refer to our basis functions as *availability-tracking basis functions* because we will impose the condition that  $\varphi_j(\mathbf{x})$  takes the value of zero if the vector of resource capacities  $\mathbf{x}$  does not provide enough availability to serve a request for product  $j$ . Following is the full definition of availability-tracking basis functions.

**Definition 1** (Availability-Tracking Bases). The collection  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is called a collection of availability-tracking basis functions if it satisfies the following conditions:

a. Availability tracking: For each  $j \in \mathcal{J}$  and  $\mathbf{x} \in \mathcal{D}$ , we have  $\varphi_j(\mathbf{x}) = 0$  whenever  $x_i = 0$  for some  $i \in A^j$ . That is,  $\varphi_j(\mathbf{x}) \leq \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}}$ .

b. Limited dependence: For each  $j \in \mathcal{J}$  and  $\mathbf{x} \in \mathcal{D}$ , the basis function  $\varphi_j(\mathbf{x})$  only depends on  $(x_i : i \in A^j)$ . That is, for any  $\mathbf{x}, \mathbf{y} \in \mathcal{D}$  with  $x_i = y_i$  for all  $i \in A^j$ , we have  $\varphi_j(\mathbf{x}) = \varphi_j(\mathbf{y})$ .

c. Normalization: For each  $j \in \mathcal{J}$ , we have  $\varphi_j(\mathbf{C}) = 1$ .

As discussed right before the definition, the range of  $\varphi_j$  is the interval  $[0, 1]$ . The first property ensures that  $\varphi_j(\mathbf{x})$  takes the value zero if the resource capacities  $\mathbf{x}$  are not sufficient to serve a request for product  $j$ . The second property ensures that  $\varphi_j(\mathbf{x})$  is independent of the capacities of the resources that are not used by product  $j$ . The third property ensures that  $\varphi_j(\mathbf{x})$  is one when the resource capacities  $\mathbf{x}$  are at their largest possible values. For example, the minimum basis function  $\varphi_j(\mathbf{x}) = \min_{i \in A^j} \frac{x_i}{C_i}$  and the polynomial basis function  $\varphi_j(\mathbf{x}) = \prod_{i \in A^j} \frac{x_i}{C_i}$  satisfy the three properties in the Definition 1. One can check that if  $\varphi_j(\mathbf{x})$  is linear in  $\mathbf{x}$ , then it cannot satisfy the three properties. Therefore, linear basis functions are not availability tracking.

Given a collection  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  of availability-tracking basis functions, we define the *maximum scaled incremental contribution*  $\Delta_{\mathcal{B}}$  from a unit of resource as

$$\Delta_{\mathcal{B}} = \max_{j \in \mathcal{J}, i \in \mathcal{L}} \max_{\mathbf{x} \in \mathcal{D} : x_i \geq 1} C_i \times \left( \varphi_j(\mathbf{x}) - \varphi_j(\mathbf{x} - e_i) \right).$$

A collection with a smaller value of  $\Delta_{\mathcal{B}}$  will yield an approximate policy with a better performance guarantee, but  $\Delta_{\mathcal{B}}$  can never be smaller than one, as shown in the next lemma.

**Lemma 1** (Bound on Change in the Value of Bases). If  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is a collection of availability-tracking basis functions, then we must have  $\Delta_{\mathcal{B}} \geq 1$ .

**Proof.** For any product  $j$  and resource  $i \in A^j$ , by parts (a) and (c) of Definition 1, we have  $\varphi_j(\mathbf{C} - C_i e_i) = 0$  and  $\varphi_j(\mathbf{C}) = 1$ . Therefore, a telescoping sum yields

$$\begin{aligned} 1 = \varphi_j(\mathbf{C}) - \varphi_j(\mathbf{C} - C_i e_i) &= \sum_{h=1}^{C_i} \left[ \varphi_j \left( \sum_{s \neq i} C_s e_s + h e_i \right) \right. \\ &\quad \left. - \varphi_j \left( \sum_{s \neq i} C_s e_s + (h-1) e_i \right) \right] \leq \sum_{h=1}^{C_i} \frac{\Delta_{\mathcal{B}}}{C_i} = \Delta_{\mathcal{B}}, \end{aligned}$$

where the inequality holds because  $\Delta_{\mathcal{B}} \geq C_i \times (\varphi_j(\sum_{s \neq i} C_s e_s + h e_i) - \varphi_j(\sum_{s \neq i} C_s e_s + (h-1) e_i))$  by the definition of  $\Delta_{\mathcal{B}}$ .  $\square$

As shown in the next two examples, the lower bound in Lemma 1 is tight.

**Example 1** (Minimum). Let  $\varphi_j(\mathbf{x}) = \min_{i \in A^j} \frac{x_i}{C_i}$ . It is simple to verify that  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  satisfies the three properties in Definition 1, so  $\Delta_{\mathcal{B}} \geq 1$  by Lemma 1. Moreover,

$$\begin{aligned} \varphi_j(\mathbf{x}) - \varphi_j(\mathbf{x} - e_i) &= \min_{\ell \in A^j} \left\{ \frac{x_\ell}{C_\ell} \right\} - \min_{\ell \in A^j} \left\{ \frac{x_\ell - \mathbf{1}_{\{\ell=i\}}}{C_\ell} \right\} \\ &= \begin{cases} \frac{1}{C_i} & \text{if } \frac{x_i}{C_i} = \min_{\ell \in A^j} \frac{x_\ell}{C_\ell} \\ \min_{\ell \in A^j} \frac{x_\ell}{C_\ell} - \frac{x_i-1}{C_i} & \text{if } \frac{x_i}{C_i} > \min_{\ell \in A^j} \frac{x_\ell}{C_\ell} \geq \frac{x_i-1}{C_i} \\ 0 & \text{if } \frac{x_i-1}{C_i} > \min_{\ell \in A^j} \frac{x_\ell}{C_\ell}. \end{cases} \end{aligned}$$

The quantity in all three cases on the right side of the equation is no larger than  $1/C_i$ , so that  $C_i(\varphi_j(\mathbf{x}) - \varphi_j(\mathbf{x} - e_i)) \leq 1$ , in which case, we get  $\Delta_{\mathcal{B}} \leq 1$ . Thus, we must have  $\Delta_{\mathcal{B}} = 1$ .

**Example 2** (Polynomial). Let  $\varphi_j(\mathbf{x}) = \prod_{i \in A^j} \frac{x_i}{C_i}$ . It is simple to verify that  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  satisfies the three properties in Definition 1, so  $\Delta_{\mathcal{B}} \geq 1$  by Lemma 1. Also,

$$\begin{aligned} \varphi_j(\mathbf{x}) - \varphi_j(\mathbf{x} - e_i) &= \left( \prod_{\ell \in A^j} \frac{x_\ell}{C_\ell} \right) - \left( \frac{x_i-1}{C_i} \cdot \prod_{\ell \in A^j \setminus \{i\}} \frac{x_\ell}{C_\ell} \right) \\ &= \frac{1}{C_i} \prod_{\ell \in A^j \setminus \{i\}} \frac{x_\ell}{C_\ell} \leq \frac{1}{C_i}, \end{aligned}$$

where the last inequality follows because  $x_\ell \leq C_\ell$ . Therefore, we have  $C_i(\varphi_j(\mathbf{x}) - \varphi_j(\mathbf{x} - e_i)) \leq 1$ , indicating that  $\Delta_{\mathcal{B}} \leq 1$ . So, we get  $\Delta_{\mathcal{B}} = 1$ .

By the discussion so far, if we set  $\varphi_j(\mathbf{x}) = \min_{i \in A^j} \frac{x_i}{C_i}$  or  $\varphi_j(\mathbf{x}) = \prod_{i \in A^j} \frac{x_i}{C_i}$  for all  $j \in \mathcal{J}$ , then the collection of basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is availability tracking with  $\Delta_{\mathcal{B}} = 1$ . There are many other availability-tracking basis functions. In Appendix A in the e-companion, we give three results that allow us to construct rich collections of availability-tracking basis functions and guide us to choose a collection of availability-tracking basis functions. First, we show that if  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is a collection of availability-tracking basis functions and  $f : [0, 1] \rightarrow [0, 1]$  is a nondecreasing and differentiable function with  $f(0) = 0$  and  $f(1) = 1$ , then letting  $\rho_j(\mathbf{x}) = f(\varphi_j(\mathbf{x}))$  for all  $j \in \mathcal{J}$ , the collection of basis functions  $\mathcal{D} = \{\rho_j : j \in \mathcal{J}\}$  is also availability tracking and  $\Delta_{\mathcal{D}} \leq \Delta_{\mathcal{B}} \max_{a \in [0, 1]} f'(a)$ . For example, using this result with  $\varphi_j(\mathbf{x}) = \prod_{i \in A^j} \frac{x_i}{C_i}$  and  $f(a) = (1 - e^{-a}) / (1 - e^{-1})$ , it follows that the collection of basis functions  $\mathcal{D} = \{\rho_j : j \in \mathcal{J}\}$  with

$$\rho_j(\mathbf{x}) = f(\varphi_j(\mathbf{x})) = \frac{1 - e^{-\prod_{i \in A^j} x_i / C_i}}{1 - e^{-1}}$$

is availability tracking. It might be difficult to see at a first glance that the previous basis function is

availability tracking. Second, we show that if, for each  $i \in \mathcal{L}$ ,  $g_i : [0, 1] \rightarrow [0, 1]$  is a nondecreasing and differentiable function with  $g_i(0) = 0$  and  $g_i(1) = 1$ , then letting  $\rho_j(x) = \min_{i \in A^j} g_i(x_i/C_i)$  or  $\rho_j(x) = \prod_{i \in A^j} g_i(x_i/C_i)$  for all  $j \in \mathcal{J}$ , the collection of basis functions  $\mathcal{D} = \{\rho_j : j \in \mathcal{J}\}$  is also availability-tracking and  $\Delta_{\mathcal{D}} \leq \max_{i \in \mathcal{L}} \max_{a \in [0, 1]} g_i'(a)$ . For example, using this result with  $g_i(a) = (1 - e^{-a})/(1 - e^{-1})$ , it follows that the collection of basis functions  $\{\rho_j : j \in \mathcal{J}\}$  with  $\rho_j(x) = \min_{i \in A^j} (1 - e^{-x_i/C_i})/(1 - e^{-1})$  or  $\rho_j(x) = \prod_{i \in A^j} (1 - e^{-x_i/C_i})/(1 - e^{-1})$  is availability tracking. By using these two results, we can construct rich collections of availability-tracking basis functions. Third, Definition 1 does not require the availability-tracking basis functions to be nondecreasing, but we show that if  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is any collection of availability-tracking basis functions, then we can always construct a collection of nondecreasing availability-tracking basis functions  $\mathcal{D} = \{\rho_j : j \in \mathcal{J}\}$  such that  $\Delta_{\mathcal{D}} \leq \Delta_{\mathcal{B}}$ . Since a collection of basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  with a smaller value of  $\Delta_{\mathcal{B}}$  will yield an approximate policy with a better performance guarantee, this result indicates that it is preferable to use nondecreasing basis functions. In Appendix A in the e-companion, we also give an example of a nonmonotone basis function.

The discussion so far focuses on the basis functions  $\{\varphi_j : j \in \mathcal{J}\}$  in the value function approximations in Equation (2). Next, we discuss computing the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ .

### 3.2. Tuning the Coefficients

To compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  in the value function approximations in Equation (2), we use the following backward recursion over the time periods:

- Initialization: Let  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  be any collection of availability-tracking basis functions and  $\theta \geq \Delta_{\mathcal{B}}$  be a tuning parameter. Initialize  $\gamma_j^{T+1} = 0$  for all  $j \in \mathcal{J}$ .
- Coefficient computation: For each  $t = T, T-1, \dots, 1$ , use the coefficients  $\{\gamma_j^{t+1} : j \in \mathcal{J}\}$  to compute  $\{\gamma_j^t : j \in \mathcal{J}\}$  as

$$\gamma_j^t = \lambda_j^t \left[ r_j - \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right]^+ + \gamma_j^{t+1}. \quad (3)$$

We shortly provide an intuitive interpretation for the presented algorithm. The algorithm allows us to compute  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ , which specifies the value function approximations  $\{H^t : t \in \mathcal{T}\}$  given in Equation (2). In our approximate policy, we make the decisions for the product request at time period  $t$  by replacing the optimal value function  $V^{t+1}$  on the right side of Equation (1) with  $H^{t+1}$ . Thus, given that the state of the system at time period  $t$  is  $x$ , if  $r_j \geq H^{t+1}(x) - H^{t+1}(x - \sum_{i \in A^j} e_i)$ ,

then our approximate policy accepts a request for product  $j$  as long as we have capacity on all resources used by product  $j$ . Next, we give an intuitive interpretation for the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  and the backward recursion in Equation (3) that we use to compute these coefficients.

To give an intuitive interpretation for the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ , we note that the optimal total expected revenue over the selling horizon is given by  $V^1(C)$ . Therefore,  $H^1(C)$  is our approximation to the optimal total expected revenue over the selling horizon. Since  $\varphi_j(C) = 1$  for all  $j \in \mathcal{J}$  by part (c) of Definition 1, we have  $H^1(C) = \sum_{j \in \mathcal{J}} \gamma_j^1 \varphi_j(C) = \sum_{j \in \mathcal{J}} \gamma_j^1$ , which implies that  $\sum_{j \in \mathcal{J}} \gamma_j^1$  is our approximation to the optimal total expected revenue over the selling horizon. Thus, we can interpret  $\gamma_j^1$  as our approximation to the optimal total expected revenue that we can extract from the requests for product  $j$  over the selling horizon. In this case, extending the argument intuitively to a generic time period, we will also interpret  $\gamma_j^t$  as our approximation to the optimal total expected revenue that we can extract from the requests for product  $j$  over time periods  $t, t+1, \dots, T$ . Since  $\gamma_j^t$  is our approximation to the optimal total expected revenue that we can extract from the requests for product  $j$  over time periods  $t, t+1, \dots, T$ , whereas  $\gamma_j^{t+1}$  is our approximation to the optimal total expected revenue that we can extract from the requests for product  $j$  over time periods  $t+1, t+2, \dots, T$ , we naturally expect that  $\gamma_j^t \geq \gamma_j^{t+1}$ . Noting that  $[a]^+ \geq 0$  for any  $a \in \mathbb{R}$ , by Equation (3), we indeed have  $\gamma_j^t \geq \gamma_j^{t+1}$ .

To understand the intuitive idea behind the backward recursion in Equation (3), recall that we view  $V^{t+1}(x) - V^{t+1}(x - \sum_{i \in A^j} e_i)$  as the opportunity cost of the capacities that are used by a request for product  $j$  that we accept at time period  $t$ . Using our value function approximations, we can approximate this opportunity cost by  $H^{t+1}(x) - H^{t+1}(x - \sum_{i \in A^j} e_i)$ . Considering the value function approximation  $H^{t+1}(x) = \sum_{j \in \mathcal{J}} \gamma_j^{t+1} \varphi_j(x)$  with any collection of availability-tracking basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$ , it is possible to show that we can upper bound the opportunity cost  $H^{t+1}(x) - H^{t+1}(x - \sum_{i \in A^j} e_i)$  as  $H^{t+1}(x) - H^{t+1}(x - \sum_{i \in A^j} e_i) \leq \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}$ , as long as  $\theta \geq \Delta_{\mathcal{B}}$ . We show this result in Lemma 3 when we analyze the performance of our approximate policy. Note that the upper bound  $\theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}$  does not depend on the system state. Thus, we use  $\theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}$  as a state-independent approximation to the opportunity cost of the capacities that are used by a request for product  $j$  that we accept at time period  $t$ . In this case, going back to Equation (3), if we accept a request for product  $j$  at time period  $t$ , then we obtain a revenue of  $r_j$  and we approximate the opportunity cost of the capacities

that we consume by using  $\theta \sum_{i \in A^t} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}$ . We accept a request for product  $j$  only if the revenue from the product exceeds the opportunity cost of the capacities consumed by product  $j$ . Thus,  $[r_j - \theta \sum_{i \in A^t} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}]^+$  is our approximation to the net revenue contribution from a request for product  $j$  at time period  $t$ . By the discussion in the previous paragraph,  $\gamma_j^t$  is an approximation to the optimal total expected revenue from product  $j$  over time periods  $t, t+1, \dots, T$ . Since we have a request for product  $j$  at time period  $t$  with probability  $\lambda_j^t$ , the recursion in Equation (3) states the following. The approximation to the optimal total expected revenue from product  $j$  over time periods  $t, t+1, \dots, T$  is given by the approximation to the expected net revenue contribution from product  $j$  at time period  $t$  plus the approximation to the optimal total expected revenue from product  $j$  over time periods  $t+1, t+2, \dots, T$ .

In the next section, we show that our approximate policy is guaranteed to obtain at least  $1/(1 + \theta L)$  fraction of the optimal total expected revenue. The proof of this performance guarantee involves two steps. In the first step, we use a linear program whose objective value provides an upper bound on the optimal total expected revenue. In particular, assuming that the total number of requests for each product takes on its expected value, this linear program finds the number of requests to accept for each product. The total expected number of requests for product  $j$  over the selling horizon is  $\sum_{t \in \mathcal{T}} \lambda_j^t$ . Letting  $\Lambda_j = \sum_{t \in \mathcal{T}} \lambda_j^t$  for notational brevity and using the decision variable  $z_j$  to capture the number of requests for product  $j$  that we plan to accept over the selling horizon, the linear program has the form

$$\begin{aligned} Z_{LP}^* = \max \left\{ \sum_{j \in \mathcal{J}} r_j z_j : \sum_{j \in \mathcal{J}} \mathbf{1}_{\{i \in A^j\}} z_j \leq C_i \quad \forall i \in \mathcal{L}, \right. \\ \left. 0 \leq z_j \leq \Lambda_j \quad \forall j \in \mathcal{J} \right\}. \end{aligned} \quad (4)$$

The objective function accounts for the total revenue over the selling horizon. The first constraint ensures that the total capacity consumptions of the resources do not exceed the initial capacities. The second constraint ensures that the numbers of requests that we serve for the products do not exceed the expected demands. We show that we can use the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  that are computed through Equation (3) to construct a feasible solution to the dual of the previous linear program and this feasible solution provides an objective value of at most  $(1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1$ . Since the dual is a minimization problem, we get  $(1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1 \geq Z_{LP}^* \geq V^1(C)$ , where the second inequality uses the fact that the optimal objective value of the linear program is an upper bound on the optimal total expected revenue.

In the second step, we use backward induction over the time periods to show that the total expected revenue obtained by our approximate policy is at least  $\sum_{j \in \mathcal{J}} \gamma_j^1$ . The intuition behind this result is that the recursion in Equation (3) uses the upper bound  $\theta \sum_{i \in A^t} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1}$  on the opportunity cost of the capacities used by each product  $j$ . So, intuitively speaking,  $\gamma_j^1$  computed through the recursion in Equation (3) turns out to be a pessimistic approximation to the optimal total expected revenue obtained from product  $j$  over the selling horizon. In this case, we will show that the total expected revenue obtained by our approximate policy is at least as large as the pessimistic approximation of the optimal total expected revenue  $\sum_{j \in \mathcal{J}} \gamma_j^1$ . Thus, letting APP be the total expected revenue obtained by our approximate policy, we have  $\text{APP} \geq \sum_{j \in \mathcal{J}} \gamma_j^1$ . Using the chain of inequalities at the end of the previous paragraph, we get  $\text{APP} \geq \sum_{j \in \mathcal{J}} \gamma_j^1 \geq \frac{1}{1 + \theta L} Z_{LP}^* \geq \frac{1}{1 + \theta L} V^1(C)$ , establishing that the total expected revenue obtained by our approximate policy is guaranteed to be at least  $1/(1 + \theta L)$  fraction of the optimal total expected revenue.

In the next section, we follow the two steps discussed in the previous two paragraphs to show this performance guarantee for our approximate policy.

### 3.3. Performance Guarantee for the Approximate Policy

Recall that once we compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  using Equation (3), in our approximate policy, given that the state of the resources at time period  $t$  is  $x$ , we accept a request for product  $j$  as long as  $r_j \geq H^{t+1}(x) - H^{t+1}(x - \sum_{i \in A^j} e_i)$  and we have capacity on all resources used by product  $j$ . To formally state our approximate policy, we use  $u_j^{\text{App},t} : \mathcal{Q} \rightarrow \{0, 1\}$  to denote the decision function of our approximate policy at time period  $t$ . Given that the state of the resources at time period  $t$  is  $x$ , we have  $u_j^{\text{App},t}(x) = 1$  if our approximate policy accepts a request for product  $j$  at time period  $t$ . Otherwise, we have  $u_j^{\text{App},t}(x) = 0$ . Therefore,  $u_j^{\text{App},t}(x)$  is given by

$$u_j^{\text{App},t}(x) = \begin{cases} \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} & \text{if } r_j \geq H^{t+1}(x) - H^{t+1}\left(x - \sum_{i \in A^j} e_i\right), \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

In the next theorem, we give a performance guarantee for this policy as a function of the tuning parameter  $\theta$  in Equation (3) and the maximum number of resources  $L$  used by a product.

**Theorem 1 (Performance).** *If the tuning parameter  $\theta$  satisfies  $\theta \geq \Delta_B$ , then the total expected revenue obtained by the approximate policy is at least  $1/(1 + \theta L)$  fraction of the optimal.*

We devote the rest of this section to giving a proof of Theorem 1, but before we go into the proof, we make three remarks on this theorem. First, to obtain the best performance guarantee, we need to choose the tuning parameter  $\theta$  as small as possible and the smallest possible value of  $\theta$  in the theorem is  $\Delta_{\mathcal{B}}$ . By Lemma 1, we have  $\Delta_{\mathcal{B}} \geq 1$ , but as shown in Examples 1 and 2, there are choices of basis functions under which  $\Delta_{\mathcal{B}} = 1$ . Therefore, working with these basis functions, we can choose the tuning parameter  $\theta$  as one and obtain an approximate policy whose total expected revenue is at least  $1/(1+L)$  fraction of the optimal total expected revenue. In many network revenue management problems arising in settings such as airlines and hotels, each product uses only a small number of resources. Therefore, even though the number of resources can be large, as long as the number of resources used by a product is uniformly bounded, we obtain a constant-factor approximation guarantee. Second, although we obtain the best performance guarantee by choosing  $\theta$  at its smallest possible value of  $\Delta_{\mathcal{B}}$ , our computational experiments indicate that increasing  $\theta$  beyond  $\Delta_{\mathcal{B}}$  can improve the total expected revenue of the approximate policy. We view the tuning parameter  $\theta$  as a knob that provides flexibility in the implementation of our approximate policy. Third, as is the case for most constant-factor approximation algorithms, the performance guarantee in Theorem 1 is a worst-case guarantee. In our computational experiments, we compare the total expected revenue obtained by our approximate policy with a computationally tractable upper bound on the optimal total expected revenue, and demonstrate that the approximate policy performs substantially better than its worst-case performance guarantee. Next, we turn to the proof of Theorem 1. The proof uses a sequence of lemmas.

In the next lemma, we show that  $(1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1$  is an upper bound on  $V^1(\mathbf{C})$ . The proof is based on using  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  to construct a feasible solution to the dual of problem (4).

**Lemma 2** (Upper Bound on the Optimal Total Expected Revenue). *If the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  are computed through Equation (3), then we have  $V^1(\mathbf{C}) \leq (1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1$ .*

**Proof.** It is a well-known result in the network revenue management literature that the optimal objective value of problem (4) is an upper bound on the optimal total expected revenue; see Bertsimas and Popescu (2003). Thus, we have  $Z_{LP}^* \geq V^1(\mathbf{C})$ . Problem (4) is feasible and bounded because setting  $z_j = 0$  for all  $j \in \mathcal{J}$  provides a feasible solution and the decision variables have upper bounds. Thus, the optimal objective value of the dual of problem (4) is also  $Z_{LP}^*$ . Using the vectors of dual

variables  $\mu = (\mu_i : i \in \mathcal{L})$  and  $\sigma = (\sigma_j : j \in \mathcal{J})$ , the dual of problem (4) is  $\min_{(\mu, \sigma) \in \mathbb{R}_+^{|\mathcal{L}|+|\mathcal{J}|}} \{\sum_{i \in \mathcal{L}} C_i \mu_i + \sum_{j \in \mathcal{J}} \Lambda_j \sigma_j : \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \mu_i + \sigma_j \geq r_j \quad \forall j \in \mathcal{J}\}$ . The decision variable  $\sigma_j$  has a nonnegative objective function coefficient, so it takes on its smallest possible value in an optimal solution. Writing the constraints in the dual problem as  $\sigma_j \geq r_j - \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \mu_i$  for all  $j \in \mathcal{J}$  and noting that the decision variable  $\sigma_j$  is nonnegative, the smallest possible value of  $\sigma_j$  is  $[r_j - \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \mu_i]^+$ . In this case, replacing the decision variable  $\sigma_j$  by its value in an optimal solution, we can write the dual of problem (4) as

$$Z_{LP}^* = \min_{\mu \in \mathbb{R}_+^{|\mathcal{L}|}} \left\{ \sum_{i \in \mathcal{L}} C_i \mu_i + \sum_{j \in \mathcal{J}} \Lambda_j \left[ r_j - \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \mu_i \right]^+ \right\}. \quad (6)$$

We define the solution  $\hat{\mu} = (\hat{\mu}_i : i \in \mathcal{L})$  to the problem in Equation (6) by setting  $\hat{\mu}_i = \frac{\theta}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^1$  for all  $i \in \mathcal{L}$ . Noting Equation (3), since  $[a]^+ \geq 0$  for any  $a \in \mathbb{R}$ , we have  $\gamma_j^1 \geq \gamma_j^2 \geq \dots \geq \gamma_j^T \geq \gamma_j^{T+1} = 0$  for all  $j \in \mathcal{J}$ , which implies that  $\hat{\mu}_i \geq 0$  for all  $i \in \mathcal{L}$ . Thus,  $\hat{\mu}$  is a feasible solution to problem (6). In this case, the objective value provided by this solution is at least  $Z_{LP}^*$ , yielding

$$\begin{aligned} Z_{LP}^* &\leq \sum_{i \in \mathcal{L}} C_i \hat{\mu}_i + \sum_{j \in \mathcal{J}} \Lambda_j \left[ r_j - \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \hat{\mu}_i \right]^+ \\ &= \sum_{i \in \mathcal{L}} C_i \frac{\theta}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^1 \\ &\quad + \sum_{j \in \mathcal{J}} \Lambda_j \left[ r_j - \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^j\}} \frac{\theta}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^1 \right]^+ \\ &= \theta \sum_{k \in \mathcal{J}} \sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^1 \\ &\quad + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \Lambda_j^t \left[ r_j - \theta \sum_{i \in \mathcal{L}} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^1 \right]^+ \\ &\leq \theta L \sum_{k \in \mathcal{J}} \gamma_k^1 + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \Lambda_j^t \\ &\quad \cdot \left[ r_j - \theta \sum_{i \in \mathcal{L}} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right]^+ \\ &= \theta L \sum_{k \in \mathcal{J}} \gamma_k^1 + \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} [\gamma_j^t - \gamma_j^{t+1}] \\ &= (1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1, \end{aligned}$$

where the second equality uses the fact that  $\Lambda_j = \sum_{t \in \mathcal{T}} \Lambda_j^t$ , the second inequality uses the fact that  $\sum_{i \in \mathcal{L}} \mathbb{1}_{\{i \in A^k\}} = |A^k| \leq L$  and  $\gamma_k^1 \geq \gamma_k^2 \geq \dots \geq \gamma_k^T \geq \gamma_k^{T+1}$  for all  $k \in \mathcal{J}$ , the third equality holds because  $\gamma_j^t - \gamma_j^{t+1} = [r_j - \theta \sum_{i \in \mathcal{L}} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^{t+1}]^+$  by Equation (3), and the fourth equality follows by noting that  $\sum_{t \in \mathcal{T}} [\gamma_j^t - \gamma_j^{t+1}] = \gamma_j^1$ . Thus, we have  $(1 + \theta L) \sum_{j \in \mathcal{J}} \gamma_j^1 \geq Z_{LP}^*$ , in which case, the desired result follows from the fact that  $Z_{LP}^* \geq V^1(\mathbf{C})$ .  $\square$

In the next lemma, we bound the opportunity cost of the capacities that are consumed by product  $j$  under our value function approximations.

**Lemma 3** (Bound on the Opportunity Cost). *For a collection of availability-tracking basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$ , let  $H^t(\mathbf{x}) = \sum_{k \in \mathcal{J}} \gamma_k^t \varphi_k(\mathbf{x})$ , where the coefficients  $\{\gamma_k^t : k \in \mathcal{J}\}$  satisfy  $\gamma_k^t \geq 0$  for all  $k \in \mathcal{J}$ . Then, for each  $j \in \mathcal{J}$  and  $\mathbf{x} \in \mathcal{Q}$  such that  $\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i \geq 0$ , we have*

$$H^t(\mathbf{x}) - H^t\left(\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i\right) \leq \Delta_{\mathcal{B}} \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t.$$

**Proof.** We prove the result using induction on the cardinality of  $A^j$ . Consider the base case where  $|A^j| = 1$  so that we have  $A^j = \{i\}$  for some  $i \in \mathcal{L}$ . In this case, we get

$$\begin{aligned} H^t(\mathbf{x}) - H^t(\mathbf{x} - \mathbf{e}_i) &= \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t (\varphi_k(\mathbf{x}) - \varphi_k(\mathbf{x} - \mathbf{e}_i)) \\ &\leq \frac{\Delta_{\mathcal{B}}}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t, \end{aligned} \quad (7)$$

where the equality holds because  $\varphi_k(\mathbf{x}) - \varphi_k(\mathbf{x} - \mathbf{e}_i) = 0$  whenever  $i \notin A^k$  by part (b) of Definition 1 and the inequality follows from the definition of  $\Delta_{\mathcal{B}}$ . Thus, the base case holds. Suppose that the result holds for any  $|A^j| \leq s$ . Consider a case in which  $|A^j| = s + 1$ , so  $A^j = B \cup \{\ell\}$  for some  $B \subseteq \mathcal{L}$  with  $|B| = s$  and  $\ell \in \mathcal{L}$  with  $\ell \notin B$ . Letting  $\mathbf{y} = \mathbf{x} - \sum_{i \in B} \mathbf{e}_i$ , we obtain

$$\begin{aligned} H^t(\mathbf{x}) - H^t\left(\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i\right) &= H^t(\mathbf{x}) - H^t\left(\mathbf{x} - \sum_{i \in B} \mathbf{e}_i - \mathbf{e}_{\ell}\right) \\ &= H^t(\mathbf{x}) - H^t\left(\mathbf{x} - \sum_{i \in B} \mathbf{e}_i\right) + H^t\left(\mathbf{x} - \sum_{i \in B} \mathbf{e}_i\right) \\ &\quad - H^t\left(\mathbf{x} - \sum_{i \in B} \mathbf{e}_i - \mathbf{e}_{\ell}\right) \\ &\leq \Delta_{\mathcal{B}} \sum_{i \in B} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t + H(\mathbf{y}) - H(\mathbf{y} - \mathbf{e}_{\ell}) \\ &\leq \Delta_{\mathcal{B}} \sum_{i \in B} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t + \frac{\Delta_{\mathcal{B}}}{C_{\ell}} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{\ell \in A^k\}} \gamma_k^t \\ &= \Delta_{\mathcal{B}} \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^t, \end{aligned}$$

where the first inequality follows from the induction assumption, the second inequality follows from the base case, and the last equality is by the fact that  $A^j = B \cup \{\ell\}$ .  $\square$

We obtain the term  $\mathbb{1}_{\{i \in A^k\}}$  on the right side of Equation (7) due to part (b) of Definition 1, which, in turn, yields the term  $\mathbb{1}_{\{i \in A^k\}}$  on the right side of Equation (3). Without this term, the upper bound in Lemma 2 would involve  $(1 + \theta|\mathcal{L}|)$  rather than  $(1 + \theta L)$ . Next, we use this lemma to show that the total expected revenue obtained by our approximate policy is at least  $\sum_{j \in \mathcal{J}} \gamma_j^1$ . Let  $U^t(\mathbf{x})$  be the total expected revenue obtained by our approximate policy over time periods  $t, t+1, \dots, T$  given that the state of the resources at time period  $t$  is  $\mathbf{x}$ . Noting the decision function for the approximate policy in Equation (5), we can compute  $\{U^t : t \in \mathcal{T}\}$  through the dynamic program

$$\begin{aligned} U^t(\mathbf{x}) &= \sum_{j \in \mathcal{J}} \lambda_j^t u_j^{\text{App},t}(\mathbf{x}) \left[ r_j + U^{t+1}\left(\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i\right) \right] \\ &\quad + \left( 1 - \sum_{j \in \mathcal{J}} \lambda_j^t + \sum_{j \in \mathcal{J}} \lambda_j^t (1 - u_j^{\text{App},t}(\mathbf{x})) \right) U^{t+1}(\mathbf{x}) \\ &= U^{t+1}(\mathbf{x}) + \sum_{j \in \mathcal{J}} \lambda_j^t u_j^{\text{App},t}(\mathbf{x}) \\ &\quad \cdot \left[ r_j - U^{t+1}(\mathbf{x}) + U^{t+1}\left(\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i\right) \right], \end{aligned} \quad (8)$$

with the boundary condition that  $U^{T+1} = 0$ . In the dynamic program in Equation (8), we have a request for product  $j$  at time period  $t$  with probability  $\lambda_j^t$ , in which case, if we have  $u_j^{\text{App},t}(\mathbf{x}) = 1$  so that the approximate policy accepts this request, then we obtain a revenue of  $r_j$  and the state of the resources at the next time period is  $\mathbf{x} - \sum_{i \in A^j} \mathbf{e}_i$ . If there is a request for product  $j$  at time period  $t$ , but we have  $u_j^{\text{App},t}(\mathbf{x}) = 0$ , then the state of the resources at the next time period remains at  $\mathbf{x}$ . With probability  $1 - \sum_{j \in \mathcal{J}} \lambda_j^t$ , there is no request, in which case, the state of the resources at the next time period also remains at  $\mathbf{x}$ . The second equality is by arranging the terms.

The total expected revenue obtained by our approximate policy is  $U^1(\mathbf{C})$ . In the next lemma, we show that this total expected revenue is at least  $\sum_{j \in \mathcal{J}} \gamma_j^1$ .

**Lemma 4** (Lower Bound on Performance). *If the tuning parameter  $\theta$  satisfies  $\theta \geq \Delta_{\mathcal{B}}$  and the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  are computed through Equation (3), then we have  $U^1(\mathbf{C}) \geq \sum_{j \in \mathcal{J}} \gamma_j^1$ .*

**Proof.** We use induction over the time periods to show that  $U^t(\mathbf{x}) \geq H^t(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{Q}$  and  $t \in \mathcal{T}$ , where  $H^t(\mathbf{x}) = \sum_{j \in \mathcal{J}} \gamma_j^t \varphi_j(\mathbf{x})$  with  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  computed through Equation (3). Consider the base case at time period  $T+1$ . Since  $U^{T+1} = 0$  and  $\gamma_j^{T+1} = 0$ , the base case holds. Suppose that the result holds at time

period  $t + 1$ , so  $U^{t+1}(x) \geq H^{t+1}(x)$  for all  $x \in \mathcal{Q}$ . On the right side of Equation (8), the coefficients of  $U^{t+1}(x)$  and  $U^{t+1}(x - \sum_{i \in A^j} e_i)$  are, respectively,  $1 - \sum_{j \in \mathcal{J}} \lambda_j^t$ ,  $u_j^{\text{App},t}(x)$  and  $\lambda_j^t u_j^{\text{App},t}(x)$ , which are nonnegative. Since  $U^{t+1} \geq H^{t+1}$  by the induction assumption, replacing  $U^{t+1}$  on the right side of Equation (8) with  $H^{t+1}$ , the right side of Equation (8) gets smaller. Thus, we have

$$\begin{aligned}
U^t(x) &\geq H^{t+1}(x) + \sum_{j \in \mathcal{J}} \lambda_j^t u_j^{\text{App},t}(x) \\
&\quad \cdot \left[ r_j - H^{t+1}(x) + H^{t+1} \left( x - \sum_{i \in A^j} e_i \right) \right] \\
&= H^{t+1}(x) + \sum_{j \in \mathcal{J}} \lambda_j^t \left( \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \\
&\quad \cdot \left[ r_j - H^{t+1}(x) + H^{t+1} \left( x - \sum_{i \in A^j} e_i \right) \right]^+ \\
&\geq H^{t+1}(x) + \sum_{j \in \mathcal{J}} \lambda_j^t \left( \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \\
&\quad \cdot \left[ r_j - \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbb{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right]^+ \\
&= H^{t+1}(x) + \sum_{j \in \mathcal{J}} \left( \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \left( \gamma_j^t - \gamma_j^{t+1} \right) \\
&\geq H^{t+1}(x) + \sum_{j \in \mathcal{J}} \varphi_j(x) \left( \gamma_j^t - \gamma_j^{t+1} \right) = H^t(x). \quad (9)
\end{aligned}$$

Here, the first equality follows by the definition of the decision function in Equation (5) as we have  $u_j^{\text{App},t}(x) = 1$  if and only if  $\prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} = 1$  and  $r_j - H^{t+1}(x) + H^{t+1}(x - \sum_{i \in A^j} e_i) \geq 0$ . The second inequality follows from Lemma 3 and the fact that  $\theta \geq \Delta_{\mathcal{B}}$ . The second equality uses the definition of  $\gamma_j^t$  in Equation (3). The third inequality follows from part (a) of Definition 1, along with the fact that  $\gamma_j^t - \gamma_j^{t+1} \geq 0$ . The last equality holds because we have  $H^{t+1}(x) = \sum_{j \in \mathcal{J}} \gamma_j^{t+1} \varphi_j(x)$ . So,  $U^t(x) \geq H^t(x)$ , completing the induction argument. This inequality at  $t = 1$  and  $x = C$ , along with part (c) of Definition 1, yields  $U^1(C) \geq H^1(C) = \sum_{j \in \mathcal{J}} \gamma_j^1 \varphi_j(C) = \sum_{j \in \mathcal{J}} \gamma_j^1$ .  $\square$

Note that due to part (a) of Definition 1, we can use  $\varphi_j(x)$  in the last inequality in Equation (9) as a lower bound substitute on the feasibility condition  $\prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}}$  for accepting a request for product  $j$ . The proof of Theorem 1 directly follows by combining Lemmas 2 and 4.

**Proof of Theorem 1.** The optimal total expected revenue is  $V^1(C)$ , whereas the total expected revenue obtained by our approximate policy is  $U^1(C)$ . In this case, by Lemmas 2 and 4, we have  $U^1(C) \geq \sum_{j \in \mathcal{J}} \gamma_j^1 \geq V^1(C)/(1 + \theta L)$ .  $\square$

### 3.4. Tightness of the Analysis

In this section, we give a class of problem instances to demonstrate that the performance guarantee in Theorem 1 is tight for any  $L$ . Fix two positive integers  $K$  and  $\beta$ . We consider a problem instance with  $K$  resources and  $K + 1$  products. We index the resources by  $\mathcal{L} = \{1, \dots, K\}$  and the products by  $\mathcal{J} = \{1, \dots, K, K + 1\}$ . The initial capacities of the resources are  $C_i = \beta + 1$  for all  $i \in \mathcal{L}$ . For  $j = 1, \dots, K$ , product  $j$  uses only one unit of resource  $j$ . Product  $K + 1$  uses all of the  $K$  resources. That is,  $A^j = \{j\}$  for all  $j = 1, \dots, K$  and  $A^{K+1} = \mathcal{L}$ . Note that  $L = \max_{j \in \mathcal{J}} |A^j| = K$ . The revenues of the products are given by  $r_j = \frac{1}{\beta+1} (1 - \frac{1}{\beta})$  for all  $j = 1, \dots, K$  and  $r_{K+1} = 1$ . There are  $K\beta + 1$  time periods in the selling horizon indexed by  $\mathcal{T} = \{1, \dots, K\beta + 1\}$ . For  $j = 1, \dots, K$ , the requests for product  $j$  arrive only at time periods  $(j-1)\beta + 1, (j-1)\beta + 2, \dots, j\beta$ . At the last time period, we have a request of product  $K + 1$ . In particular, for  $j = 1, \dots, K$ , we have

$$\begin{aligned}
\lambda_j^t &= \begin{cases} 1 & \text{if } (j-1)\beta + 1 \leq t \leq j\beta, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \\
\lambda_{K+1}^t &= \begin{cases} 1 & \text{if } t = K\beta + 1, \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

Even if we accept all of the product requests, we do not run out of the capacities of any of the resources. Thus, the optimal policy accepts all requests, in which case, the total expected revenue of the optimal policy is  $\text{OPT} = K\beta r_1 + r_{K+1} = \frac{K\beta}{\beta+1} (1 - \frac{1}{\beta}) + 1$ .

We consider our approximate policy using the collection of availability-tracking basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  with  $\varphi_j(x) = \min_{i \in A^j} \frac{x_i}{C_i}$  for all  $j \in \mathcal{J}$ . In Example 1, we show that we have  $\Delta_{\mathcal{B}} = 1$  for this collection of basis functions. Therefore, we can choose  $\theta = 1$  in the recursion in Equation (3). We proceed to computing the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ . Since  $\gamma_{K+1}^{K\beta+2} = 0$ , by Equation (3), we have  $\gamma_{K+1}^{K\beta+1} = \lambda_{K+1}^{K\beta+1} r_{K+1} = 1$ . Also, noting that  $\lambda_{K+1}^t = 0$  for all  $t \in \mathcal{T} \setminus \{K\beta + 1\}$ , we have  $\gamma_{K+1}^1 = \gamma_{K+1}^2 = \dots = \gamma_{K+1}^{K\beta+1} = 1$ . For  $j = 1, \dots, K$ , by Equation (3), we have

$$\begin{aligned}
\gamma_j^t &= \lambda_j^t \left[ r_j - \frac{\gamma_j^{t+1} + \gamma_{K+1}^{t+1}}{C_j} \right]^+ + \gamma_j^{t+1} \\
&= \lambda_j^t \left[ r_j - \frac{\gamma_j^{t+1} + \gamma_{K+1}^{t+1}}{\beta + 1} \right]^+ + \gamma_j^{t+1}.
\end{aligned}$$

Since  $\gamma_j^{K\beta+2} = 0$ , noting that  $\lambda_j^{K\beta+1} = 0$ , we obtain  $\gamma_j^{K\beta+1} = 0$  by the earlier recursion. Also, since  $\gamma_{K+1}^1 = \gamma_{K+1}^2 = \dots = \gamma_{K+1}^{K\beta+1} = 1$ , for all  $t \in \mathcal{T} \setminus \{K\beta + 1\}$ , we have

$$r_j - \frac{\gamma_j^{t+1} + \gamma_{K+1}^{t+1}}{\beta + 1} \leq r_j - \frac{\gamma_{K+1}^{t+1}}{\beta + 1} = \frac{1}{\beta + 1} \left( 1 - \frac{1}{\beta} \right) - \frac{1}{\beta + 1} \leq 0.$$

Therefore, we get

$$\left[ r_j - \frac{\gamma_j^{t+1} + \gamma_{K+1}^{t+1}}{\beta + 1} \right]^+ = 0$$

for all  $t \in \mathcal{T} \setminus \{K\beta + 1\}$ , in which case, noting the earlier recursion, we have  $\gamma_j^1 = \gamma_j^2 = \dots = \gamma_j^{K\beta+1} = 0$ . Thus, by the discussion in this paragraph, we have  $\gamma_{K+1}^t = 1$  for all  $t \in \mathcal{T}$ . Also, for  $j = 1, \dots, K$ , we have  $\gamma_j^t = 0$  for all  $t \in \mathcal{T}$ . In this case, we get  $H^t(x) = \sum_{j \in \mathcal{J}} \gamma_j^t \min_{i \in A^j} \frac{x_i}{C_i} = \gamma_{K+1}^t \min_{i \in \mathcal{L}} \frac{x_i}{C_i} = \frac{1}{\beta+1} \cdot \min_{i \in \mathcal{L}} x_i$  for all  $t \in \mathcal{T}$ .

Consider the decisions made by our approximate policy. Initially, the capacities of the resources are  $C_i = \beta + 1$  for all  $i \in \mathcal{L}$ . At the first time period, we have a request for product 1. Since  $r_1 = \frac{1}{\beta+1}(1 - \frac{1}{\beta}) < \frac{1}{\beta+1} = H^2(C) - H^2(C - e_1)$ , we do not accept the request. Since  $r_1 = r_2 = \dots = r_K$ , a similar argument shows that we reject all of the requests at time periods 1, 2, ...,  $K\beta$ . However, we accept the request for product  $K + 1$  at the last time period, because  $H^{K\beta+2}(x) = 0$ , and thus,  $r_{K+1} = 1 > 0 = H^{K\beta+2}(C) - H^{K\beta+2}(C - \sum_{i \in A^{K+1}} e_i)$ . Therefore, the total expected revenue from our approximate policy is  $\text{APP} = r_{K+1} = 1$ . So, the ratio between the total expected revenues of our approximate policy and the optimal policy is  $\frac{\text{APP}}{\text{OPT}} = 1/(1 + \frac{K\beta}{\beta+1}(1 - \frac{1}{\beta}))$ . If we choose  $\beta$  arbitrarily large, then  $\frac{\text{APP}}{\text{OPT}}$  becomes arbitrarily close to  $1/(1 + K)$ , which is equal to  $1/(1 + L)$ .

## 4. Extensions

We extend our approximate policy to the cases in which the customers choose among the offered products, and a product can use more than one unit of the capacity of a resource. We also discuss leveraging a linear program to build value function approximations.

### 4.1. Customer Choice Behavior

In the model in Section 2, each customer enters the system with a request for a particular product. We decide whether to accept or reject the request for this product. In this section, we extend our model and performance guarantee to a case in which we offer a subset of products to each arriving customer, and the customer chooses among the offered products or decides to leave without a purchase. Therefore, the customer does not arrive with a request for a particular product, and the product that the customer ends up choosing may depend on the subset of products that we offer. The notation that we use closely follows the one introduced in Section 2. We only discuss the additional notation that we need. If we offer the subset  $S \subseteq \mathcal{J}$  of products to a customer arriving at time period  $t$ , then the customer chooses product  $j \in S$  with probability  $\phi_j^t(S)$ . Naturally, we have  $\phi_j^t(S) = 0$  for all  $j \notin S$ . Note that the choices of the customers at different time periods may be governed

by different purchase probabilities. Therefore, we allow nonstationarities in the choice process of the customers. We refer to a subset of products that we offer to the customers as an assortment. We use  $\mathcal{F} \subseteq 2^{\mathcal{J}}$  to denote the set of feasible assortments that we can offer to an arriving customer. We impose the following mild assumption on the choice probabilities and the set of feasible assortments that we can offer to the customers.

**Assumption 1** (Substitutability and Feasibility). *For all  $t \in \mathcal{T}$ ,  $S \in \mathcal{F}$ ,  $j \in S$ , and  $k \notin S$ , we have  $\phi_j^t(S) \geq \phi_j^t(S \cup \{k\})$ . Also, if  $S \in \mathcal{F}$ , then we have  $R \in \mathcal{F}$  for all  $R \subseteq S$*

The first part of the assumption ensures that if we introduce an additional product into the assortment  $S$ , then the choice probability of a product that is already in the assortment  $S$  does not increase. This property holds for all choice models that are based on the random utility maximization principle. The second part of the assumption ensures that if we remove products from a feasible assortment, then the assortment remains feasible. To formulate the problem as a dynamic program, we let  $V^t(x)$  be the optimal total expected revenue over time periods  $t, t+1, \dots, T$  given that the capacities of the resources at time period  $t$  is  $x$ . We can compute the optimal value functions  $\{V^t(x) | x \in \mathcal{Q}, t \in \mathcal{T}\}$  using the dynamic program

$$\begin{aligned} V^t(x) &= \max_{S \in \mathcal{F}} \left\{ \sum_{j \in \mathcal{J}} \phi_j^t(S) \left( \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \right. \\ &\quad \cdot \left[ r_j + V^{t+1} \left( x - \sum_{i \in A^j} e_i \right) \right] \\ &\quad + \left( 1 - \sum_{j \in \mathcal{J}} \phi_j^t(S) + \sum_{j \in \mathcal{J}} \phi_j^t(S) \left( 1 - \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \right) \\ &\quad \left. \cdot V^{t+1}(x) \right\} \\ &= V^{t+1}(x) + \max_{S \in \mathcal{F}} \left\{ \sum_{j \in \mathcal{J}} \phi_j^t(S) \left( \prod_{i \in A^j} \mathbb{1}_{\{x_i \geq 1\}} \right) \right. \\ &\quad \cdot \left. \left[ r_j - V^{t+1}(x) + V^{t+1} \left( x - \sum_{i \in A^j} e_i \right) \right] \right\}, \end{aligned} \quad (10)$$

with the boundary condition that  $V^{T+1} = 0$ . If  $|A^j| = 1$  for all  $j \in \mathcal{J}$  so that each product uses exactly one resource, then our model becomes equivalent to the one in Golrezaei et al. (2014). In Golrezaei et al. (2014), there are multiple customer types, but multiple customer types do not bring any complication. Under multiple customer types, the result of the max operator in Equation (10) depends on the customer type and we simply take an expectation over the arriving customer type.

In the first equality in Equation (10), if we offer the assortment  $S$  at time period  $t$ , then the arriving customer chooses product  $j$  with probability  $\phi_j^t(S)$ . If we have sufficient resource capacities to serve product  $j$ , so that  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} = 1$ , then the customer purchases product  $j$ , in which case, we generate a revenue of  $r_j$  and the state of the resources at the next time period is  $\mathbf{x} - \sum_{i \in A^j} e_i$ . On the other hand, the arriving customer does not choose any product and decides to leave without a purchase with probability  $1 - \sum_{j \in \mathcal{J}} \phi_j^t(S)$ , in which case, the state of the resources at the next time period remains at  $\mathbf{x}$ . Lastly, the arriving customer chooses product  $j$  with probability  $\phi_j^t(S)$ , but if we do not have sufficient resource capacities to serve product  $j$ , so that  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} = 0$ , then the customer leaves without a purchase, in which case, the state of the resources at the next time period remains at  $\mathbf{x}$  as well. In Equation (10), we allow offering a product for which we lack sufficient resource capacities to serve. If the customer ends up choosing such a product, then the customer leaves without a purchase. Offering a product for which we lack sufficient resource capacities to serve may not sound realistic, but it is simple to argue that there exists an optimal policy that never offers such a product anyway. To establish this result, note that in the optimal solution to the second maximization problem in Equation (10), if we attempt to offer products for which we do not have enough resource capacity to serve, then we can drop all such products from the assortment, along with each product  $j$  such that  $r_j - V^{t+1}(\mathbf{x}) + V^{t+1}(\mathbf{x} - \sum_{i \in A^j} e_i) < 0$ , in which case, by Assumption 1, the choice probabilities of all remaining products in the assortment do not decrease, yielding another assortment that provides an objective value that is as large as the original one.

Similar to our earlier model, computing the optimal value functions  $\{V^t : t \in \mathcal{T}\}$  is intractable. We use a value function approximation of the form  $H^t(\mathbf{x}) = \sum_{j \in \mathcal{J}} \gamma_j^t \varphi_j(\mathbf{x})$ , where  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  is a collection of availability-tracking basis functions. We compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  in the value function approximations using a slight variation of our earlier algorithm.

- Initialization: Let  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  be any collection of availability-tracking basis functions and  $\theta \geq \Delta_{\mathcal{B}}$  be a tuning parameter. Initialize  $\gamma_j^{T+1} = 0$  for all  $j \in \mathcal{J}$ .

- Coefficient computation: For each  $t = T, T-1, \dots, 1$ , use the coefficients  $\{\gamma_j^{t+1} : j \in \mathcal{J}\}$  to compute the assortment  $\hat{S}^t \in \mathcal{F}$  at time period  $t$  as

$$\hat{S}^t = \arg \max_{S \in \mathcal{F}} \left\{ \sum_{j \in \mathcal{J}} \phi_j^t(S) \left( r_j - \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right) \right\}. \quad (11)$$

Then, use the coefficients  $\{\gamma_j^{t+1} : j \in \mathcal{J}\}$  and the assortment  $\hat{S}^t$  computed previously to compute  $\{\gamma_j^t : j \in \mathcal{J}\}$  as

$$\gamma_j^t = \phi_j^t(\hat{S}^t) \left( r_j - \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right) + \gamma_j^{t+1}. \quad (12)$$

The algorithm in Equation (12) specifies the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ , which, in turn, specify the approximate value functions  $\{H^t : t \in \mathcal{T}\}$ .

Given that the state of the resources at time period  $t$  is  $\mathbf{x}$ , we solve the maximization problem on the right side of the second equality in Equation (10) to find the optimal assortment to offer. We construct our approximate policy by replacing  $V^{t+1}$  in this problem with  $H^{t+1}$ . Thus, given that the state of the resources at time period  $t$  is  $\mathbf{x}$ , our approximate policy offers the assortment

$$S^{\text{App},t}(\mathbf{x}) = \arg \max_{S \in \mathcal{F}} \left\{ \sum_{j \in \mathcal{J}} \phi_j^t(S) \left( \prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}} \right) \cdot \left( r_j - H^{t+1}(\mathbf{x}) + H^{t+1} \left( \mathbf{x} - \sum_{i \in A^j} e_i \right) \right) \right\}. \quad (13)$$

An optimal solution to this problem can be viewed as the decision function of our approximate policy under customer choice behavior. By the next theorem, our approximate policy enjoys the same performance guarantee as in Theorem 1. The proof of this theorem uses a technique similar to the one in Section 3.3. We defer the proof to Appendix B in the e-companion.

**Theorem 2** (Performance under Choice). *If the choice probabilities and the feasible assortments satisfy Assumption 1 and the tuning parameter  $\theta$  satisfies  $\theta \geq \Delta_{\mathcal{B}}$ , then the total expected revenue obtained by the approximate policy is at least  $1/(1 + \theta L)$  fraction of the optimal.*

## 4.2. Multiple Units of Capacity Consumption

In this section, we extend our approach to allow products to use multiple units of a resource. Once again, the notation that we use closely follows the one introduced in Section 2. We only discuss the additional notation that we need. For each product  $j$  and resource  $i$ , we use  $a_{ij}$  to denote the number of units of resource  $i$  used by product  $j$ . In our earlier model, we have  $a_{ij} \in \{0, 1\}$  for all  $i \in \mathcal{L}, j \in \mathcal{J}$ . In this section, we consider the case where  $a_{ij}$  can be any nonnegative integer. As before,  $A^j = \{i \in \mathcal{L} : a_{ij} \geq 1\}$  denotes the set of resources used by product  $j$ , and  $L = \max_{j \in \mathcal{J}} |A^j|$  denotes the maximum number of resources used by a product. We use  $m_i = \max_{j \in \mathcal{J}} a_{ij}$  to denote the maximum number of units of resource  $i$  that is used by any product. Without loss of generality, we assume that the initial capacity of each resource  $i$  satisfies  $C_i \geq m_i$ .

Otherwise, there is a product that uses more units than the initial capacity of a resource, in which case, we can drop such a product. To find the optimal policy, we can use a dynamic program that is similar to the one in Equation (1). All we need to do is to replace all occurrences of  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}}$  with  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq a_{ij}\}}$  and all occurrences of  $\sum_{i \in A^j} e_i$  with  $\sum_{i \in A^j} a_{ij} e_i$  in the dynamic program.

We modify our basis functions as follows. For each product  $j \in \mathcal{J}$ , let  $G_j : \mathcal{Q} \rightarrow \mathcal{Q}$  be a mapping such that for each  $x \in \mathcal{Q}$ ,  $G_j(x) = (x_i \mathbf{1}_{\{x_i \geq a_{ij}\}} : i \in \mathcal{L})$ . Thus,  $G_j(x)$  leaves the  $i$ th component of  $x$  unchanged when the value of the component exceeds the amount of resource  $i$  consumed by product  $j$ ; otherwise,  $G_j(x)$  sets the component to zero. For a collection of availability-tracking basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$ , we use the value function approximation  $H^t$  given by

$$H^t(x) = \sum_{j \in \mathcal{J}} \gamma_j^t \varphi_j(G_j(x)). \quad (14)$$

We compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  in the previous value function approximation using the following algorithm:

- Initialization: Let  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  be any collection of availability-tracking basis functions and  $\theta \geq \Delta_{\mathcal{B}}$  be a tuning parameter. Initialize  $\gamma_j^{T+1} = 0$  for all  $j \in \mathcal{J}$ .
- Coefficient computation: For each  $t = T, T-1, \dots, 1$ , use the coefficients  $\{\gamma_j^{t+1} : j \in \mathcal{J}\}$  to compute  $\{\gamma_j^t : j \in \mathcal{J}\}$  as

$$\gamma_j^t = \lambda_j^t \left[ r_j - \theta \sum_{i \in A^j} \frac{2m_i - 1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right]^+ + \gamma_j^{t+1}. \quad (15)$$

To construct our approximate policy, we use the decision function in Equation (5) after replacing  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq 1\}}$  with  $\prod_{i \in A^j} \mathbf{1}_{\{x_i \geq a_{ij}\}}$  and  $\sum_{i \in A^j} e_i$  with  $\sum_{i \in A^j} a_{ij} e_i$ . This decision function provides the decisions made by our approximate policy given that the state of the resources at time period  $t$  is  $x$ . The following theorem gives a performance guarantee for our approximate policy when a product can consume multiple units of a resource. The proof is in Appendix C in the e-companion.

**Theorem 3** (Performance under Multiple Units Consumption). *Let  $M = \max_{i \in \mathcal{L}, j \in \mathcal{J}} a_{ij}$  be the maximum number of units of a resource used by a product. If the tuning parameter  $\theta$  satisfies  $\theta \geq \Delta_{\mathcal{B}}$ , then the total expected revenue obtained by the approximate policy is at least  $1/(1 + \theta(2M - 1)L)$  fraction of the optimal.*

To intuitively see the origin of the term  $2M - 1$ , assume that all products use one unit of a resource. If the capacity of the resource goes down to one, then we can accept at most one unit of any product. In contrast,

assume that all products use  $M$  units of a resource. If the capacity of the resource goes down to  $2M - 1$ , then we can accept at most one unit of any product.

#### 4.3. Leveraging a Linear Programming Approximation

Noting that  $\Lambda_j = \sum_{t \in \mathcal{T}} \lambda_j^t$  in the second constraint in problem (4) corresponds to the total expected number of requests for product  $j$  over the selling horizon, we can view the linear program in Equation (4) as an approximation to the network revenue management problem that is formulated under the assumption that the numbers of requests for the products take on their expected values. It is well-known that the optimal objective value of the linear program in Equation (4) provides an upper bound on the optimal total expected revenue; see Bertsimas and Popescu (2003). In practice, this upper bound becomes useful when assessing the optimality gaps of various heuristics. In this section, we show that we can leverage an optimal solution to the linear program in Equation (4) when constructing our value function approximations. We explain the idea using the model in Section 2, but we can incorporate customer choice behavior as shown in Section 4.1, and allow for products consuming multiple units of a resource as shown in Section 4.2. For a collection of availability-tracking basis functions  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$ , we approximate the optimal value functions  $\{V^t : t \in \mathcal{T}\}$  using value function approximations  $\{H^t : t \in \mathcal{T}\}$  of the form  $H^t(x) = \sum_{j \in \mathcal{J}} \gamma_j^t \varphi_j(x)$ . To compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ , we use the following algorithm.

- Initialization: Let  $\mathcal{B} = \{\varphi_j : j \in \mathcal{J}\}$  be any collection of availability-tracking basis functions,  $\theta \geq \Delta_{\mathcal{B}}$  be a tuning parameter, and  $(z_j^* : j \in \mathcal{J})$  be an optimal solution to the linear program in Equation (4). Initialize  $\gamma_j^{T+1} = 0$  for all  $j \in \mathcal{J}$ .
- Coefficient computation: For each  $t = T, T-1, \dots, 1$ , use the coefficients  $\{\gamma_j^{t+1} : j \in \mathcal{J}\}$  to compute  $\{\gamma_j^t : j \in \mathcal{J}\}$  as

$$\gamma_j^t = \frac{z_j^*}{\Lambda_j} \lambda_j^t \left[ r_j - \theta \sum_{i \in A^j} \frac{1}{C_i} \sum_{k \in \mathcal{J}} \mathbf{1}_{\{i \in A^k\}} \gamma_k^{t+1} \right]^+ + \gamma_j^{t+1}. \quad (16)$$

Once we construct the value function approximations  $\{H^t : t \in \mathcal{T}\}$  using the previous algorithm, we use the same decision function in Equation (5) in our approximate policy. The following theorem gives a performance guarantee for our approximate policy. The proof is in Appendix D in the e-companion.

**Theorem 4** (Performance with Linear Programming Approximation). *If the tuning parameter  $\theta$  satisfies  $\theta \geq \Delta_{\mathcal{B}}$ , then the total expected revenue obtained by the approximate policy is at least  $1/(1 + \theta L)$  fraction of the optimal.*

Intuitively, if product  $j$  is unlikely to contribute to the optimal total expected revenue, then we expect  $z_j^*$  to be close to zero. In this case, noting Equation (16), the coefficients  $\{\gamma_j^t : t \in \mathcal{T}\}$  for this product do not contribute significantly to the value function approximation.

## 5. Computational Experiments

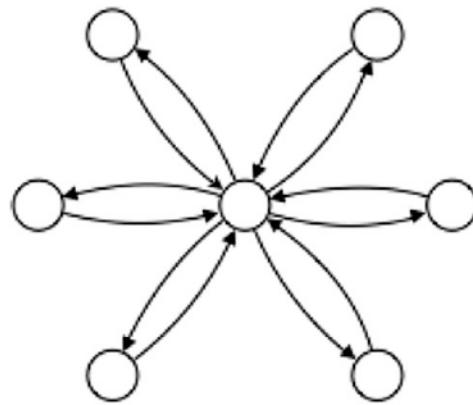
In this section, we discuss computational experiments that we conducted to assess the numerical performance of our approximate policy. During the course of our computational experiments, we also experiment with different availability-tracking basis functions.

### 5.1. Experimental Setup

In our computational experiments, we use the test problems in Topaloglu (2009). A number of other papers, including Hu et al. (2013), Brown and Smith (2014), Vossen and Zhang (2015a, b), and Kunnumkal and Talluri (2016a), used these test problems in their computational experiments as well. The test problems in Topaloglu (2009) originate from the airline setting, where the resources correspond to the flight legs and the products correspond to the itineraries. In our test problems, the airline network has  $N + 1$  locations. One location is the hub and the remaining  $N$  locations are the spokes. There is a flight leg from each spoke to the hub and a flight leg from the hub to each spoke. Therefore, the number of flight legs is  $2N$ . In Figure 1, we show the structure of the airline network with  $N = 6$ . We vary  $N$  in our computational experiments. Note that there are  $N$  origin-destination pairs that connect the hub to a spoke,  $N$  origin-destination pairs that connect a spoke to the hub, and  $N(N - 1)$  origin-destination pairs that connect a spoke to another spoke, resulting in  $2N + N(N - 1)$  origin-destination pairs. There is a high-fare and a low-fare itinerary that connects each origin-destination pair. Thus, the number of itineraries is  $2(2N + N(N - 1))$ . For a certain origin-destination pair, the revenue associated with the high-fare itinerary connecting this origin-destination pair is  $\kappa$  times the revenue associated with the corresponding low-fare itinerary, where  $\kappa$  captures the revenue difference between the high-fare and low-fare itineraries. We vary  $\kappa$  in our computational experiments as well.

The arrival probabilities for the itinerary requests are generated in such a way that the requests for the high-fare itineraries tend to arrive later in the selling horizon. To generate the arrival probabilities  $\{\lambda_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  for the itinerary requests, for each origin-destination pair  $(o, d)$ , we sample  $B_{o,d}$  from the uniform distribution over  $[0, 1]$  and  $\tau_{o,d}$  from the uniform distribution over  $\{\frac{1}{2}T, \dots, \frac{2}{3}T\}$ . In this case, the probability that we have a request for an itinerary that

**Figure 1.** Structure of the Airline Network with  $N = 6$  Spokes



connects the origin-destination pair  $(o, d)$  is proportional to  $B_{o,d}$ . The probability that we have a request for the low-fare itinerary that connects the origin-destination pair  $(o, d)$  decreases over time, whereas the probability that we have a request for the corresponding high-fare itinerary is zero until time period  $\tau_{o,d}$ , but this probability increases over time after time period  $\tau_{o,d}$ . Therefore, the requests for the high-fare itinerary start appearing only after time period  $\tau_{o,d}$ . To be precise, after generating  $B_{o,d}$  and  $\tau_{o,d}$ , for all  $t \in \mathcal{T}$ , we set  $\eta_{o,d}^{\text{t,low}}$  and  $\eta_{o,d}^{\text{t,high}}$  as

$$\eta_{o,d}^{\text{t,low}} = B_{o,d} \frac{T + 1 - t}{T} \quad \text{and}$$

$$\eta_{o,d}^{\text{t,high}} = \begin{cases} 0 & \text{if } t \leq \tau_{o,d} \\ B_{o,d} \frac{t - \tau_{o,d}}{T - \tau_{o,d}} & \text{otherwise.} \end{cases}$$

Letting  $\mathcal{D}$  be the set of all origin-destination pairs, if itinerary  $j$  is the low-fare itinerary for origin-destination pair  $(o, d)$ , then we set  $\lambda_j^t = \eta_{o,d}^{\text{t,low}} / \sum_{(s,r) \in \mathcal{D}} (\eta_{s,r}^{\text{t,low}} + \eta_{s,r}^{\text{t,high}})$ , but if itinerary  $j$  is the corresponding high-fare itinerary, then we set  $\lambda_j^t = \eta_{o,d}^{\text{t,high}} / \sum_{(s,r) \in \mathcal{D}} (\eta_{s,r}^{\text{t,low}} + \eta_{s,r}^{\text{t,high}})$ . Since the requests for the high-fare itineraries tend to arrive later in the selling horizon, it becomes important to reserve the capacities for the high-fare itinerary requests by rejecting the requests early in the selling horizon. The total expected demand for the capacity on flight leg  $i$  is  $\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \mathbb{1}_{\{i \in A^t\}} \lambda_j^t$ , so the initial capacity of flight leg  $i$  is set to be  $C_i = \frac{1}{\alpha} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \mathbb{1}_{\{i \in A^t\}} \lambda_j^t$ . Thus, larger values for  $\alpha$  yield tighter capacities. We vary  $\alpha$  in our computational experiments.

Letting  $N$ ,  $\kappa$ , and  $\alpha$  be as previously, and recalling that  $T$  is the length of the selling horizon, we vary  $T \in \{200, 600\}$ ,  $N \in \{4, 5, 6, 8\}$ ,  $\kappa \in \{2, 4\}$ , and  $\alpha \in \{1.0, 1.2, 1.6\}$ , to get 48 test problems.

## 5.2. Benchmark Methods

In our computational experiments, we work with six benchmarks, one of which is our approximate policy. We proceed to describing our benchmarks.

**5.2.1. Approximate Policy (APP).** This benchmark is our approximate policy with the decision function in Equation (5). We experimented with different basis functions. The basis function  $\varphi_j(x) = \min_{i \in A^j}(1 - e^{-x_i/C_i})/(1 - e^{-1})$ , which is availability tracking by the discussion in Section 3.1, provided consistent improvements over the others that we experimented with. Thus, we use this basis function. One can check that  $\Delta_{\mathcal{B}} = 1/(1 - e^{-1})$  for this basis function. In our computational results, we discuss our experimentation with different basis functions. Note that our basis function resembles the tradeoff function  $\Psi(x) = (1 - e^{-x})/(1 - e^{-1})$  in Golrezaei et al. (2014).

In our practical implementation of APP, we make two modifications. First, we divide the selling horizon into five equal segments and reconstruct our value function approximations at the beginning of each segment. In particular, the beginning of segment  $k$  corresponds to time period  $(k-1)\frac{T}{5} + 1$ . If the remaining capacities on the flight legs at the beginning of segment  $k$  are given by the vector  $x$ , then we replace  $C_i$  in the recursion in Equation (3) with  $x_i$  and use this recursion over time periods  $T, T-1, \dots, (k-1)\frac{T}{5} + 1$  to compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t = (k-1)\frac{T}{5} + 1, \dots, T\}$ . These coefficients specify the value function approximations that we use when making the decisions over segment  $k$ . When we reach the beginning of the next segment, we reconstruct our value function approximations in a similar fashion. Second, we calibrate the value for the tuning parameter  $\theta$  at the beginning of each segment. The values of the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  in Equation (3) depend on  $\theta$ , which, in turn, implies that the total expected revenue obtained by APP also depends on  $\theta$ . When reconstructing our value function approximations at the beginning of each segment, we search for the best tuning parameter over the interval  $\mathcal{J} = [1/(1 - e^{-1}), 15]$  with a precision of 0.01. Given that we use the tuning parameter  $\theta$  when reconstructing our value function approximations at the beginning of segment  $k$ , let  $U^{k,\theta}(x)$  be the total expected revenue obtained by APP over time periods  $(k-1)\frac{T}{5} + 1, \dots, T$  starting with the capacities  $x$  for the flight legs. Computing the total expected revenue  $U^{k,\theta}(x)$  exactly is intractable, because computing this quantity requires solving a dynamic program similar to the one in Equation (8), but we estimate this quantity using simulation. At the beginning of segment  $k$ , we choose the value of the tuning parameter  $\theta$  as  $\arg \max\{U^{k,\theta}(x) : \theta \in \mathcal{J} \cap \{0.01 \times \ell : \ell = 0, 1, \dots\}\}$ . We

use this value for the tuning parameter until we reach the beginning of the next segment. The theoretical performance guarantee that we give in Theorem 1 will be the strongest when we use the smallest possible value for  $\theta$  with  $\theta \geq \Delta_{\mathcal{B}}$ , but choosing a different value for the tuning parameter may actually provide better practical performance for APP.

**5.2.2. Bid Price Policy (BPP).** This benchmark is the well-known bid price policy; see section 3.3 in Talluri and van Ryzin (2005). The idea behind BPP is to use the optimal values of the dual variables associated with the first constraint in the linear program in Equation (4) to estimate the value of a unit of capacity on each flight leg. In this case, if the revenue from a certain itinerary exceeds the value of the capacities used by this itinerary, then we accept the request for the itinerary. To be specific, letting  $(\mu_i^* : i \in \mathcal{L})$  be the optimal values of the dual variables associated with the first constraint in problem (4), BPP accepts a request for itinerary  $j$  if and only if  $r_j \geq \sum_{i \in A^j} \mu_i^*$  and there are sufficient capacities to serve a request for itinerary  $j$ . In our practical implementation of BPP, we divide the selling horizon into five equal segments and resolve problem (4) at the beginning of each segment. In particular, if the remaining capacities on the flight legs at the beginning of segment  $k$  are given by the vector  $x$ , then we replace  $C_i$  with  $x_i$  and  $\Lambda_j$  with  $\sum_{t=(k-1)\frac{T}{5}+1}^T \lambda_j^t$  in problem (4). Letting  $(\mu_i^* : i \in \mathcal{L})$  be the optimal values of the dual variables associated with the first constraint, we use these values of the dual variables during segment  $k$ . Using a similar approach, we recompute the policy parameters for all other benchmarks at the beginning of each segment, but for economy of space, we do not discuss recomputation any longer.

**5.2.3. Randomized Linear Program (RLP).** In this benchmark, we use the realizations of the total numbers of itinerary requests over the selling horizon, as opposed to their expected values, to capture the distribution information for the total numbers of itinerary requests. Using the random variable  $D_j$  to denote the total number of requests for itinerary  $j$  over the selling horizon, we replace  $\Lambda_j$  on the right side of the second constraint in problem (4) with  $D_j$ . As a function of  $D = (D_j : j \in \mathcal{J})$ , letting  $(\mu_i^*(D) : i \in \mathcal{L})$  be the optimal values of the dual variables associated with the first constraint in problem (4), we use the  $\mathbb{E}\{\mu^*(D)\}$  to estimate the value of a unit of capacity on flight leg  $i$ . In this case, RLP accepts a request for itinerary  $j$  as long as  $r_j \geq \sum_{i \in A^j} \mathbb{E}\{\mu_i^*(D)\}$ ; see Talluri and van Ryzin (1999). Computing the expectation  $\mathbb{E}\{\mu^*(D)\}$  exactly is intractable, so we estimate this expectation using simulation.

**5.2.4. Finite Differences (DIF).** Here, we use the optimal objective value of problem (4) to estimate the value of the capacities used by an itinerary. As a function of  $C = (C_i : i \in \mathcal{L})$ , we let  $Z_{LP}^*(C)$  be the optimal objective value of problem (4), in which case, we estimate the value of the capacities used by itinerary  $j$  as  $Z_{LP}^*(C) - Z_{LP}^*(C - \sum_{i \in A^j} e_i)$ . Thus, DIF accepts a request for itinerary  $j$  as long as  $r_j \geq Z_{LP}^*(C) - Z_{LP}^*(C - \sum_{i \in A^j} e_i)$ ; see Bertsimas and Popescu (2003).

**5.2.5. Dynamic Programming Decomposition (DEC).** The idea behind this benchmark is to decompose the dynamic programming formulation of the problem by the flight legs, in which case, we can solve dynamic programs with scalar state variables to obtain value function approximations; see section 4.2 in Zhang and Adelman (2009). To our knowledge, DEC is one of the strongest heuristics in practice, but it does not have a performance guarantee.

**5.2.6. Online Packing Policy (OPP).** This benchmark uses a linear program similar to the one in Equation (4) to construct a policy for online packing problems; see Kesselheim et al. (2014). If the arrivals are stationary, then this policy has the competitive ratio of  $1 - O(\sqrt{(\log L)/c_{\min}})$  that we discuss in the introduction section, but in Appendix E in the e-companion, we give a simple example to show that

this policy can perform arbitrarily poorly under non-stationary arrivals.

### 5.3. Computational Results

Table 1 shows our computational results on the test problems with  $T = 200$  time periods in the selling horizon, whereas Table 2 shows our computational results on the test problems with  $T = 600$ . The layouts of the two tables are identical. In the first column, we show the parameter configuration for each test problem using the tuple  $(T, N, \kappa, \alpha)$ , where  $N$ ,  $\kappa$ , and  $\alpha$  are as discussed in our experimental setup. In the second column, we show the upper bound on the optimal total expected revenue provided by the optimal objective value of problem (4). In the third to eighth columns, we show the total expected revenues obtained by APP, BPP, RLP, DIF, DEC, and OPP. We estimate these total expected revenues by simulating the performance of each benchmark over 100 sample paths. In the ninth to thirteenth columns, we give the percent gaps between the total expected revenues obtained by APP and the remaining five benchmarks.

Considering the results in Table 1, the performance of APP is better than that of BPP with a substantial margin. RLP and DIF both perform better than BPP, but this improvement is not enough to catch up with APP. We underline two trends by comparing the

**Table 1.** Computational Results for the Test Problems with  $T = 200$  Time Periods in the Selling Horizon

Parameters ( $T, N, \kappa, \alpha$ )	Upper bound	Total expected revenue						Percent gap with APP				
		APP	BPP	RLP	DIF	DEC	OPP	BPP	RLP	DIF	DEC	OPP
(200, 4, 4, 1.0)	21,531	20,013	19,377	19,937	19,608	20,076	19,601	3.18	0.38	2.02	-0.32	2.06
(200, 4, 4, 1.2)	19,882	18,386	17,140	17,964	17,529	18,538	17,276	6.78	2.29	4.66	-0.82	6.04
(200, 4, 4, 1.6)	17,530	15,993	14,474	15,712	14,996	16,185	13,507	9.50	1.76	6.24	-1.20	15.55
(200, 4, 8, 1.0)	34,571	32,655	30,692	32,447	31,305	32,845	31,858	6.01	0.64	4.14	-0.58	2.44
(200, 4, 8, 1.2)	32,922	31,020	27,324	30,310	28,384	31,284	28,025	11.91	2.29	8.50	-0.85	9.65
(200, 4, 8, 1.6)	30,570	28,704	24,062	27,890	25,461	28,861	21,901	16.17	2.83	11.30	-0.55	23.70
(200, 5, 4, 1.0)	22,144	20,984	20,197	20,796	20,488	21,139	20,478	3.75	0.90	2.36	-0.74	2.41
(200, 5, 4, 1.2)	21,263	19,565	18,462	19,225	18,933	19,716	18,463	5.64	1.73	3.23	-0.77	5.63
(200, 5, 4, 1.6)	18,870	17,037	15,406	16,676	16,090	17,260	14,925	9.57	2.12	5.56	-1.31	12.39
(200, 5, 8, 1.0)	35,387	33,943	31,844	33,519	32,491	34,219	33,048	6.18	1.25	4.28	-0.81	2.63
(200, 5, 8, 1.2)	34,495	32,318	29,232	31,551	30,395	32,653	29,702	9.55	2.37	5.95	-1.04	8.10
(200, 5, 8, 1.6)	32,081	29,666	24,971	28,943	26,800	30,068	24,074	15.82	2.44	9.66	-1.36	18.85
(200, 6, 4, 1.0)	22,300	20,595	19,819	20,496	20,110	20,699	19,943	3.77	0.48	2.35	-0.51	3.17
(200, 6, 4, 1.2)	20,932	19,049	17,927	18,753	18,463	19,174	17,793	5.89	1.55	3.08	-0.65	6.60
(200, 6, 4, 1.6)	18,592	16,595	15,220	16,302	15,863	16,786	14,235	8.29	1.76	4.41	-1.15	14.22
(200, 6, 8, 1.0)	35,544	33,338	31,132	33,205	31,946	33,644	32,186	6.62	0.40	4.17	-0.92	3.46
(200, 6, 8, 1.2)	34,172	31,623	28,504	31,107	29,737	32,004	28,603	9.86	1.63	5.96	-1.21	9.55
(200, 6, 8, 1.6)	31,824	29,191	24,923	28,459	26,702	29,551	22,880	14.62	2.51	8.53	-1.23	21.62
(200, 8, 4, 1.0)	20,052	18,359	17,508	17,875	17,742	18,421	17,634	4.63	2.63	3.36	-0.34	3.95
(200, 8, 4, 1.2)	18,952	16,936	15,753	16,354	16,188	17,054	15,728	6.98	3.43	4.41	-0.70	7.13
(200, 8, 4, 1.6)	16,833	14,676	13,371	14,161	14,019	14,831	12,439	8.89	3.51	4.48	-1.05	15.24
(200, 8, 8, 1.0)	31,835	29,742	27,378	28,779	28,058	29,890	28,346	7.95	3.24	5.66	-0.50	4.69
(200, 8, 8, 1.2)	30,727	28,232	24,793	27,116	25,956	28,412	25,218	12.18	3.95	8.06	-0.64	10.68
(200, 8, 8, 1.6)	28,608	25,913	21,844	24,837	23,616	26,116	19,869	15.70	4.15	8.87	-0.78	23.32
Average								8.73	2.09	5.47	-0.83	9.71

**Table 2.** Computational Results for the Test Problems with  $T = 600$  Time Periods in the Selling Horizon

Parameters ( $T, N, \kappa, \alpha$ )	Upper bound	Total expected revenue						Percent gap with APP				
		APP	BPP	RLP	DIF	DEC	OPP	BPP	RLP	DIF	DEC	OPP
(600, 4, 4, 1.0)	32,409	30,596	29,751	30,209	29,969	30,724	29,988	2.76	1.27	2.05	-0.42	1.99
(600, 4, 4, 1.2)	29,852	27,990	26,100	27,455	26,546	28,233	26,138	6.75	1.91	5.16	-0.87	6.62
(600, 4, 4, 1.6)	26,324	24,414	22,009	24,060	22,567	24,668	20,474	9.85	1.45	7.57	-1.04	16.14
(600, 4, 8, 1.0)	52,086	49,909	47,249	49,168	47,796	50,194	48,553	5.33	1.49	4.23	-0.57	2.72
(600, 4, 8, 1.2)	49,529	47,214	41,751	46,069	43,748	47,677	42,223	11.57	2.43	7.34	-0.98	10.57
(600, 4, 8, 1.6)	46,001	43,637	36,632	42,941	38,256	44,008	33,108	16.05	1.59	12.33	-0.85	24.13
(600, 5, 4, 1.0)	33,299	31,928	30,763	31,592	31,027	32,095	31,217	3.65	1.05	2.82	-0.52	2.23
(600, 5, 4, 1.2)	31,943	29,813	28,290	29,473	28,832	30,147	28,206	5.11	1.14	3.29	-1.12	5.39
(600, 5, 4, 1.6)	28,343	25,903	23,847	25,713	24,582	26,405	22,724	7.94	0.73	5.10	-1.94	12.27
(600, 5, 8, 1.0)	53,285	51,563	48,653	50,871	49,263	51,923	50,250	5.64	1.34	4.46	-0.70	2.55
(600, 5, 8, 1.2)	51,904	49,207	44,958	48,639	46,078	49,737	45,280	8.64	1.15	6.36	-1.08	7.98
(600, 5, 8, 1.6)	48,283	45,221	38,966	44,582	40,784	45,911	36,476	13.83	1.41	9.81	-1.52	19.34
(600, 6, 4, 1.0)	26,873	25,369	24,405	24,777	24,576	25,445	24,661	3.80	2.33	3.13	-0.30	2.79
(600, 6, 4, 1.2)	25,184	23,325	21,945	22,553	22,455	23,517	21,841	5.92	3.31	3.73	-0.82	6.36
(600, 6, 4, 1.6)	22,274	20,327	18,542	19,622	19,259	20,571	17,466	8.78	3.47	5.25	-1.20	14.07
(600, 6, 8, 1.0)	42,865	41,102	38,411	39,830	38,935	41,305	39,768	6.55	3.09	5.27	-0.49	3.24
(600, 6, 8, 1.2)	41,166	38,936	34,848	37,267	36,205	39,238	35,128	10.50	4.29	7.02	-0.77	9.78
(600, 6, 8, 1.6)	38,252	35,845	30,536	34,430	32,492	36,238	28,181	14.81	3.95	9.35	-1.09	21.38
(600, 8, 4, 1.0)	24,167	22,332	21,241	21,601	21,616	22,466	21,546	4.89	3.27	3.20	-0.60	3.52
(600, 8, 4, 1.2)	22,755	20,539	19,074	19,726	19,668	20,710	19,077	7.13	3.96	4.24	-0.83	7.12
(600, 8, 4, 1.6)	20,228	17,852	16,411	17,221	17,083	18,076	15,002	8.07	3.53	4.31	-1.25	15.97
(600, 8, 8, 1.0)	38,395	36,299	33,270	34,870	34,159	36,521	34,599	8.34	3.94	5.89	-0.61	4.68
(600, 8, 8, 1.2)	36,976	34,376	30,063	32,791	31,831	34,614	30,603	12.55	4.61	7.40	-0.69	10.97
(600, 8, 8, 1.6)	34,449	31,589	26,876	30,209	28,856	31,840	23,875	14.92	4.37	8.65	-0.79	24.42
Average								8.47	2.55	5.75	-0.88	9.84

performance of APP with that of RLP, but similar observations hold when we compare the performance of APP with that of BPP or DIF. We focus on comparing APP with RLP, because RLP already performs noticeably better than BPP and DIF, which are the two other benchmarks that are based on the linear program in Equation (4). The first trend is that as the parameter  $\alpha$  increases and the capacities on the flight legs get tighter, the performance gaps between APP and RLP get larger. Considering the test problems with  $\alpha = 1.0$ ,  $\alpha = 1.2$ , and  $\alpha = 1.6$  separately, the average percent gaps between the total expected revenues obtained by APP and RLP are, respectively, 1.24%, 2.41%, and 2.64%. As the capacities on the flight legs get tighter, it becomes more important to protect the capacity for the high-fare itinerary requests that tend to arrive later in the selling horizon. It appears that APP does a better job of capturing this tradeoff. The second trend is that as the parameter  $\kappa$  increases and the revenue difference between the high-fare and low-fare itineraries increases, the performance gap between APP and RLP increases as well. Considering the test problems with  $\kappa = 2$  and  $\kappa = 4$  separately, the average percent gaps between the total expected revenues obtained by APP and RLP are, respectively, 1.88% and 2.31%. When the revenue difference between the high-fare and low-fare itineraries increases, the opportunity cost of not having the capacity to serve a high-fare itinerary request also increases and

APP seems to do a better job of reserving the capacity for high-fare itinerary requests. APP lags behind DEC with a small but consistent margin. Over all test problems, the average gap between the performance of APP and DEC is 0.83%. As mentioned earlier, to our knowledge, DEC is one of the strongest heuristics for network revenue management problems in practice, but we emphasize that DEC does not have a theoretical performance guarantee. Similar to the trends discussed earlier in this paragraph, the performance gap between APP and DEC gets larger as the capacities on the flight legs get tighter or the revenue difference between the high-fare and low-fare itineraries increases. The performance of OPP is not competitive to APP, which is not surprising because OPP is designed to deal with stationary arrivals. Similar observations hold for the results in Table 2. APP consistently performs better than BPP. Although RLP and DIF both perform better than BPP, they do not catch up with APP. APP lags behind DEC by a small but consistent margin. OPP significantly lags behind APP.

The calibrated value of the tuning parameter  $\theta$  that we use for APP depends on the problem parameters. In Table 3, we show the value of the tuning parameter for all of our test problems when we calibrate the tuning parameter at the beginning of the selling horizon. Our results indicate that the calibrated value of the tuning parameter gets smaller as  $\alpha$  increases and the capacities on the flight legs get tighter. Intuitively

speaking, as the capacities on the flight legs get tighter and the resources become more scarce, we expect the value of a unit of capacity to increase. Indeed, we numerically observed that if we decrease the value of  $\theta$  in the recursion in Equation (3), then values of the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  computed through this recursion tend to increase, in which case, the value of a unit of capacity also increases. We carried out all of our computational experiments using Java 1.8.0 on 2.8 GHz Intel Xeon E5-2680 CPU with 1 GB of RAM.

**Table 3.** Calibrated Values of the Tuning Parameter  $\theta$

Parameters ( $T, N, \kappa, \alpha$ )	$\theta$
(200, 4, 4, 1.0)	1.91
(200, 4, 4, 1.2)	1.60
(200, 4, 4, 1.6)	1.59
(200, 4, 8, 1.0)	5.50
(200, 4, 8, 1.2)	4.75
(200, 4, 8, 1.6)	3.76
(200, 5, 4, 1.0)	2.23
(200, 5, 4, 1.2)	1.65
(200, 5, 4, 1.6)	1.59
(200, 5, 8, 1.0)	6.33
(200, 5, 8, 1.2)	4.84
(200, 5, 8, 1.6)	3.99
(200, 6, 4, 1.0)	1.78
(200, 6, 4, 1.2)	1.62
(200, 6, 4, 1.6)	1.59
(200, 6, 8, 1.0)	5.64
(200, 6, 8, 1.2)	4.75
(200, 6, 8, 1.6)	3.98
(200, 8, 4, 1.0)	1.61
(200, 8, 4, 1.2)	1.59
(200, 8, 4, 1.6)	1.59
(200, 8, 8, 1.0)	5.03
(200, 8, 8, 1.2)	4.37
(200, 8, 8, 1.6)	3.40
(600, 4, 4, 1.0)	1.82
(600, 4, 4, 1.2)	1.62
(600, 4, 4, 1.6)	1.59
(600, 4, 8, 1.0)	5.66
(600, 4, 8, 1.2)	4.86
(600, 4, 8, 1.6)	3.69
(600, 5, 4, 1.0)	1.97
(600, 5, 4, 1.2)	1.60
(600, 5, 4, 1.6)	1.60
(600, 5, 8, 1.0)	6.20
(600, 5, 8, 1.2)	4.84
(600, 5, 8, 1.6)	3.93
(600, 6, 4, 1.0)	1.60
(600, 6, 4, 1.2)	1.59
(600, 6, 4, 1.6)	1.60
(600, 6, 8, 1.0)	5.65
(600, 6, 8, 1.2)	4.81
(600, 6, 8, 1.6)	3.93
(600, 8, 4, 1.0)	1.71
(600, 8, 4, 1.2)	1.59
(600, 8, 4, 1.6)	1.59
(600, 8, 8, 1.0)	5.50
(600, 8, 8, 1.2)	4.38
(600, 8, 8, 1.6)	3.45

For the largest test problems with  $T = 600$  time periods and  $N = 8$  spokes, the average CPU time to compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$  for a fixed value of  $\theta$  was about 0.05 seconds.

In our implementation of APP, we experimented with five other availability-tracking basis functions:

$$\begin{aligned}\varphi_j^{\text{prd-exp}}(x) &= \prod_{i \in A^j} \frac{\left(1 - e^{-x_i/C_i}\right)}{\left(1 - e^{-1}\right)}, \\ \varphi_j^{\text{min}}(x) &= \min_{i \in A^j} \frac{x_i}{C_i}, \\ \varphi_j^{\text{prd}}(x) &= \prod_{i \in A^j} \frac{x_i}{C_i}, \\ \varphi_j^{\text{exp-sum}}(x) &= \exp \left( \sum_{i \in A^j} \cdot \left(1 - \frac{C_i}{x_i}\right) \right) \text{ and} \\ \varphi_j^{\text{recip-sum}}(x) &= \frac{|A^j|}{\sum_{i \in A^j} \frac{C_i}{x_i}}.\end{aligned}$$

In Table 4, for economy of space, we consider a subset of our test problems and show the performance of APP with the five basis functions, along with  $\varphi_j^{\text{min-exp}}(x) = \min_{i \in A^j} (1 - e^{-x_i/C_i}) / (1 - e^{-1})$  used in our earlier computational results. In the first column, we show the parameter configuration for each test problem. In the second to seventh columns, we show the total expected revenues obtained by APP with the six basis functions. In the eighth to twelfth columns, we show the percent gap between the performance of APP with the basis function  $\varphi_j^{\text{min-exp}}(x) = \min_{i \in A^j} (1 - e^{-x_i/C_i}) / (1 - e^{-1})$  and the remaining five basis functions. Our results indicate that the performance of APP is somewhat robust to the choice of basis functions, but by experimenting with different basis functions, we can improve the performance by about 1.20% on average.

In all of our test problems, the maximum number of resources used by a product is two. The theoretical performance guarantee that we give for APP in Theorem 1 depends on the maximum number of resources used by a product. In Appendix F in the e-companion, we provide computational experiments in the hotel network revenue management setting, where we systematically vary the maximum number of resources used by a product. Our results demonstrate that APP maintains its edge over BPP, RLP, DIF, and OPP, and the performance gap between APP and DEC remains stable. Lastly, in our computational experiments, the arrivals for the product requests are nonstationary. OPP has a competitive ratio under stationary arrivals. In Appendix G in the e-companion, we provide computational experiments to test the performance of OPP under stationary arrivals. Under stationary arrivals, the performance of OPP gets better, but APP still provides significant improvements

**Table 4.** Performance of APP with Different Basis Functions

Parameters ( $T, N, \kappa, \alpha$ )	Total expected revenue						Percent gap with Min-Exp				
	Min-Exp	Prd-Exp	Min	Prd	Exp-Sum	Recip-Sum	Prd-Exp	Min	Prd	Exp-Sum	Recip-Sum
(200, 6, 4, 1.0)	20,595	20,420	20,447	20,086	20,284	20,427	0.85	0.72	2.47	1.51	0.82
(200, 6, 4, 1.2)	19,049	18,847	18,820	18,341	18,586	18,874	1.06	1.20	3.72	2.43	0.92
(200, 6, 4, 1.6)	16,595	16,621	16,459	16,149	16,337	16,539	-0.16	0.82	2.69	1.56	0.33
(200, 6, 8, 1.0)	33,338	33,248	33,123	32,544	32,914	33,088	0.27	0.64	2.38	1.27	0.75
(200, 6, 8, 1.2)	31,623	31,619	31,425	30,816	31,059	31,192	0.01	0.63	2.55	1.78	1.36
(200, 6, 8, 1.6)	29,191	29,265	28,928	28,654	29,046	29,090	-0.25	0.90	1.84	0.50	0.34

over OPP. An examination of Theorem 3 and Lemma 7 in Kesselheim et al. (2014) shows that the competitive ratio of OPP is  $\max\{1 - 45\sqrt{(1 + \log_2 L)/c_{\min}}, 1/(8e(2L)^{1/(c_{\min}-1)})\}$  for  $c_{\min} \geq 2$ , which, although it approaches one as the capacities of the resource get large, can be substantially less than one for practical instances. For example, for  $c_{\min} = 100$  and  $L = 2$ , this competitive ratio is about 0.045, which is significantly less than the performance guarantee of  $1/(1 + L) = 1/3$  for APP. Note that OPP does not use forecasts of the numbers of requests for different products, which can partly explain its poor performance. The practical performance of OPP is not competitive, but it is remarkable that OPP has an asymptotic performance guarantee without using forecasts.

In all of our computational experiments, we recompute the policy parameters five times over the selling horizon. In Appendix H in the e-companion, we test the performance of the benchmarks when we compute the policy parameters only once at the beginning of the selling horizon. Lastly, we also implemented the approach in Section 4.3, which uses an optimal solution to problem (4) to compute the coefficients  $\{\gamma_j^t : j \in \mathcal{J}, t \in \mathcal{T}\}$ . For our test problems, this approach did not provide noticeable improvements in the performance of APP.

## 6. Conclusion

We developed an approximate policy with a performance guarantee for network revenue management problems, which, to our knowledge, is unique as it works under nonstationary arrivals. In the paper, we pointed out several extensions of our approximate policy, but there are still other extensions that are possible. For example, in some papers in the online packing literature, there are multiple service modes with different revenues and resource consumptions. If we choose to accept a product request, then we decide which mode to use to serve the request; see, for example, Feldman et al. (2010) and Kesselheim et al. (2014). We can incorporate multiple service modes into our approximate policy by using our extension to the customer choice behavior given in Section 4.1. In particular, we use  $\mathcal{K}$  to denote the set of service modes.

At time period  $t$ , we have a request for product  $j$  with probability  $\lambda_j^t$ . If we use mode  $k$  to serve a request for product  $j$ , then we generate a revenue of  $r_{j,k}$  and consume one unit of capacity for each resource in the set  $A^{j,k} \subseteq \mathcal{L}$ . We can reformulate the problem with multiple service modes equivalently as an instance of the problem in Section 4.1. In our reformulation, we refer to each product and service mode combination as a meta-product. At each time period, we choose a subset of meta-products to offer to the customers. Offering meta-product  $(j, k) \in \mathcal{J} \times \mathcal{K}$  corresponds to being willing to use mode  $k$  to serve a request for product  $j$ . Not offering any of the meta-products  $\{(j, k) : k \in \mathcal{K}\}$  corresponds to not being willing to accept a request for product  $j$ . The feasible subsets of meta-products that we can offer to the customers at a particular time period is given by  $\mathcal{F} = \{S \subseteq \mathcal{J} \times \mathcal{K} : |S \cap \{(j, k) : k \in \mathcal{K}\}| \leq 1 \forall j \in \mathcal{J}\}$ , meaning that for each product  $j$ , we can offer at most one meta-product of the form  $(j, k)$ . In this way, we ensure that if we are willing to accept a request for product  $j$ , then we choose one mode to serve it. If we offer the subset  $S$  of meta-products at time period  $t$  such that  $(j, k) \in S$  for service mode  $k$  and we have a request for product  $j$ , then the arriving customer chooses meta-product  $(j, k)$ . Therefore, we have the choice probability  $\phi_{j,k}^t(S) = \lambda_j^t$  if  $(j, k) \in S$ . Otherwise,  $\phi_{j,k}^t(S) = 0$ . The revenue of meta-product  $(j, k)$  is  $r_{j,k}$ . If we sell one unit of meta-product  $(j, k)$ , then we consume the capacities of the resources in the set  $A^{j,k}$ . Thus, replacing the products in the formulation in Section 4.1 with the meta-products, we can check that the feasible subsets of meta-products and the choice probabilities given previously satisfy Assumption 1. So, we can use the formulation in Section 4.1 to come up with an approximate policy to decide which meta-products to offer at each time period to maximize the total expected revenue, which, in turn, yields an approximate policy to decide which product requests to accept and which mode to use for the accepted requests.

By the preceding discussion, the extension in Section 4.1 allows us to incorporate multiple service modes, but there are other variants of the network revenue management problem that are difficult to address using

our approach. For example, if overbooking is allowed, then the dynamic programming formulation of the problem is fundamentally different from the one that we use, because the state variable needs to keep track of the number of accepted bookings for each product, which ultimately may or may not show up. Thus, extending our work to handle overbooking is not straightforward. Also, the choice of our basis functions and the algorithm that we use to construct the coefficients of the basis functions strictly exploit the structure of the network revenue management problem. Extending our approach to a broader class of dynamic programs is certainly worthwhile, but such extensions appear to be nontrivial to us at this point. Another important point is that our choice of the basis functions was based on experimentation. The definition of availability-tracking basis functions provides some guidance on the choice of the basis functions, but a more systematic approach for choosing the basis functions is a useful research direction. Moreover, although our approach has a performance guarantee, this performance guarantee stays away from one. It would be useful if we can establish that our approach becomes asymptotically optimal in some regime, such as the one where the resource capacities and the expected demands for the products increase linearly with the same rate. Lastly, the dynamic programming decomposition approach, which we used as a benchmark in our computational experiments, is one of the strongest heuristics for network revenue management problems. To our knowledge, however, this approach does not have a performance guarantee. It would be useful to understand whether it is possible to give a performance guarantee for this approach.

### Acknowledgments

The authors gratefully acknowledge the comments of the area editor, associate editor, and three referees, which substantially improved the paper's exposition, technical results, and computational experiments.

### References

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

Agrawal S, Wang Z, Ye Y (2014) A dynamic near-optimal algorithm for online linear programming. *Oper. Res.* 62(4):876–890.

Asadpour A, Nazerzadeh H (2016) Maximizing stochastic monotone submodular functions. *Management Sci.* 62(8):2374–2391.

Bertsimas D, Popescu I (2003) Revenue management in a dynamic network environment. *Transportation Sci.* 37(3):257–277.

Bront JJM, Mendez Diaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Oper. Res.* 57(3):769–784.

Brown DB, Smith JE (2014) Information relaxations, duality, and convex stochastic dynamic programs. *Oper. Res.* 62(6):1394–1415.

Buchbinder N, Naor JS (2009) The design of competitive online algorithms via a primal: Dual approach. *Foundations Trends Theoret. Comput. Sci.* 3(2–3):93–263.

Chan CW, Farias VF (2009) Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Math. Oper. Res.* 34(2):333–350.

Chaneton JM, Vulcano G (2011) Computing bid prices for revenue management under customer choice behavior. *Manufacturing Service Oper. Management* 13(4):452–470.

Cooper WL (2002) Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.* 50(4):720–727.

Cooper WL, Homem-de-Mello T (2007) Some decomposition methods for revenue management. *Transportation Sci.* 41(3):332–353.

Devanur NR, Hayes TP (2009) The adwords problem: Online keyword matching with budgeted bidders under random permutations. Fortnow L, Pu P, eds. *Proc. ACM Conf. Electronic Commerce* (ACM, New York), 71–78.

Devanur NR, Jain K, Sivan B, Wilkens CA (2011) Near optimal online algorithms and fast approximation algorithms for resource allocation problems. Chen Y, Roughgarden T, eds. *Proc. 12th ACM Conf. Electronic Commerce* (ACM, New York), 29–38.

Feldman J, Henzinger M, Korula N, Mirrokni VS, Stein C (2010) Online stochastic packing applied to display ad allocation. de Berg M, Meyer U, eds. *Proc. 18th Annual Eur. Conf. Algorithms: Part I* (Springer-Verlag, Berlin, Heidelberg), 182–194.

Gallego G, Iyengar G, Phillips R, Dubey A (2004) Managing flexible products on a network. CORC Technical Report TR-2004-01, Columbia University, New York.

Gallego G, Li A, Truong VA, Wang X (2016) Online bipartite matching. Working paper, Hong Kong University of Science and Technology, Hong Kong.

Goel A, Mahdian M, Nazerzadeh H, Saberi A (2010) Advertisement allocation for generalized second-pricing schemes. *Oper. Res. Lett.* 38(6):571–576.

Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-time optimization of personalized assortments. *Management Sci.* 60(6):1532–1551.

Hu X, Caldentey R, Vulcano G (2013) Revenue sharing in airline alliances. *Management Sci.* 59(5):1177–1195.

Jasin S, Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Math. Oper. Res.* 37(2):313–345.

Jasin S, Kumar S (2013) Analysis of deterministic LP-based booking limit and bid price controls for revenue management. *Oper. Res.* 61(6):1312–1320.

Kesselheim T, Tönnis A, Radke K, Vöcking B (2014) Primal beats dual on online packing LPs in the random-order model. Shmoys D, ed. *Proc. 46th Annual ACM Sympos. Theory Comput.* (ACM, New York), 303–312.

Kirshner SN, Nediak M (2015) Scalable dynamic bid prices for network revenue management in continuous time. *Production Oper. Management* 24(10):1621–1635.

Kunnumkal S, Talluri K (2016a) On a piecewise-linear approximation for network revenue management. *Math. Oper. Res.* 41(1):72–91.

Kunnumkal S, Talluri K (2016b) Technical note: A note on relaxations of choice network revenue management dynamic program. *Oper. Res.* 64(1):158–166.

Kunnumkal S, Topaloglu H (2008) A refined deterministic linear program for the network revenue management problem with customer choice behavior. *Naval Res. Logist. Quart.* 55(6):563–580.

Kunnumkal S, Topaloglu H (2010a) Computing time-dependent bid prices in network revenue management problems. *Transportation Sci.* 44(1):38–62.

Kunnumkal S, Topaloglu H (2010b) A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior. *Production Oper. Management* 19(5):575–590.

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

Maglaras C, Meissner J (2006) Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management* 8(2):136–148.

Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *J. ACM* 54(5):22:1–22:19.

Meissner J, Strauss A, Talluri K (2012) An enhanced concave program relaxation for choice network revenue management. *Production Oper. Management* 22(1):71–87.

Mendez-Diaz I, Bront JJM, Vulcano G, Zabala P (2010) A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Appl. Math.* 36:383–390.

Molinaro M, Ravi R (2014) The geometry of online packing linear programs. *Math. Oper. Res.* 39(1):46–59.

Rusmevichientong P, Sumida M, Topaloglu H (2020) Dynamic assortment optimization for reusable products with random usage durations. *Management Sci.* ePub ahead of print January 29, <https://doi.org/10.1287/mnsc.2019.3346>.

Simpson RW (1989) *Using Network Flow Techniques to Find Shadow Prices for Market Demands and Seat Inventory Control* (MIT Press, Cambridge, MA).

Strauss AK, Talluri K (2017) Tractable consideration set structures for network revenue management. *Production Oper. Management* 26(7):1359–1368.

Talluri K (2014) New formulations for choice network revenue management. *INFORMS J. Comput.* 26(2):401–413.

Talluri KT, van Ryzin GJ (1998) An analysis of bid-price controls for network revenue management. *Management Sci.* 44(11-part-1): 1577–1593.

Talluri KT, van Ryzin GJ (1999) A randomized linear programming method for computing network bid prices. *Transportation Sci.* 33(2):207–216.

Talluri KT, van Ryzin GJ (2005) *The Theory and Practice of Revenue Management* (Kluwer Academic Publishers, Boston).

Tan B, Srikant R (2012) Online advertisement, optimization and stochastic networks. *IEEE Trans. Automatic Control* 57(11): 2854–2868.

Tong C, Topaloglu H (2013) On the approximate linear programming approach for network revenue management problems. *INFORMS J. Comput.* 26(1):121–134.

Topaloglu H (2008) A stochastic approximation method to compute bid prices in network revenue management problems. *INFORMS J. Comput.* 20(4):596–610.

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

van Ryzin G, Vulcano G (2008a) Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing Service Oper. Management* 10(3):448–467.

van Ryzin G, Vulcano G (2008b) Simulation-based optimization of virtual nesting controls for network revenue management. *Oper. Res.* 56(4):865–880.

Vossen TWM, Zhang D (2015a) A dynamic disaggregation approach to approximate linear programs for network revenue management. *Production Oper. Management* 24(3):469–487.

Vossen TWM, Zhang D (2015b) Reductions of approximate linear programs for network revenue management. *Oper. Res.* 63(6): 1352–1371.

Wang X, Truong VA, Bank D (2016) Online advance admission scheduling for services, with customer preferences. Working paper, Columbia University, New York.

Williamson EL (1992) Airline network seat control. PhD thesis, Massachusetts Institute of Technology, Cambridge.

Zhang D (2011) An improved dynamic programming decomposition approach for network revenue management. *Manufacturing Service Oper. Management* 13(1):35–52.

Zhang D, Adelman D (2009) An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Sci.* 43(3):381–394.

Zhang D, Cooper WL (2005) Revenue management for parallel flights with customer choice behavior. *Oper. Res.* 53(3):415–431.

Zhang D, Cooper WL (2009) Pricing substitutable flights in airline revenue management. *Eur. J. Oper. Res.* 197(3):848–861.

**Yuhang Ma** is a PhD student in the School of Operations Research and Information Engineering at Cornell University. Her research interests are assortment optimization and revenue management.

**Paat Rusmevichientong** is a professor of data sciences and operations at the University of Southern California Marshall School of Business. His research interests include revenue management, choice modeling, pricing, assortment optimization, and approximate dynamic programming.

**Mika Sumida** is a PhD student in the School of Operations Research and Information Engineering at Cornell University. Her research interests include algorithm design, resource allocation, and revenue management with applications in online marketplaces and the sharing economy.

**Huseyin Topaloglu** is a professor in the School of Operations Research and Information Engineering at Cornell University. His research focuses on revenue management, assortment optimization, transportation logistics, and approximate dynamic programming.