A Randomized Linear Programming Method for Network Revenue Management with Product-Specific No-Shows

Sumit Kunnumkal
Indian School of Business, Gachibowli, Hyderabad, 500032, India
sumit_kunnumkal@isb.edu

Kalyan Talluri
ICREA and Universitat Pompeu Fabra, Barcelona, 08005, Spain
kalyan.talluri@upf.edu

Huseyin Topaloglu
School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu

May 20, 2011

Abstract

Revenue management practices often include overbooking capacity to account for customers who make reservations but do not show up. In this paper, we consider the network revenue management problem with no-shows and overbooking, where the show-up probabilities are specific to each product. No-show rates differ significantly by product (for instance, each itinerary and fare combination for an airline) as sale restrictions and the demand characteristics vary by product. However, models that consider no-show rates by each individual product are difficult to handle as the state-space in dynamic programming formulations (or the variable space in approximations) increases significantly. In this paper, we propose a randomized linear program to jointly make the capacity control and overbooking decisions with product-specific no-shows. We establish that our formulation gives an upper bound on the optimal expected total profit and our upper bound is tighter than a deterministic linear programming upper bound that appears in the existing literature. Furthermore, we show that our upper bound is asymptotically tight in a regime where the leg capacities and the expected demand is scaled linearly with the same rate. We also describe how the randomized linear program can be used to obtain a bid price control policy. Computational experiments indicate that our approach is quite fast, able to scale to industrial problems and can provide significant improvements over standard benchmarks.

Keywords: Network revenue management, linear programming, simulation, overbooking, no-shows.

Introduction

Revenue management controls the sale of a perishable product to a heterogeneous population of customers with different valuations for the same product. The physical product could be hotel rooms, airline seats or media advertising slots sold at a specific price with sale restrictions. Typically, the products are defined over a network and firms have to control the sale of multiple products that consume different bundles of resources. For example, airlines products are itineraries that span different flight legs, while hotel customers stay for multiple nights using the inventory over different days.

Revenue management practices often include overbooking capacity to account for customers who make reservations but do not show up. No-show rates differ significantly by product as sale restrictions, time-of-purchase and demand characteristics vary by product. However, models that consider no-show rates by each individual product are difficult to handle as the state-space in dynamic programming formulations or the variable space in approximations increases significantly. In this paper, we develop tractable models for jointly making the capacity control and overbooking decisions in network revenue management problems with product-specific no-shows.

While the problem setting is applicable to a wide range of industries, we use airline terminology throughout the paper for concreteness. Thus, products are itinerary and fare-class combinations, where a fare-class represents a revenue management fare-product corresponding to a fare combined with some restrictions. The resources are the seats on the flight legs. The airline gets requests for the products for a future date and it has to decide in real time which of these product requests to accept and which to reject. In making this decision, the airline not only has to consider the uncertainty in the customer arrivals but also the fact that not all customers who make purchases show up at the time of departure of the flight. Because of these no-shows, the airline may choose to overbook and accept more itinerary requests than the capacity of the flight leg. However, by overbooking it also runs the risk of denying seats to customers with reservations, if the number that show up at the time of flight departure exceeds the seating capacity of the flight leg. Since the number of dimensions of the state variable quickly gets large in the dynamic programming formulations of practical problem instances, computing the optimal capacity control policy is generally intractable and one has to resort to heuristics that approximate the solution to the dynamic problem.

In this paper, we propose a randomized linear programming method for jointly making the capacity control and overbooking decisions in network revenue management problems with product-specific no-shows. In our approach, we generate samples of the demands for the products and the show-ups, and solve a two-stage linear stochastic program, where the first stage decisions are the number of reservation requests to accept and the second stage decisions are the number of denied boardings. It thus extends the randomized linear programming method of Talluri and van Ryzin (1999) to overbooking decisions with product-specific no-shows.

The approach that we propose has a number of appealing features. To begin with, it yields an upper bound on the optimal expected total profit and we show that this upper bound is tighter than the one obtained by the deterministic linear program of Bertsimas and Popescu (2003). Having a tight upper bound on the optimal expected total profit becomes valuable when trying to assess the performance of approximate control policies. Moreover, there is evidence that methods obtaining tight upper bounds also tend to yield policies with good profit performance; see for example Topaloglu (2009*b*) and Talluri (2009). Also, by using samples of the random variables rather than their expected values, we are able to better capture the stochastic nature of the network revenue management problem; for instance, our model can accommodate arbitrary probability distributions for the demand random variables. As a result, we expect the randomized linear program to yield good control policies.

Another appealing feature of our approach is that the deterministic linear program proposed by Bertsimas and Popescu (2003) is, to our knowledge, the only tractable method to obtain upper bounds on the expected total profit. However, this approach assumes that all random variables take on their expected values and it does not consider the random nature of the demand and show-ups. Our randomized linear program closes this gap by providing a tractable method to obtain upper bounds on the optimal expected total profit, while considering the random nature of the demand and show-ups. Finally, the method that we propose requires solving only linear programs, leveraging the speed, robustness and parallelization of modern solvers. Therefore, our approach can be easily implemented using commercially available solvers and modeling languages with minimal customized coding, an attractive proposition for practicing revenue managers.

A control policy that is widely used in practice is a bid price control. In bid price control, we have a bid price for each flight leg, which is essentially a proxy for the expected marginal value of capacity for that flight leg. In this case, a bid price policy accepts a product request only if its fare exceeds the sum of the bid prices of the flight legs used by it. Traditionally, bid prices are computed by using the optimal values of the dual variables of the seat availability constraints in the deterministic linear programming formulation. By using samples of the random variables, our randomized linear program is able to obtain better estimates of the expected marginal value of capacity than the deterministic linear program. Consequently, we expect better profit performance from our randomized linear program. This is indeed the case in our computational experiments.

To summarize, we make the following research contributions in this paper. 1) We propose a new method to jointly make the capacity control and overbooking decisions in network revenue management problems with no-shows. Our method is sampling-based and thus flexible enough to model a wide variety of probability distributions for the random demand arrival processes. 2) We show that our method yields an upper bound on the optimal expected total profit and this upper bound is tighter than that obtained by the deterministic linear program. 3) We show that the upper bound on the expected total profit provided by our approach is asymptotically tight as we scale the capacities on the flight legs and the expected amount of demand linearly with the same rate. Furthermore, our proof also implies that the upper bound provided by the standard deterministic linear program is tight in the same asymptotic regime, which is a result that has not been established within the overbooking context. 4) Our computational experiments compare our randomized linear programming method with existing methods for a range of demand scenarios to test the upper bounds and revenue performance. Overall, our approach is fast, easy to implement, numerically robust and scalable even for industrial problems and computational experiments indicate that the bid price policy obtained by our method can generate significantly higher profits than standard benchmarks.

1 Related work

A commonly used method to make the capacity control decisions over an airline network is the deterministic linear program; see Simpson (1989) and Williamson (1992) for a model without overbooking and Bertsimas and Popescu (2003) for an extension to handle no-shows and cancellations. The underlying assumptions of the deterministic linear program are that all random quantities take on their expected values and the number of reservation requests accepted and the number of denied boardings can be fractional. In this linear program, there is one constraint for each flight leg, which ensures that the total number of passengers that are eventually boarded does not exceed the capacity of the flight. The optimal values of the dual variables associated with the flight leg capacity constraints are usually used as proxies for the expected marginal values of each unit of capacity. These dual variables are used to control sales through a bid price policy. Erdelyi and Topaloglu (2009) show that the optimal objective value of the deterministic linear program is an upper bound on the optimal expected total profit.

Kleywegt (2001) develops a joint pricing and overbooking model over an airline network assuming that the reservation requests are deterministic. He solves the model by using duality and decomposition ideas. Karaesmen and van Ryzin (2004*b*) describe a capacity allocation and overbooking model that is useful when dealing with multiple flight legs that can serve as substitutes of each other. Similar to ours, their approach has a two stage stochastic programming flavor, but it is not clear how to extend their approach to general airline networks. Karaesmen and van Ryzin (2004*a*) develop a joint capacity allocation and overbooking model by using the deterministic linear program to estimate the revenue from the accepted reservations. Their approach uses a sequence of approximating assumptions, where the authors assume that the overbooking cost of a passenger can be prorated over the different flight legs and a reservation for an itinerary shows up for each flight in the itinerary independently.

Kunnumkal and Topaloglu (2008) and Erdelyi and Topaloglu (2009) formulate the overbooking problem over a network as a dynamic program with a high dimensional state vector. They approximate the value functions by separable functions and use the separable approximations to obtain control policies. Erdelyi and Topaloglu (2010) work with the same dynamic programming formulation and propose a dynamic programming decomposition approach to decompose the network problem into a number of single flight leg problems. However, since the single flight leg problem is intractable when there are product-specific no-shows, they have to resort to heuristics to approximately solve the single flight leg problems. Kunnumkal and Topaloglu (2011) propose a stochastic approximation algorithm to obtain bid prices. Their approach visualizes the expected total profit as a function of the bid prices and uses sample path-based gradients of the expected total profit function to search for a good set of bid prices in a stochastic approximation algorithm. This stochastic approximation algorithm eventually computes high quality bid prices, but its run times can be long, it involves tuneable parameters that need to be adjusted with trial and error and it does not have a well-defined stopping condition. We use the approaches proposed by Karaesmen and van Ryzin (2004*a*) and Kunnumkal and Topaloglu (2011) as benchmarks in our computational experiments. Finally, the book by Talluri and van Ryzin (2004) contains background and details on revenue management, specifically the chapters on overbooking and network revenue management.

The rest of the paper is organized as follows. Section 2 formulates the network revenue management problem with no-shows as a dynamic program. Section 3 describes the deterministic linear program proposed by Bertsimas and Popescu (2003). Section 4 builds on this linear program to develop our randomized linear program. Section 5 establishes the asymptotic tightness of the upper bounds provided by our randomized linear program. Section 6 presents our computational experiments.

2 Problem Formulation

An airline network consists of a set of m flights and n products (itinerary-fare class combinations). The physically available capacity on flight i is c_i. The booking horizon consists of time periods $1, 2, \ldots, \tau$ and all flights depart at time period $\tau + 1$. We make the standard assumption that the time periods are fine enough so that there is at most one product request in each time period. The probability that there is a request for product j in time period t is p_{jt}. Product j has a revenue f_j associated with it and we denote the flights in the product by $i \in j$. Throughout, we index flights by i, products by j and time by t. We have to decide, in an on-line fashion, whether to accept a request for product j to obtain revenue f_j or reject the request in anticipation of future higher revenue requests.

When we decide to accept a request, it becomes a reservation. Not all reservations show up at the time of departure of the flight legs and we let q_{jt} be the probability that a reservation for product j made at time period t shows up at the time of departure. Knowing that only a portion of the reservations show up at the departure time of the flight legs, the airline overbooks. That is, it accepts more reservations than the capacity of the flight. If more reservations show up than the capacity of the flight, then the airline has to decide which of these reservations to deny boarding.

Product j can potentially consume (if a reservation for it is allowed boarding) one unit of capacity on all flights $i \in j$, and we let $a_{ij} = 1$ if $i \in j$ and $a_{ij} = 0$ otherwise. On the other hand, if we deny boarding to a reservation for product j, then we incur a denied-service penalty cost of θ_j. We have to decide which of the product requests to accept during the booking period and which of the accepted reservations to deny boarding at the time of departure of the flight legs, the goal being to maximize expected total profits. The expected total profit is given by the difference between the expected revenue obtained by accepting the product requests and the expected penalty cost of denied service.

We make the following assumptions in our model. We assume that the demands for the different products and the show-up decisions of the different reservations are independent of each other and across time periods. We assume that there are no cancellations (reservations that cancel *prior* to τ) and we do not give refunds to the reservations that do not show up at the departure time. This is for ease of notation and all the development in the paper goes through with minor modifications in the presence of cancellations and refunds, provided that the cancellation decisions are independent across reservations and time periods. Finally, we assume that $f_j \leq \theta_j q_{jt}$ for each product j. This is again without loss of generality since if we have $f_j > \theta_j q_{jt}$, then we can always make a profit in expectation by accepting a request for product j at time period t and denying it boarding at the time of departure.

The decision problem is to determine (online, without knowing future demands) the product requests to accept, and at departure time, to determine which of the confirmed reservations to deny boarding. We let x_{jt} be a binary variable equal to 1 if we accept a request for product j at time period t and 0 otherwise. Since a reservation for product j at time period t shows up with probability q_{jt}, the show-up decision for a product j request at time period t can be written as $S_{jt}x_{jt}$, where S_{jt} is a Bernoulli random variable with success probability q_{jt}. Letting s_{jt} be a realization of S_{jt} and $s = \{s_{jt} : \forall j, t\}$ and $x = \{x_{jt} : \forall j, t\}$, we represent the total denied-service cost as the solution to

$$\Pi(s, x) = \quad \min \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} \theta_j w_{jt} \tag{1}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} a_{ij} [s_{jt} x_{jt} - w_{jt}] \leq c_i \quad \forall i \tag{2}$$

$$w_{jt} \leq s_{jt} x_{jt} \quad \forall j, t \tag{3}$$

$$w_{jt} \in \{0, 1\} \quad \forall j, t. \tag{4}$$

In the above linear integer program, the decision variable w_{jt} indicates whether or not we deny boarding to the confirmed reservation for product j purchased at time period t. The first set of constraints ensures that the total number of reservations that are eventually allowed to board does not exceed the capacity of the flight. The second set of constraints ensures that we can deny boarding only if the corresponding reservation shows up at the time of departure. Notice that we assume the airline has the ability to first observe the show-up demand for all the reservations and then decide which reservations to deny boarding. This is a somewhat stylized model of the actual deny-service process, where the airline may have to make the deny-service decisions online with partial information and many more restrictions. Nevertheless, following Bertsimas and Popescu (2003) and Erdelyi and Topaloglu (2009), we use problem (1)-(4) to approximately capture the overall denied-service costs.

Let $\overline{x}_t = \{x_{js} : \forall j, s = 1, \ldots, t-1\}$ denote the state of reservations in the system at the beginning of time period $t = 2, \ldots, \tau + 1$ and let $\overline{x}_t \oplus e_j$ denote the state of reservations in the system at the beginning of time period $t+1$ given that we had $\overline{x}_t$ reservations in the system at the beginning of time period t and we accepted a request for product j at time period t. Similarly, let $\overline{x}_t \oplus 0$ denote the state of reservations in the system at the beginning of time period $t+1$ given that we had $\overline{x}_t$ reservations in the system at the beginning of time period t and we did not accept a request for any product at time period t. The value functions $\{V_t(\cdot) : \forall t\}$ are given by the optimality equation

$$V_t(\overline{x}_t) \quad = \quad \sum_{j=1}^{n} p_{jt} \max \big\{ f_j + V_{t+1}(\overline{x}_t \oplus e_j), V_{t+1}(\overline{x}_t \oplus 0) \big\} + \big[1 - \sum_{j=1}^{n} p_{jt} \big] V_{t+1}(\overline{x}_t \oplus 0) \tag{5}$$

with the boundary condition that $V_{\tau+1}(\overline{x}_{\tau+1}) = -\mathbb{E}\{\Pi(S, \overline{x}_{\tau+1})\}$. In this case, $V_1(\bar{0})$ denotes the optimal expected total profit at the beginning of the booking horizon, where $\bar{0}$ is an n-dimensional vector of zeros representing the fact that we start with 0 reservations. If the state of reservations at the beginning of time period t is $\overline{x}_t$, then it follows from (5) that it is optimal to accept a request for product j at time period t provided $f_j + V_{t+1}(\overline{x}_t \oplus e_j) \geq V_{t+1}(\overline{x}_t \oplus 0)$.

In the optimality equation (5), the dimensionality of the state space increases exponentially with the number of products and the number of time periods. Therefore, solving this optimality equation quickly becomes computationally intractable. In the next two sections, we describe approximate methods that can be used to jointly make the capacity control and overbooking decisions.

3 Deterministic Linear Program

The deterministic linear program with overbooking and no-shows, proposed by Bertsimas and Popescu (2003), is given as

$$z_{DLP} = \max \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} [f_j y_{jt} - \theta_j w_{jt}] \tag{6}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} a_{ij} [q_{jt} y_{jt} - w_{jt}] \leq c_i \quad \forall i \tag{7}$$

$$y_{jt} \leq p_{jt} \quad \forall j, t \tag{8}$$

$$w_{jt} \leq q_{jt} y_{jt} \quad \forall j, t \tag{9}$$

$$y_{jt}, w_{jt} \geq 0 \quad \forall j, t, \tag{10}$$

where y_{jt} represents the number of requests accepted for product j at time period t and w_{jt} represents the number of these reservations that are denied boarding. The deterministic linear program assumes that of the y_{jt} requests accepted, exactly $q_{jt} y_{jt}$ requests show up at the time of departure. The first set of constraints ensures that the numbers of reservations that we allow boarding do not exceed the capacities of the flight legs. The second set of constraints ensures that the number of requests for product j that we accept at time period t does not exceed the expected number of product requests. The third set of constraints ensures that the numbers of reservations that we deny boarding do not exceed the expected numbers of reservations that show up at the time of departure. Thus, problem (6)-(10) assumes that all random quantities take on their expected values.

There are two uses of problem (6)-(10). First, Erdelyi and Topaloglu (2009) show that the optimal objective value of problem (6)-(10) provides an upper bound on the optimal expected total profit. That is, we have $V_1(\bar{0}) \leq z_{DLP}$. Upper bounds are useful when assessing the optimality gap of suboptimal policies. Second, we can use the optimal values of the dual variables corresponding to the flight leg capacity constraints as the bid prices. That is, letting $\mu = \{\mu_i : \forall i\}$ be the optimal values of the dual variables corresponding to constraints (7), we accept a request for product j at time period t if

$$f_j \geq q_{jt} \sum_{i=1}^{m} a_{ij} \mu_i. \tag{11}$$

This is the bid price policy used by Bertsimas and Popescu (2003). Noting that μ_i is an estimate of the marginal value of capacity on flight leg i and q_{jt} is the show-up probability, we can interpret the decision rule in (11) as accepting a product request only if its revenue exceeds the total expected marginal value of the capacities that it uses.

4 A Randomized Linear Program

The randomized linear program proposed by Talluri and van Ryzin (1999) is a tractable and attractive approach for making the capacity control decisions in the absence of no-shows. It is only natural that we try to extend the approach for jointly making the capacity control and overbooking decisions when we have no-shows. We propose solving the following optimization problem to obtain an upper bound on the optimal expected total profit. Let D_{jt} be the random variable that denotes the number of requests for product j at time period t. Note that we have $\mathbb{E}\{D_{jt}\} = p_{jt}$. As before, let S_{jt} be a Bernoulli(q_{jt}) random variable, s_{jt} be a realization of S_{jt} and $S = \{S_{jt} : \forall j, t\}$, $s = \{s_{jt} : \forall j, t\}$. To compute the denied service cost as a function of the show-ups and accepted reservations, we let

$$\tilde{\Pi}(s, y) = \quad \min \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} \theta_j w_{jt} \tag{12}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} a_{ij} [s_{jt} y_{jt} - w_{jt}] \leq c_i \quad \forall i \tag{13}$$

$$w_{jt} \leq s_{jt} y_{jt} \quad \forall j, t \tag{14}$$

$$w_{jt} \geq 0 \quad \forall j, t, \tag{15}$$

in which case, we approximate the expected total profit under a realization d of demand arrivals as

$$z_{RLP}(d) = \quad \max \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} f_j y_{jt} - \mathbb{E}\{\tilde{\Pi}(S, y)\} \tag{16}$$

$$\text{subject to} \quad 0 \leq y_{jt} \leq d_{jt} \quad \forall j, t, \tag{17}$$

where $d = \{d_{jt} : \forall j, t\}$ is a realization of the random variables $D = \{D_{jt} : \forall j, t\}$. Note that $\tilde{\Pi}(\cdot, \cdot)$ is the linear programming relaxation of $\Pi(\cdot, \cdot)$ and constraints (13)-(15) have the same interpretation as constraints (2)-(4). Note also that the optimal objective values of problems (12)-(15) and (16)-(17) respectively depend on the realizations of the random variables S and D. Finally, we can interpret $z_{RLP}(d)$ as being the optimal profit when we make the accept or reject decisions for the product requests after observing a realization d of the demands over the whole booking horizon.

Letting $z_{RLP} = \mathbb{E}\{z_{RLP}(D)\}$, we show below that z_{RLP} is an upper bound on the optimal expected total profit and this upper bound is tighter than that obtained from the deterministic linear program. We begin with the following observations.

Lemma 1 *1)* $\tilde{\Pi}(s, y)$ *is a convex function of* s *for a fixed* y*.*
2) $\tilde{\Pi}(s, y)$ *is a convex function of* y *for a fixed* s*.*
3) $\tilde{\Pi}(s, y) \leq \Pi(s, y)$ *for all* s, y*.*

Proof Parts (1) and (2) follow from standard linear programming theory as we can write $\tilde{\Pi}(s, y)$ with all $s_{jt}y_{jt}$ terms on the right-hand side. Part (3) follows from the fact that $\tilde{\Pi}(\cdot, \cdot)$ is the linear programming relaxation of $\Pi(\cdot, \cdot)$.

□

The next result shows that z_{RLP} gives an upper bound on the optimal expected total profit.

Proposition 2 *We have $V_1(\bar{0}) \leq z_{RLP}$.*

Proof Let $\hat{\pi}$ be an optimal policy to decide whether to accept or reject a product request at a time period and let $x^{\hat{\pi}}(d) = \{x^{\hat{\pi}}_{jt}(d) \in \{0, 1\} : \forall j, t\}$ denote the number of product requests that this optimal policy accepts along a sample path, where the argument d emphasizes that the number of accepted product requests depends on the sample $d = \{d_{jt} : \forall j, t\}$. We have

$$V_1(\bar{0}) = \mathbb{E}\Big\{\sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j x^{\hat{\pi}}_{jt}(D) - \mathbb{E}\{\Pi(S, x^{\hat{\pi}}(D)) \,|\, D\}\Big\} \leq \mathbb{E}\Big\{\sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j x^{\hat{\pi}}_{jt}(D) - \mathbb{E}\{\tilde{\Pi}(S, x^{\hat{\pi}}(D)) \,|\, D\}\Big\}$$
$$\leq \mathbb{E}\{z_{RLP}(D)\} = z_{RLP},$$

where the first equality follows from the optimality of the policy $\hat{\pi}$, the first inequality uses the third part of Lemma 1 and the last inequality holds since $x^{\hat{\pi}}(d)$ is feasible but not necessarily optimal to problem (16)-(17).

□

We next show that the upper bound obtained by the randomized linear program is tighter than that obtained by the deterministic linear program. We need the following intermediate result. Let

$$\begin{aligned} \zeta(d, s) = \qquad \max \quad & \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j y_{jt} - \tilde{\Pi}(s, y) \\ \text{subject to} \quad & 0 \leq y_{jt} \leq d_{jt} \quad \forall j, t, \end{aligned}$$

where $d = \{d_{jt} : \forall j, t\}$ and $s = \{s_{jt} : \forall j, t\}$. Note that we have $\zeta(p, q) = z_{DLP}$, where $p = \{p_{jt} : \forall j, t\}$ and $q = \{q_{jt} : \forall j, t\}$ and we use the fact that $-\min\{x\} = \max\{-x\}$.

Lemma 3 *For a fixed s, $\zeta(d, s)$ is a concave function of d.*

Proof Let $0 \leq \alpha \leq 1$, and y^1 and y^2 be the optimal solutions to $\zeta(d^1, s)$ and $\zeta(d^2, s)$, respectively. We have that $\alpha y^1 + (1-\alpha)y^2$ is feasible to $\zeta(\alpha d^1 + (1-\alpha)d^2, s)$. Moreover, by the second part of Lemma 1, $\tilde{\Pi}(s, \alpha y^1 + (1-\alpha)y^2) \leq \alpha\tilde{\Pi}(s, y^1) + (1-\alpha)\tilde{\Pi}(s, y^2)$. It follows that $\zeta(\alpha d^1 + (1-\alpha)d^2, s) \geq \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j[\alpha y^1_{jt} + (1-\alpha)y^2_{jt}] - \tilde{\Pi}(s, \alpha y^1 + (1-\alpha)y^2) \geq \alpha\zeta(d^1, s) + (1-\alpha)\zeta(d^2, s)$.

□

The following proposition shows that the randomized linear program obtains a tighter upper bound than the deterministic linear program.

Proposition 4 *We have $z_{RLP} \leq z_{DLP}$.*

Proof The first part of Lemma 1 and Jensen's inequality imply that we have that $\mathbb{E}\{\tilde{\Pi}(S,y)\} \geq \tilde{\Pi}(\mathbb{E}\{S\},y) = \tilde{\Pi}(q,y)$, where $q = \{q_{jt} : \forall j,t\}$ and we use the fact that S_{jt} is Bernoulli(q_{jt}). Therefore, we have

$$\begin{aligned} \zeta(d,q) \;\geq \quad & \max \quad \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j y_{jt} - \mathbb{E}\{\tilde{\Pi}(S,y)\} \quad = \; z_{RLP}(d) \\ & \text{subject to} \quad 0 \leq y_{jt} \leq d_{jt} \quad \forall j,t, \end{aligned}$$

where the equality is by the definition of $z_{RLP}(d)$ in problem (16)-(17). It follows that

$$z_{RLP} = \mathbb{E}\{z_{RLP}(D)\} \leq \mathbb{E}\{\zeta(D,q)\} \leq \zeta(\mathbb{E}\{D\},q) = z_{DLP},$$

where the second inequality uses Lemma 3 and Jensen's inequality and the last equality holds since $\mathbb{E}\{D_{jt}\} = p_{jt}$.

□

As computing the upper bound $\mathbb{E}\{z_{RLP}(D)\}$ analytically is difficult, we propose a simulation-based optimization scheme to approximate $\mathbb{E}\{z_{RLP}(D)\}$. We generate K samples of the random variables $D = \{D_{jt} : \forall j,t\}$ using Monte Carlo simulation. For each sample, we generate a further L samples of the random variables $S = \{S_{jt} : \forall j,t\}$. That is, letting $d^k = \{d^k_{jt} : \forall j,t\}$ denote the kth sample of D, we generate samples $s^{kl} = \{s^{kl}_{jt} : \forall j,t\}$ for $l = 1,\ldots,L$. We solve

$$z^k_{RLP} = \quad \max \quad \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j y^k_{jt} - \frac{1}{L}\sum_{l=1}^{L}\sum_{j=1}^{n}\sum_{t=1}^{\tau} \theta_j w^{kl}_{jt} \tag{18}$$

$$\text{subject to} \quad \sum_{j=1}^{n}\sum_{t=1}^{\tau} a_{ij}[s^{kl}_{jt} y^k_{jt} - w^{kl}_{jt}] \leq c_i \quad \forall i,l \tag{19}$$

$$w^{kl}_{jt} \leq s^{kl}_{jt} y^k_{jt} \quad \forall j,t,l \tag{20}$$

$$y^k_{jt} \leq d^k_{jt} \quad \forall j,t \tag{21}$$

$$y^k_{jt} \geq 0 \quad \forall j,t \tag{22}$$

$$w^{kl}_{jt} \geq 0 \quad \forall j,t,l \tag{23}$$

for the kth demand sample and use $\sum_{k=1}^{K} z^k_{RLP}/K$ as an estimate of z_{RLP}. Furthermore, letting $\{\rho^{kl}_i : \forall i,l\}$ be the optimal values of the dual variables corresponding to constraints (19), we use $\rho_i = \sum_{k=1}^{K}\sum_{l=1}^{L} \rho^{kl}_i/K$ as the bid price for flight leg i and we accept a request for product j at time period t only if

$$f_j \geq \sum_{i=1}^{m} a_{ij}\rho_i. \tag{24}$$

In Appendix A, we show that z_{RLP} is a concave function of the flight leg capacities by showing that it has a subgradient. Furthermore, we show that $\rho = \{\rho_i : \forall i\}$ is an estimate of the subgradient that we obtain from our sample. Therefore, we can interpret ρ_i as an estimate of the expected marginal value of capacity on flight leg i.

Note that comparing our acceptance rule in (24) with the acceptance rule in (11) of Bertsimas and Popescu (2003), (24) does not have the term q_{jt} in it. This leads to increased robustness when the actual deny-service process does not have full knowledge of no-shows. The acceptance rule in (11) can potentially accept very low fares which have low show-up probabilities. This is fine if we know that we can reject them in case of excess reservations, but this can turn out to be dangerous if at departure time we have to make deny decisions without knowing all the no-shows.

We close this section with two observations. First, since we have at most one product request arriving at each time period, we have at most one of the $\{d^k_{jt} : \forall j\}$ being nonzero for each time period t. Consequently constraints (21) imply that at most τ of $\{y^k_{jt} : \forall j, t\}$ are nonzero. Constraints (20) then imply that for each $l = 1, \ldots, L$, we have at most τ of $\{w^{kl}_{jt} : \forall j, t\}$ nonzero. Therefore, problem (18)-(23) can be reduced to a linear program that has $\tau + \tau L$ variables and $(m + \tau)L + \tau$ constraints. Second, it is possible to come up with an equivalent formulation of problem (18)-(23) by aggregating the denied boarding decisions for each product. In particular, we let $w^{kl}_j = \sum_{t=1}^{\tau} w^{kl}_{jt}$ and replace constraints (19) with the constraints $\sum_{j=1}^{n} a_{ij} [\sum_{t=1}^{\tau} s^{kl}_{jt} y^k_{jt} - w^{kl}_j] \leq c_i$ and constraints (20) with the constraints $w^{kl}_j \leq \sum_{t=1}^{\tau} s^{kl}_{jt} y^k_{jt}$. The resulting formulation has $\tau + nL$ decision variables and $(m + n)L + \tau$ constraints. The alternative formulation is more attractive when $n \leq \tau$.

5 Asymptotic Optimality

In this section, we show that the upper bounds on the optimal expected total profit provided by our randomized linear program is asymptotically tight as the capacities on the flight legs and the expected demand scales linearly with the same rate. For models where all reservations show up at the departure time and overbooking is not possible, similar results have been shown for the deterministic linear program in Talluri and van Ryzin (1998) and for the randomized linear program in Topaloglu (2009*a*). Our results in this section can be visualized as generalizations of those results to the overbooking setting, but these generalizations are nontrivial due to the difficulties brought out by the possibility of overbooking.

We define a family of network revenue management problems $\{\mathcal{P}^\kappa : \kappa \in \mathbb{Z}_+\}$ indexed by the parameter κ with the following properties. (1) Problem $\mathcal{P}^\kappa$ has $\kappa\tau$ time periods in the booking horizon. (2) In problem $\mathcal{P}^\kappa$, the probability that we get a request for product j at time period t is given by $p_{j\lceil t/\kappa \rceil}$, where $\lceil \cdot \rceil$ is the round up function. For notational brevity, we let $p^\kappa_{jt} = p_{j\lceil t/\kappa \rceil}$. (3) In problem $\mathcal{P}^\kappa$, a reservation for product j accepted at time period t shows up with probability $q_{j\lceil t/\kappa \rceil}$. Similar to p^κ_{jt}, for notational brevity, we let $q^\kappa_{jt} = q_{j\lceil t/\kappa \rceil}$. As before, we assume that the arrivals and show-up decisions are independent across time periods. Furthermore, the show-up decisions of different reservations are independent. (4) In problem $\mathcal{P}^\kappa$, the capacity on flight leg i is κc_i.

With this definition of problem $\mathcal{P}^\kappa$, we observe that problem $\mathcal{P}^1$ corresponds to the original network revenue management problem that we have been working with throughout the paper. The capacities on the flight legs in problem $\mathcal{P}^\kappa$ are κ times the capacities on the flight legs in problem $\mathcal{P}^1$. Similarly, the length of the booking horizon in problem $\mathcal{P}^\kappa$ is κ times the length of the booking horizon in problem $\mathcal{P}^1$. In addition, the probability of getting a request for product j at time periods $1, 2, \ldots, \kappa$ in problem $\mathcal{P}^\kappa$ is the same as the probability of getting a request for product j at time period 1 in problem $\mathcal{P}^1$. A similar observation holds for blocks of successive κ time periods over the booking horizon of problem $\mathcal{P}^\kappa$. In particular, for a fixed $\ell = 1, \ldots, \tau$, the probability of getting a request for product j at time periods $1+(\ell-1)\kappa, 2+(\ell-1)\kappa, \ldots, \kappa+(\ell-1)\kappa$ in problem $\mathcal{P}^\kappa$ is the same as the probability of getting a request for product j at time period ℓ in problem $\mathcal{P}^1$. Thus, the expected total demand for product j in problem $\mathcal{P}^\kappa$ is

$$\sum_{t=1}^{\kappa\tau} p_{jt}^\kappa = \sum_{t=1}^{\kappa\tau} p_{j\lceil t/\kappa\rceil} = \kappa \sum_{t=1}^{\tau} p_{jt},$$

which implies that the expected total demand for product j in problem $\mathcal{P}^\kappa$ is κ times the expected total demand for product j in problem $\mathcal{P}^1$. Consequently, problem $\mathcal{P}^\kappa$ is a scaled version of problem $\mathcal{P}^1$, where the leg capacities and the expected demand is scaled by the same factor κ. Intuitively, the parameter κ is a measure of how large the problem is and our goal is to show that the upper bound provided by our randomized linear program becomes tight as the problem gets larger.

We consider the deterministic linear program given by problem (6)-(10) for the network revenue management problem $\mathcal{P}^\kappa$. Letting z_{DLP}^κ be the optimal objective value of the deterministic linear program for problem $\mathcal{P}^\kappa$, we have

$$z_{DLP}^\kappa = \quad \max \quad \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau}[f_j y_{jt}^\kappa - \theta_j w_{jt}^\kappa] \tag{25}$$

$$\text{subject to} \quad \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} a_{ij}[q_{jt}^\kappa y_{jt}^\kappa - w_{jt}^\kappa] \leq \kappa c_i \quad \forall i \tag{26}$$

$$y_{jt}^\kappa \leq p_{jt}^\kappa \quad \forall j,t \tag{27}$$

$$w_{jt}^\kappa \leq q_{jt}^\kappa y_{jt}^\kappa \quad \forall j,t \tag{28}$$

$$y_{jt}^\kappa, w_{jt}^\kappa \geq 0 \quad \forall j,t. \tag{29}$$

We let $V_1^\kappa(\bar{0})$ denote the optimal expected total profit for problem $\mathcal{P}^\kappa$ that we obtain by solving the corresponding dynamic program. Erdelyi and Topaloglu (2009) show that the optimal objective value of the deterministic linear program provides an upper bound on the optimal expected total profit. Therefore, we have $z_{DLP}^\kappa \geq V_1^\kappa(\bar{0}) \geq 0$, where the last inequality follows from the fact that the optimal expected total profit is nonnegative since rejecting all requests is a feasible policy with an optimal expected total profit of zero. This implies that if $z_{DLP}^\kappa = 0$, then we also have $V_1^\kappa(\bar{0}) = 0$. Otherwise, $V_1^\kappa(\bar{0})/z_{DLP}^\kappa \leq 1$. In the next proposition, we show that this ratio converges to 1 as κ goes to infinity. In other words, as the problem size, measured by κ, increases, the optimal objective value of the deterministic linear program becomes a sharper and sharper estimate of the optimal expected total profit. We defer the proof of this result to Appendix B.

Proposition 5 *We have* $\lim_{\kappa\to\infty} V_1^\kappa(\bar{0})/z_{DLP}^\kappa = 1$.

It is interesting to observe that the proof of Proposition 5 in Appendix B also gives a policy to accept or reject the product requests and the ratio between the expected total profits obtained by this policy and the optimal policy converges to 1. In particular, letting $\{\hat{y}^\kappa_{jt} : \forall j,t\}$ be the optimal values of the decision variables $\{y^\kappa_{jt} : \forall j,t\}$ in problem (25)-(29), we can accept a request for product j at time period t with probability $\hat{y}^\kappa_{jt}/p^\kappa_{jt}$. Noting constraints (27), $\hat{y}^\kappa_{jt}/p^\kappa_{jt} \in [0,1]$ and we can indeed use $\hat{y}^\kappa_{jt}/p^\kappa_{jt}$ as a probability. After making the acceptance decisions, we decide which show-ups to deny boarding at the departure time by a specially constructed coin flip given in Appendix B. In this case, letting Prof^κ be the expected total profit obtained by this policy for problem $\mathcal{P}^\kappa$, the proof of Proposition 5 in Appendix B also shows that $\lim_{\kappa\to\infty} \mathsf{Prof}^\kappa / V^\kappa_1(\bar{0}) = 1$. Therefore, not only the upper bound provided by the deterministic linear program is asymptotically tight, but one can also derive a policy from the deterministic linear program whose performance becomes asymptotically optimal as κ increases.

By Proposition 4, the upper bounds provided by the randomized linear program are tighter than those provided by the deterministic linear program. Therefore, using z^κ_{RLP} to denote the optimal objective value of the randomized linear program formulated for the network revenue management problem $\mathcal{P}^\kappa$, the previous proposition immediately implies the following result.

Corollary 6 *We have* $\lim_{\kappa\to\infty} V^\kappa_1(\bar{0})/z^\kappa_{RLP} = 1$.

6 Computational Experiments

In this section, we compare the upper bounds and the expected total profits obtained by the randomized linear program with four benchmark strategies. We begin by describing the benchmark strategies and the experimental setup.

6.1 Benchmark Strategies

Deterministic Linear Program (DLP) This is the solution method described in Section 3. In our practical implementation, we divide the planning horizon into 10 equal segments and resolve problem (6)-(10) at the beginning of each segment to obtain a fresh set of bid prices. In particular, if the state of the system at the beginning of segment s is given by $\overline{x}_{\tau(s-1)/10+1}$, then we solve problem (6)-(10) after replacing the constraints (8) for $t = 1, \ldots, \tau(s-1)/10$ with the constraints $y_{jt} = x_{jt}$ to reflect the fact that we have already made the acceptance decisions for the product requests up to time period $\tau(s-1)/10$. Letting $\mu = \{\mu_i : \forall i\}$ be the optimal values of the dual variables associated with constraints (7), we accept a request for product j at time period t according to the decision rule in (11). We continue to use this decision rule until the beginning of the next segment where we resolve problem (6)-(10).

Randomized Linear Program (RLP) This is the solution method described in Section 4. Similar to the deterministic linear program, we divide the planning horizon into 10 equal segments and resolve problem (18)-(23) for K demand samples at the beginning of each segment to obtain a fresh set of bid prices. In particular, if the state of the system at the beginning of segment s is given by $\overline{x}_{\tau(s-1)/10+1}$, we solve problem (18)-(23) after replacing constraints (21) for $t = 1, \ldots, \tau(s-1)/10$ with $y^k_{jt} = x_{jt}$. This again reflects the fact that we have already made the acceptance decisions for the product requests up to time period $\tau(s-1)/10$. Using $\{\rho_i^{kl} : \forall i, l\}$ to denote the optimal values of the dual variables corresponding to constraints (19), we let $\rho_i = \sum_{k=1}^{K} \sum_{l=1}^{L} \rho_i^{kl}/K$ and accept a request for product j at time period t according to the decision rule in (24). We continue to use the above decision rule until the beginning of the next segment, at which point we resolve problem (18)-(23).

Partially Randomized Linear Program (PRLP) This solution method is similar to the randomized linear program, but instead of simulating both the demands for the products and the show-ups, we only simulate the demands for the products and assume that the numbers of show-ups take on their expected values. The main motivation for this method is that its running time is significantly less than RLP, striking a middle ground between DLP and RLP. PRLP generates K samples of the demands for the products and solves the problem

$$z_{PRLP}^k = \quad \max \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} f_j y_{jt}^k - \sum_{j=1}^{n} \sum_{t=1}^{\tau} \theta_j w_{jt}^k \tag{30}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{t=1}^{\tau} a_{ij} [q_{jt} y_{jt}^k - w_{jt}^k] \leq c_i \quad \forall i \tag{31}$$

$$w_{jt}^k \leq q_{jt} y_{jt}^k \quad \forall j, t \tag{32}$$

$$y_{jt}^k \leq d_{jt}^k \quad \forall j, t \tag{33}$$

$$y_{jt}^k, w_{jt}^k \geq 0 \quad \forall j, t \tag{34}$$

for each sample, where $d^k = \{d_{jt}^k : \forall j, t\}$ is the kth sample of the random variables $\{D_{jt} : \forall j, t\}$. Note that this approach has a lower computational burden than the randomized linear program. Using the fact that there is at most one product request in each time period, the partially randomized linear programming approach involves solving K linear programs, each having 2τ variables and $m + 2\tau$ constraints. It is possible to show that the partially randomized linear program yields an upper bound on the optimal expected total profit that lies in between the upper bounds obtained by the randomized and the deterministic linear programs. We use $\sum_{k=1}^{K} z_{PRLP}^k / K$ as an estimate of this upper bound. We obtain a bid price control policy by using the optimal values of the dual variables associated with constraints (31) as the bid prices. As with DLP and RLP, PRLP divides the planning horizon into 10 equal segments. Letting $\overline{x}_{\tau(s-1)/10+1}$ be the state of the system at the beginning of segment s, we solve problem (30)-(34) for K demand samples after replacing constraints (33) for $t = 1, \ldots, \tau(s-1)/10$ with $y_{jt}^k = x_{jt}$. This reflects the fact that we have already made the acceptance decisions for the product requests up to time period $\tau(s-1)/10$. Letting $\{\lambda_i^k : \forall i\}$ be the optimal values of the dual variables associated with constraints (31) and $\lambda_i = \sum_{k=1}^{K} \lambda_i^k / K$, we accept a request for product j at time period t only if $f_j \geq q_{jt} \sum_{i=1}^{m} a_{ij} \lambda_i$. We continue to use the above decision rule until the beginning of the next segment, at which point we resolve problem (30)-(34).

Virtual Capacities with an Economic Model (VCE) VCE is proposed by Karaesmen and van Ryzin (2004*a*). VCE chooses a virtual capacity u_i for each flight leg i so that while accepting reservations we pretend that the capacity of flight leg i is u_i instead of c_i. To get a tractable model, VCE makes the following three assumptions. First, we can make the deny boarding decisions for a reservation that uses multiple flight legs independently across each flight leg. That is, we can allow boarding to a reservation over one flight leg while denying it boarding over another flight leg. Second, the show-up probability and the denied boarding penalty cost depends only on the flight leg and is the same for all products using that flight leg. Third, VCE assumes that the number of seats sold on flight leg i is exactly equal to its virtual capacity u_i. In this case, letting Q_i be the show-up probability and Γ_i be the denied boarding penalty on flight leg i, the three assumptions imply that the number of show-ups on flight leg i, $Y_i(u_i)$, is a binomial (u_i, Q_i) random variable, while the expected denied boarding penalty cost is $\Gamma_i \mathbb{E}\{[Y_i(u_i) - c_i]^+\}$, where we use $[a]^+ = \max\{a, 0\}$. Using z_j to denote the number of product j requests that we accept over the booking horizon, VCE solves the problem

$$\begin{aligned}
\max \quad & \sum_{j=1}^{n} f_j z_j - \sum_{i=1}^{m} \Gamma_i \mathbb{E}\{[Y_i(u_i) - c_i]^+\} \\
\text{subject to} \quad & \sum_{j=1}^{n} a_{ij} z_j - u_i = 0 \quad \forall i \\
& z_j \leq \sum_{t=1}^{\tau} p_{jt} \quad \forall j \\
& z_j \geq 0 \quad \forall j \\
& u_i \geq 0 \quad \forall i.
\end{aligned}$$

We use linear interpolations of $\mathbb{E}\{[Y_i(u_i) - c_i]^+\}$ to compute the above objective function at noninteger values of u_i. VCE uses the optimal values of the dual variables associated with the first set of constraints as the bid prices and accepts a request for product j at time period t according to the decision rule in (11). In our computational experiments, we set $Q_i = \sum_{j=1}^{n} a_{ij} (\sum_{t=1}^{\tau} q_{jt}/\tau) / \sum_{j=1}^{n} a_{ij}$ so that Q_i is the average of the show-up probabilities of the products that use flight leg i. We use the same logic in setting Γ_i. In particular, we let $\bar{\theta}_j = \theta_j / \sum_{l=1}^{m} a_{lj}$ to evenly distribute the penalty cost associated with product j across the flight legs it uses and set $\Gamma_i = \sum_{j=1}^{n} a_{ij} \bar{\theta}_j / \sum_{j=1}^{n} a_{ij}$. This is one of the several choices for Q_i and Γ_i that Karaesmen and van Ryzin (2004*a*) propose and all of their proposed choices appeared to perform similarly. Finally, similar to the other solution methods, VCE divides the planning horizon into 10 equal segments, recomputes the bid prices at the beginning of each segment and uses these bid prices until the beginning of the next segment.

Stochastic Approximation Algorithm (SAA) This solution method is proposed by Kunnumkal and Topaloglu (2011) and it is based on the observation that for a given set of bid prices, the decision rule in (11) determines the numbers of requests that we accept for the different products. Thus, we can express the expected total profit as a function of the bid prices and the idea behind SAA is to use a stochastic approximation algorithm to find a set of bid prices that maximize the expected total profit. We refer the reader to Kunnumkal and Topaloglu (2011) for further details of the stochastic approximation algorithm. We use a step size of $5/(40+k)$ in the kth iteration of the stochastic approximation algorithm and terminate the algorithm after 5000 iterations. We use the bid prices obtained by the stochastic approximation algorithm at the end of 5000 iterations in the decision rule in (11) to decide whether to accept a request for product j at time period t. Similar to the other solution methods, SAA divides the planning horizon into 10 equal segments, recomputes the bid prices at the beginning of each segment and uses these bid prices until the beginning of the next segment.

6.2 Computational Results for the Network with a Single Hub

We test the performance of our benchmark solution methods on two types of networks. The first type of network has a single hub serving multiple spokes. The second type has two hubs with half of the spokes served by the first hub and the other half served by the second hub. We begin by describing our results for the first type of network.

Computational Setup We have a hub-and-spoke network with a single hub serving N spokes. We have one flight from the hub to each spoke and one flight from each spoke to the hub so that the total number of flight legs is $2N$. Figure 1 shows the structure of the network for the case where $N = 8$. The hub and each of the spokes serve as both origins and destinations. We have a high-fare and a low-fare product connecting each origin-destination pair. Therefore, we have $2N(N+1)$ products, $4N$ of which include one flight leg and $2N(N-1)$ of which include two flight legs. The high-fare product is four times as expensive as the low-fare product for each itinerary. The probability that a reservation shows up at the time of flight departure depends on only whether it is a high-fare or a low-fare product, but not on the origin-destination locations or on the reservation time. We let q^l and q^h be the show-up probabilities for a low-fare and a high-fare product, respectively. The penalty cost of denying boarding to a reservation for product j is set as $\gamma f_j + \sigma \max_{j'=1,\dots,n}\{f_{j'}\}$, where γ and σ are two parameters that we vary. We can interpret γf_j as the component of the penalty cost that is specific to the particular product while $\sigma \max_{j'=1,\dots,n}\{f_{j'}\}$ can be interpreted as the component of the penalty cost that is common across the products. Noting that the total expected demand for the capacity on flight leg i is $\sum_{t=1}^{\tau}\sum_{j=1}^{n} a_{ij} q_{jt} p_{jt}$, we measure the tightness of the leg capacities by

$$\alpha = \frac{\sum_{i=1}^{m}\sum_{t=1}^{\tau}\sum_{j=1}^{n} a_{ij} q_{jt} p_{jt}}{\sum_{i=1}^{m} c_i}.$$

We label our test problems by $(\gamma, \sigma, q^l, q^h, \alpha)$, where $(\gamma, \sigma) \in \{(4,0),(8,0),(1,1)\}$, $q^l \in \{0.7, 0.9\}$, $q^h \in \{0.7, 0.9\}$ and $\alpha \in \{1.2, 1.6\}$. This provides 24 test problems in our experimental setup. In all of our test problems, we have 8 spokes and 360 time periods in the planning horizon. We use $K = 25$ and $L = 200$ for RLP and $K = 25$ for PRLP. We note that this set of test problems is based on that in Erdelyi and Topaloglu (2009).

Comparison of Upper Bounds Table 1 gives the upper bounds obtained by DLP, RLP and PRLP. VCE and SAA do not provide upper bounds on the optimal expected profit and these solution methods are omitted in this table. The first column in Table 1 shows the problem characteristics. The second, third and fourth columns, respectively, give the upper bounds obtained by DLP, RLP and PRLP. The next two columns respectively give the percentage gap between the upper bounds obtained by RLP and DLP, and RLP and PRLP. The "✓" in the columns emphasize that the gaps are all significant at the 95% level. The results in Table 1 indicate that RLP generates significantly tighter upper bounds than both DLP and PRLP. The average percentage gap between the upper bounds obtained by DLP and RLP is around 5%, while the average gap between the upper bounds obtained by PRLP and RLP is around 4%. While PRLP provides a slightly tighter upper bound than DLP, we can further tighten the upper bound quite significantly by using RLP. This may justify the extra computational effort in simulating the show-ups in RLP as opposed to using just the expected number of show-ups in PRLP.

Comparison of Expected Total Profits Table 2 gives the expected total profits obtained by the five solution methods. The columns have the same interpretation as in Table 1 except that they compare the expected profits obtained by DLP, RLP, PRLP, VCE and SAA. The last four columns in Table 2 give the percentage gap between the expected profits obtained by RLP and the other benchmarks. We obtain the expected profits by simulating the bid price policies obtained by the different solution methods under multiple realizations of the demand and show-up random variables. We use common random numbers in our simulations; see Law and Kelton (2000). The last four columns in Table 2 include a "✓" if RLP does better than the respective solution method at the 95% significance level, a "×" otherwise and a "⊙" if there does not exist a statistically significant difference between the two.

Comparing the expected profits in Table 2, we observe that RLP and SAA typically generate the highest profits followed by DLP, VCE and PRLP without a consistent ordering between the latter three solution methods. The performance gap between RLP and DLP is around 3% on average, but we observe performance gaps as high as 8%. RLP performs better than DLP in 23 out of the 24 test problems and the gaps are statistically significant in 21 test problems. We observe one instance where DLP performs better than RLP, but the performance gap is less than half a percent. RLP performs better than PRLP in all of the test problems and the average performance gap is around 6%. A similar observation holds when we compare RLP with VCE and the average performance gap between these two solution methods is around 5%. The profits generated by RLP and SAA are comparable with the average performance gap being around -0.19%. In 20 out of the 24 test problems, there is no statistically significant difference in the profits generated by RLP and SAA. SAA does better than RLP in three test problems, while RLP does better than SAA in one test problem. To our knowledge, SAA is one of the strongest methods to compute bid price policies for joint overbooking and capacity control problems and it is quite encouraging that the performance of RLP is comparable to that of SAA.

Despite the fact that the performance of RLP and SAA are quite close to each other, there are a number of reasons that may make RLP preferable to SAA from practical perspective. To begin with, RLP provides an upper bound on the optimal expected total profits while SAA does not. In fact, DLP is, to our knowledge, the only other computationally tractable method that can obtain upper bounds on the optimal expected total profits when overbooking is allowed. RLP significantly improves the upper bounds provided by DLP, yielding improvements up to about 8%. Tighter upper bounds on the optimal expected profits can be quite valuable when assessing the optimality gap of approximate control policies. Another attractive feature of RLP is that the implementation of SAA requires tuning a number of parameters such as the step size rule and the stopping criterion of the stochastic approximation algorithm, for which there are no hard and fast rules. The implementation of RLP tends to be easier as it involves only selecting the sample sizes K and L. While doing a reasonably exhaustive search over all step size rules and stopping criteria is virtually impossible, one can certainly test quite a few different choices of K and L for RLP and settle on a choice. It is also worthwhile to point out that RLP requires solving linear programs, minimizing the need for customized coding. Finally, an important advantage of RLP over SAA is the running time. Table 3 compares the CPU seconds to solve RLP and SAA on networks with different numbers of spokes and booking horizons with different numbers of time periods. After a few setup runs, we settle on $K = 25$ and $L = 200$ in our implementation of RLP and run SAA for 5000 iterations. The left and right portions Table 3 respectively show the CPU seconds for different numbers of spokes in the airline network and different numbers of time periods in the booking horizon. All of the computational experiments are carried out on a desktop PC running Windows XP with Intel Core 2 Duo 3 GHz CPU and 4 GB RAM. We use CPLEX 11.2 to solve all the linear programs. The running times for RLP and SAA are generally comparable and are of the order of minutes. The running time of SAA is independent of the number of time periods since SAA works with an aggregate formulation, which requires only the total number of product requests but not the sequence in which the requests arrive over time. However, the running time of SAA can grow quite rapidly as the size of the airline network measured by the number of spokes increases. DLP, PRLP and VCE take at most a few seconds to solve and we do not provide their running times in Table 3, but the performance of these solution methods is not competitive to RLP or SAA.

Comparison of Service Levels While the benchmark solution methods are all designed to maximize the expected profits, this may not necessarily be the only performance measure that is of interest to the airline. In Table 4, we compare the solution methods on two other dimensions, namely the service levels and the occupancy levels. The service level gives the fraction of the show-ups that are allowed boarding, while the occupancy level is the fraction of seats that are occupied when the flights depart. The first column in Table 4 shows the problem characteristics. The next five columns give the expected service levels achieved by DLP, RLP, PRLP, VCE and SAA respectively, while the last five columns give the corresponding occupancy levels.

We observe that PRLP has the highest service levels, followed closely by RLP and SAA. VCE and DLP tend to have slightly lower service levels associated with them. On the other hand, when we compare the occupancy levels, DLP and VCE tend to have the highest occupancy levels, followed by RLP and SAA. PRLP tends to have a significantly lower occupancy level than the remaining solution methods. The above observations suggest that PRLP tends to be conservative in accepting product requests. As a result, it is able to provide service to almost all of the reservations that show up. Its denied service cost tends to be low, but at the same time, its revenues also tend to be low because it accepts only a smaller number of product requests. The net effect is lower profits. On the other hand, DLP and VCE tend to be more aggressive in accepting product requests. As a result, their revenues and occupancy levels tend to be higher. However, they may have to deny boarding to a greater fraction of the reservations that show up, which leads to lower service levels. Therefore, although DLP and VCE obtain higher revenues, they also tend to have higher denied service costs, resulting in lower overall profits. RLP and SAA seem to achieve a good balance between the service and occupancy levels. The service and occupancy levels of RLP and SAA lie in between the other solution methods. They tend to be more selective, accepting fewer but higher-value product requests. Since they accept higher-value product requests, their revenues are comparable to DLP and VCE. On the other hand, since they accept fewer number of product requests, their denied service costs are comparable to PRLP. The net result is that their overall profits tend to be higher than DLP, VCE and PRLP.

Qualitative Behavior of RLP Figure 2 gives a feel for the problem parameters than boost the performance of RLP relative to DLP. The horizontal axis gives the problem parameters. The test problems are so arranged that two consecutive problems differ only in the tightness of the leg capacities. Blocks of eight consecutive test problems have the same penalty costs of denied boarding and the penalty costs get larger as we move from left to right. The figure indicates that the performance gap between RLP and DLP generally increases as leg capacities get tighter. We also see that the performance gap between RLP and DLP increases as the penalty costs of denied boarding increases. Problems with tight leg capacities and large penalty costs tend to be more challenging to solve because the consequences of accepting an "incorrect" product request tend to be more severe. It is encouraging that RLP seems to be an attractive alternative to DLP in such settings.

Figure 3 shows how the upper bound obtained by RLP changes with the number of demand simulations K and the number of show-up simulations L. We see that RLP is fairly robust to the number of demand and show-up simulations and we obtain stable results for $K \geq 25$ and $L \geq 100$. To be on the safe side, we use $K = 25$ and $L = 200$ in our computational experiments.

6.3 Computational Results for the Network with Two Hubs

Computational Setup We have a network with two hubs serving a total of N spokes. The first half of the spokes are connected to the first hub and the second half of the spokes are connected to the second hub. Each spoke has one flight to and one flight from the hub that it is connected to. In addition, we have one flight from the first hub to the second and another flight in the reverse direction, so that the total number of flights is $2N + 2$. Figure 4 shows the structure of the network for the case where $N = 8$. The hub and each of the spokes serve as both origins and destinations. We have a high-fare and a low-fare product connecting each origin-destination pair. We randomly sample from the set of all origin-destination pairs so that the total number of products is around 150. The remaining problem parameters are set in the same manner as for the test problems with a single hub. In particular, a high-fare product is four times as expensive as the corresponding low-fare product for each itinerary.

We label our test problems by $(\gamma, \sigma, q^l, q^h, \alpha)$, where the parameters have the same interpretation as in the test problems with a single hub. We have $(\gamma, \sigma) \in \{(4, 0), (8, 0), (1, 1)\}$, $q^l \in \{0.7, 0.9\}$, $q^h \in \{0.7, 0.9\}$ and $\alpha \in \{1.2, 1.6\}$, which provides us with 24 test problems. In all of our test problems, we have 8 spokes and 360 time periods in the booking horizon.

Comparison of Upper Bounds, Expected Total Profits and Service Levels Table 5 gives the upper bounds obtained by DLP, RLP and PRLP. The columns in this table have the same interpretation as in Table 1. The results in Table 5 indicate that RLP continues to generate significantly tighter upper bounds than DLP and PRLP. The average percentage gap between the upper bounds obtained by DLP and RLP is around 6%, while that between PRLP and RLP is around 4%.

Table 6 gives the expected total profits obtained by the five solution methods. The columns in this table have the same interpretation as in Table 2. The results generally follow the same pattern as for the test problems with a single hub. RLP and SAA generate the highest profits followed by DLP, VCE and PRLP. The average performance gap between RLP and DLP is around 4%. The gaps are statistically significant in 20 out of the 24 test problems. In the remaining four test problems, the performance gaps between RLP and DLP are not statistically significant. The average performance gap between RLP and PRLP is around 9%, while that between RLP and VCE is around 6%. The gaps are statistically significant in all of the test problems. The profits generated by RLP and SAA are very close. The average performance gap between RLP and SAA is around -0.15%. In 18 out of the 24 test problems, there is no statistically significant gap in the profits generated by RLP and SAA. SAA does better than RLP in four test problems, while RLP does better than SAA in two test problems.

Finally, Table 7 gives the service and occupancy levels achieved by the five solution methods. The average service levels of DLP, RLP, PRLP, VCE and SAA are 0.96, 0.98, 1, 0.97 and 0.99 respectively. On the other hand, the average occupancy levels of DLP, RLP, PRLP, VCE and SAA are 0.9, 0.84, 0.72, 0.89 and 0.81 respectively. We again observe that RLP and SAA seem to achieve a good balance between the service and occupancy levels.

7 Conclusions

In this paper, we developed a randomized linear program to jointly make the capacity control and overbooking decisions on an airline network. Our solution approach builds on a linear programming based formulation of the network revenue management problem, where we make the capacity control and overbooking decisions after observing a realization of the demands for the products. We establish that this formulation yields a tighter upper bound on the optimal expected total profit than the deterministic linear program. We show that the upper bound provided by the deterministic linear program is asymptotically tight, which implies that the upper bound provided by the randomized linear program is also tight. Furthermore, our proof technique for this result generates a policy from the deterministic linear program whose expected total profit is asymptotically optimal.

As it is difficult to compute the expectation of the objective value of the randomized linear program analytically, in our practical implementation, we use samples of the demand and show-ups to approximate the expected values. We solve the resulting sample average approximation and use the optimal values of the dual variables associated with the flight leg capacity constraints to get a bid price control policy. Our computational experiments indicate that our approach can generate significantly tighter upper bounds and higher profits compared to numerous standard benchmark methods.

A Appendix: Obtaining a subgradient of z_{RLP}

In this section, we show that z_{RLP} is a concave function of the capacities of the flight legs by showing that it has a subgradient. Furthermore, the expression that we obtain for the subgradient of z_{RLP} motivates the bid price policy that we derive from the randomized linear program. To emphasize the dependence on the flight leg capacities $c = \{c_i : \forall i\}$, we write z_{RLP} and $z_{RLP}(d)$ respectively as $z_{RLP}(c)$

and $z_{RLP}(d, c)$ throughout. Noting that S takes on only finitely many values and $-\min\{x\} = \max\{-x\}$, we can write problem (16)-(17) equivalently as

$$z_{RLP}(d,c) = \quad \max \quad \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j y_{jt} - \sum_{s\in\{0,1\}^{n\tau}} \Pr\{S=s\} \sum_{j=1}^{n}\sum_{t=1}^{\tau} \theta_j w_{jt}^s \tag{35}$$

$$\text{subject to} \quad \sum_{j=1}^{n}\sum_{t=1}^{\tau} a_{ij}[s_{jt} y_{jt} - w_{jt}^s] \le c_i \quad \forall i, s \tag{36}$$

$$0 \le w_{jt}^s \le s_{jt} y_{jt} \quad \forall j, t, s \tag{37}$$

$$0 \le y_{jt} \le d_{jt} \quad \forall j, t, \tag{38}$$

where $\Pr\{S = s\}$ is the probability that S takes on a value $s \in \{0,1\}^{n\tau}$ and we use w_{jt}^s to denote the number of product j reservations that are denied boarding when $S = s$. Letting $\rho^{ds}(c) = \{\rho_i^{ds}(c) : \forall i\}$ denote the optimal values of the dual variables corresponding to constraints (36), it follows from the linear programming duality theory that $z_{RLP}(d, \hat{c}) \le z_{RLP}(d, c) + \sum_{i=1}^{m} \sum_{s\in\{0,1\}^{n\tau}} \rho_i^{ds}(c)[\hat{c}_i - c_i]$. Since $z_{RLP}(c) = \mathbb{E}\{z_{RLP}(d, c)\} = \sum_{d\in\{0,1\}^{n\tau}} \Pr\{D = d\} z_{RLP}(d, c)$, we obtain

$$z_{RLP}(\hat{c}) \le z_{RLP}(c) + \sum_{i=1}^{m} [\hat{c}_i - c_i] \sum_{d\in\{0,1\}^{n\tau}} \Pr\{D=d\} \sum_{s\in\{0,1\}^{n\tau}} \rho_i^{ds}(c).$$

This implies that, if we let $\rho_i(c) = \sum_{d\in\{0,1\}^{n\tau}} \Pr(D = d) \left(\sum_{s\in\{0,1\}^{n\tau}} \rho_i^{ds}(c)\right)$, then $\rho(c) = \{\rho_i(c) : \forall i\}$ is a subgradient of $z_{RLP}(c)$ with respect to c. In our Monte Carlo estimate of this subgradient, we use $\rho_i = \sum_{k=1}^{K} \frac{1}{K} \sum_{l=1}^{L} \rho_i^{kl}$ as an estimate of $\rho_i(c)$, where K and L are respectively the number of demand and show-up simulations. This is precisely the bid price that we use in the decision rule in (24).

B Appendix: Proof of Proposition 5 and Asymptotic Tightness

In this section, we give a proof for Proposition 5. In doing so, we also derive a policy from the deterministic linear program whose performance is asymptotically optimal as the problem size, measured by κ, gets large. Therefore, not only the upper bound on the optimal expected total profit provided by the deterministic linear program is asymptotically tight, we can also derive a policy from the deterministic linear program whose performance is asymptotically optimal.

Our first observation is that if $(\hat{y}, \hat{w})$ is an optimal solution to problem (6)-(10), then $(\hat{y}^\kappa, \hat{w}^\kappa)$ with $\hat{y}_{jt}^\kappa = \hat{y}_{j\lceil t/\kappa \rceil}$ and $\hat{w}_{jt}^\kappa = \hat{w}_{j\lceil t/\kappa \rceil}$ is an optimal solution to problem (25)-(29). Thus, noting that the optimal objective values of problems (6)-(10) and (25)-(29) are respectively denoted by z_{DLP} and z_{DLP}^κ, it follows that $z_{DLP}^\kappa = \kappa\, z_{DLP}$.

Now, we consider the following policy π for problem $\mathcal{P}^\kappa$. We solve problem (25)-(29) to obtain the optimal solution $(\hat{y}^\kappa, \hat{w}^\kappa)$. For accept and reject decisions, we accept a request for product j at time period t with probability $\hat{y}_{jt}^\kappa / p_{jt}^\kappa$. At the departure time, if a reservation for product j made at time period t shows up, then we allow it to board with probability $1 - \hat{w}_{jt}^\kappa / (q_{jt}^\kappa \hat{y}_{jt}^\kappa)$, *provided* there is sufficient remaining capacity on the flight legs used by that product. That is, we first flip a coin to decide whether to allow or deny boarding to that reservation. If the outcome is to allow boarding, then we do so provided we do not violate the flight leg capacity constraints. Otherwise, we deny boarding to the reservation. In making the coin flips, we order the products j and time periods t in a predefined fashion and use the same ordering throughout. Policy π is clearly feasible.

Next, we consider another policy $\tilde{\pi}$ for problem $\mathcal{P}^\kappa$. During the booking period, $\tilde{\pi}$ makes decisions in the same manner as policy π. However, at the departure time, if a reservation for product j made at time period t shows up, then we allow it boarding with probability $1 - \hat{w}^\kappa_{jt}/(q^\kappa_{jt}\hat{y}^\kappa_{jt})$, *irrespective* of the remaining capacity on the flight legs. Thus, we deny boarding to the reservation in question with probability $\hat{w}^\kappa_{jt}/(q^\kappa_{jt}\hat{y}^\kappa_{jt})$, in which case we incur a denied boarding cost of θ_j. However, if policy $\tilde{\pi}$ allows boarding to a reservation that violates the capacity on a flight leg, then we incur a large penalty cost of $\Theta > \max_j \theta_j$ for each unit of capacity consumed in excess of the available capacity on a flight leg. Therefore, policy $\tilde{\pi}$ does not pay attention to the remaining capacities when making the boarding decisions, but it pays a high cost for each unit of violated capacity. In making the coin flips for policy $\tilde{\pi}$, we follow the same ordering between the products and time periods that we follow for policy π.

We use a standard coupling argument to establish that the profits generated on each sample path by policy $\tilde{\pi}$ is a lower bound on the profits generated by policy π; see Talluri and van Ryzin (1998). Since both policies generate the same revenues, it is sufficient to show that policy $\tilde{\pi}$ incurs higher deny costs than policy π. Consider the first time that the coin flip for policy π states that we should allow boarding to a reservation, but policy π has to deny boarding because of insufficient capacities on one or more flight legs. In this case, since policy $\tilde{\pi}$ disregards the capacities, it allows boarding to the same reservation, but since this boarding decision exceeds the capacities on the flight legs, policy $\tilde{\pi}$ incurs a penalty cost of at least $\Theta > \max_j \theta_j$. Therefore, policy $\tilde{\pi}$ incurs a higher cost and at the same time, has less remaining capacity than policy π. Therefore, when we move to the next product and next time period, policy $\tilde{\pi}$ is more at risk to run over capacity than policy π. Repeating the above argument for the subsequent deny boarding decisions, we conclude that costs incurred by policy π on each sample path are lower than the costs incurred by policy $\tilde{\pi}$. Therefore, letting P^π and $P^{\tilde{\pi}}$ respectively denote the sample path profits obtained by π and $\tilde{\pi}$, we have $P^\pi \geq P^{\tilde{\pi}}$ on every sample path.

It is possible to construct an expression for the expected total profit obtained by policy $\tilde{\pi}$ and to assess how much this expected profit deviates from the one obtained by policy π. To facilitate the discussion, let Y^κ_{jt} be the random variable taking value 1 if we sell product j at time period t under policy $\tilde{\pi}$. Using our earlier notation, let S^κ_{jt} be the random variable taking value 1 if the product j reservation made at time period t shows up at the departure time. Finally, let W^κ_{jt} be the random variable taking value 1 if the reservation made for product j at time period t is denied boarding at the departure time under policy $\tilde{\pi}$. Note that W^κ_{jt} is the result of a coin flip. Letting J_t denote the random product request at time period t, we have

$$\Pr(Y^\kappa_{jt} = 1) = \Pr(\text{accept request for product } j \text{ at time period } t \,|\, J_t = j) \times \Pr(J_t = j) = \frac{\hat{y}^\kappa_{jt}}{p^\kappa_{jt}}\, p^\kappa_{jt} = \hat{y}^\kappa_{jt}.$$

On the other hand, noting that we can deny boarding to only those reservations which were accepted (that is, $Y^\kappa_{jt} = 1$) and which show up at the time of departure (that is, $S^\kappa_{jt} = 1$), we have

$$\Pr(W^\kappa_{jt} = 1) = \Pr(W^\kappa_{jt} = 1 | Y^\kappa_{jt} = 1, S^\kappa_{jt} = 1) \times \Pr(S^\kappa_{jt} = 1 | Y^\kappa_{jt} = 1) \times \Pr(Y^\kappa_{jt} = 1) = \frac{\hat{w}^\kappa_{jt}}{q^\kappa_{jt}\hat{y}^\kappa_{jt}}\, q^\kappa_{jt}\, \hat{y}^\kappa_{jt} = \hat{w}^\kappa_{jt}.$$

In this case, the sample path profit from policy $\tilde{\pi}$ becomes

$$P^{\tilde{\pi}} = \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} f_j Y^\kappa_{jt} - \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} \theta_j W^\kappa_{jt} - \Theta \sum_{i=1}^{m} \Big[\sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} a_{ij} [S^\kappa_{jt} Y^\kappa_{jt} - W^\kappa_{jt}] - \kappa c_i \Big]^+. \tag{39}$$

In the above expression, the first term gives the total revenue, the second term gives the total cost of denied service and the last term gives the penalty cost incurred from violating the capacity constraints on the flight legs. Note that $S^\kappa_{jt} Y^\kappa_{jt} - W^\kappa_{jt}$ indicates whether a reservation for itinerary j made at time period t is allowed boarding. For the expected total revenue of policy $\tilde{\pi}$ in (39), we have

$$\mathbb{E}\Big\{ \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} f_j Y^\kappa_{jt} \Big\} = \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} f_j\, \hat{y}^\kappa_{jt} = \kappa \sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j\, \hat{y}_{jt}, \tag{40}$$

where the first equality uses the fact that $\mathbb{E}\{Y^\kappa_{jt}\} = \hat{y}^\kappa_{jt}$ and the second equality follows from the fact that $\hat{y}^\kappa_{jt} = \hat{y}_{j\lceil t/\kappa \rceil}$. Similarly, for the expected total denied service cost in (39), we have

$$\mathbb{E}\Big\{ \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} \theta_j W^\kappa_{jt} \Big\} = \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} \theta_j\, \hat{w}^\kappa_{jt} = \kappa \sum_{j=1}^{n}\sum_{t=1}^{\tau} \theta_j\, \hat{w}_{jt}. \tag{41}$$

In what follows, we concentrate on the expected penalty cost incurred from violating the capacity constraints, which corresponds to the third term in (39) and we upper bound this cost component.

For notational brevity, let $Z^\kappa_{it} = \sum_{j=1}^n a_{ij} [S^\kappa_{jt} Y^\kappa_{jt} - W^\kappa_{jt}]$ and $Z^\kappa_i = \sum_{t=1}^{\kappa\tau} Z^\kappa_{it}$. Noting that $\mathbb{E}\{Z^\kappa_{it}\} = \sum_{j=1}^n a_{ij} [q^\kappa_{jt} \hat{y}^\kappa_{jt} - \hat{w}^\kappa_{jt}]$, we obtain

$$\mathbb{E}\{Z^\kappa_i\} = \sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} a_{ij} [q_{jt} \hat{y}^\kappa_{jt} - \hat{w}^\kappa_{jt}] = \kappa \sum_{j=1}^{n}\sum_{t=1}^{\tau} a_{ij} [q_{jt} \hat{y}_{jt} - \hat{w}_{jt}] \le \kappa c_i, \tag{42}$$

where the inequality follows from the fact that $(\hat{y}, \hat{w})$ is an optimal solution to problem (6)-(10) so that it satisfies constraints (7). Now, we show that the variance of Z^κ_i scales linearly with κ. Observe that Y^κ_{jt} is a Bernoulli random variable with parameter $\hat{y}^\kappa_{jt}$ and we have $\hat{y}^\kappa_{jt} = \hat{y}_{j\lceil t/\kappa \rceil}$. Fix some $\ell = 1, \ldots, \tau$. In this case, for $t = \kappa(\ell - 1) + 1, \ldots, \kappa\,\ell$, we have $\hat{y}^\kappa_{jt} = \hat{y}_{j\ell}$, which implies that the random variables $\{Y^\kappa_{jt} : t = \kappa(\ell - 1) + 1, \ldots, \kappa\,\ell\}$ are identically distributed. Repeating the same argument, we observe that the random variables $\{S^\kappa_{jt} : t = \kappa(\ell - 1) + 1, \ldots, \kappa\,\ell\}$ are identically distributed. Similarly, the random variables $\{W^\kappa_{jt} : t = \kappa(\ell - 1) + 1, \ldots, \kappa\,\ell\}$ are identically distributed as well. Noting the definition of Z^κ_{it}, it follows that the random variables $\{Z^\kappa_{it} : t = \kappa(\ell - 1) + 1, \ldots, \kappa\,\ell\}$ are identically distributed. In addition, since the acceptance and deny decisions made by policy $\tilde{\pi}$ (that is, the coin flips) at different time periods are independent, the random variables $\{Z^\kappa_{it} : t = 1, \ldots, \kappa\tau\}$ are all independent of each other. Therefore, we obtain

$$\begin{aligned} Var(Z^\kappa_i) = Var\Big(\sum_{t=1}^{\kappa\tau} Z^\kappa_{it} \Big) &= Var\Big(\sum_{\ell=1}^{\tau}\sum_{k=1}^{\kappa} Z^\kappa_{i,\kappa(\ell-1)+k} \Big) \\ &= \sum_{\ell=1}^{\tau}\sum_{k=1}^{\kappa} Var(Z^\kappa_{i,\kappa(\ell-1)+k}) = \sum_{\ell=1}^{\tau} \kappa\, Var(Z^\kappa_{i,\kappa(\ell-1)+1}), \end{aligned} \tag{43}$$

where the second equality follows by reordering the elements of the sum, the third equality follows from the fact that the random variables $\{Z_{it}^\kappa : t = 1, \ldots, \kappa\tau\}$ are independent of each other and the fourth equality follows by noting that for fixed $\ell = 1, \ldots, \tau$, the random variables $\{Z_{it}^\kappa : t = \kappa(\ell-1)+1, \ldots, \kappa\,\ell\}$ are identically distributed. Lastly, we uniformly bound the random variable Z_{it}^κ by $|Z_{it}^\kappa| \le \sum_{j=1}^n a_{ij} |S_{jt}^\kappa Y_{jt}^\kappa - W_{jt}^\kappa| \le n \sum_{j=1}^n a_{ij} \le A$, where A is a finite upper bound on $\sum_{j=1}^n a_{ij}$ and the second inequality holds since $S_{jt}^\kappa Y_{jt}^\kappa$ and W_{jt}^κ take values 0 or 1. Thus, $Var(Z_{it}^\kappa) \le A^2$ and using this bound on the chain of equalities above, we get $Var(Z_i^k) \le \tau\kappa A^2$.

Gallego (1992) shows that for any random variable X and scalar x satisfying $\mathbb{E}\{X\} \le x$, we have $\mathbb{E}\{[X-x]^+\} \le \frac{1}{2}\sqrt{Var(X)}$. By (42), we have $\mathbb{E}\{Z_i^\kappa\} \le \kappa c_i$, in which case, we get

$$\mathbb{E}\Big\{\Big[\sum_{t=1}^{\kappa\tau}\sum_{j=1}^{n} a_{ij}[S_{jt}^\kappa Y_{jt}^\kappa - W_{jt}^\kappa] - \kappa c_i\Big]^+\Big\} = \mathbb{E}\{[Z_i^\kappa - \kappa c_i]^+\} \le \frac{1}{2}\sqrt{Var(Z_i^\kappa)} \le \frac{1}{2}A\sqrt{\tau\kappa},$$

where the second inequality uses the upper bound on $Var(Z_i^\kappa)$ that we obtain by using (43). Now, we can bound the expected penalty cost incurred from violating the capacity constraints, corresponding to the third term in (39). In particular, the last inequality above implies

$$\sum_{i=1}^{m}\mathbb{E}\Big\{\Big[\sum_{j=1}^{n}\sum_{t=1}^{\kappa\tau} a_{ij}[S_{jt}^\kappa Y_{jt}^\kappa - W_{jt}^\kappa] - \kappa c_i\Big]^+\Big\} \le \frac{1}{2}Am\sqrt{\tau\kappa}.$$

Using the last inequality along with (40) and (41) in (39), we get

$$\mathbb{E}\{P^{\tilde{\pi}}\} \ge \kappa\sum_{j=1}^{n}\sum_{t=1}^{\tau} f_j \hat{y}_{jt} - \kappa\sum_{j=1}^{n}\sum_{t=1}^{\tau}\theta_j \hat{w}_{jt} - \frac{\Theta}{2}Am\sqrt{\tau\kappa} = \kappa\, z_{DLP} - \frac{\Theta}{2}Am\sqrt{\tau\kappa} = z_{DLP}^\kappa - \frac{\Theta}{2}Am\sqrt{\tau\kappa},$$

where the last equality uses the observation that $z_{DLP}^\kappa = \kappa\, z_{DLP}$, which we established at the beginning of this section. Therefore we have

$$\kappa\, z_{DLP} - \frac{\Theta}{2}Am\sqrt{\tau\kappa} \le \mathbb{E}\{P^{\tilde{\pi}}\} \le \mathbb{E}\{P^{\pi}\} \le V_1^\kappa(\bar{0}) \le z_{DLP}^\kappa = \kappa\, z_{DLP}.$$

The second inequality follows by the coupling argument above that shows that the profits generated on each sample path by policy $\tilde{\pi}$ form a lower bound on the profits generated by policy π, so that the same ordering also holds in expectation. The third inequality is by the fact that $\mathbb{E}\{P^\pi\}$ is the expected total profit collected by policy π for problem $\mathcal{P}^\kappa$, but $V_1^\kappa(\bar{0})$ is the optimal expected total profit. The fourth inequality follows by the fact that the optimal objective value of the deterministic linear program provides an upper bound on the optimal expected total profit. Dividing by κz_{DLP} and taking the limit as κ goes to infinity in the chain of inequalities above, we get $\lim_{\kappa\to\infty} V_1^\kappa(\bar{0})/z_{DLP}^\kappa = 1$.

Since $\mathbb{E}\{P^\pi\}$ is sandwiched between $\kappa\, z_{DLP} - \frac{\Theta}{2}Am\sqrt{\tau\kappa}$ and $\kappa\, z_{DLP}$ in the last chain of inequalities, it follows that the ratio between the expected total profits collected by policy π and the optimal policy also converges to 1 as κ goes to infinity. Therefore, not only the upper bound provided by the deterministic linear program is asymptotically tight, but the performance of policy π that we derive from the deterministic linear program is also asymptotically optimal.

References

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

Erdelyi, A. and Topaloglu, H. (2009), 'Separable approximations for joint capacity control and overbooking decisions in network revenue management', *Journal of Revenue and Pricing Management* **8**(1), 3–20.

Erdelyi, A. and Topaloglu, H. (2010), 'A dynamic programming decomposition method for making overbooking decisions over an airline network', *INFORMS Journal on Computing* **22**(3), 443–456.

Gallego, G. (1992), 'A minmax distribution free procedure for the (q, r) inventory model', *Operations Research Letters* **11**, 55–60.

Karaesmen, I. and van Ryzin, G. (2004*a*), Coordinating overbooking and capacity control decisions on a network, Technical report, Columbia Business School.

Karaesmen, I. and van Ryzin, G. (2004*b*), 'Overbooking with substitutable inventory classes', *Operations Research* **52**(1), 83–104.

Kleywegt, A. J. (2001), An optimal control problem of dynamic pricing, Technical report, Georgia Institute of Technology, School of Industrial and Systems Engineering.

Kunnumkal, S. and Topaloglu, H. (2008), 'A tractable revenue management model for capacity allocation and overbooking over an airline network', *Flexible Services and Manufacturing Journal* **20**, 125–147.

Kunnumkal, S. and Topaloglu, H. (2011), 'A stochastic approximation algorithm to compute bid prices for joint capacity allocation and overbooking over an airline network', *Naval Research Logistics* (to appear).

Law, A. L. and Kelton, W. D. (2000), *Simulation Modeling and Analysis*, McGraw-Hill, Boston, MA.

Simpson, R. W. (1989), Using network flow techniques to find shadow prices for market and seat inventory control, MIT flight transportation laboratory memorandum m89-1, MIT, Cambridge, MA.

Talluri, K. (2009), On bounds for network revenue management, Working paper, Universitat Pompeu Fabra, Barcelona, Spain.

Talluri, K. and van Ryzin, G. (1998), 'An analysis of bid-price controls for network revenue management', *Management Science* **44**(11), 1577–1593.

Talluri, K. and van Ryzin, G. (1999), 'A randomized linear programming method for computing network bid prices', *Transportation Science* **33**(2), 207–216.

Talluri, K. and van Ryzin, G. (2004), *The Theory and Practice of Revenue Management*, Kluwer Academic Press.

Topaloglu, H. (2009*a*), 'On the asymptotic optimality of the randomized linear program for network revenue management', *European Journal of Operational Research* **197**(3), 884–896.

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

Williamson, E. L. (1992), Airline Network Seat Control, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

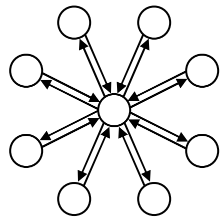

Figure 1: Structure of the network with a single hub for the case where $N = 8$.

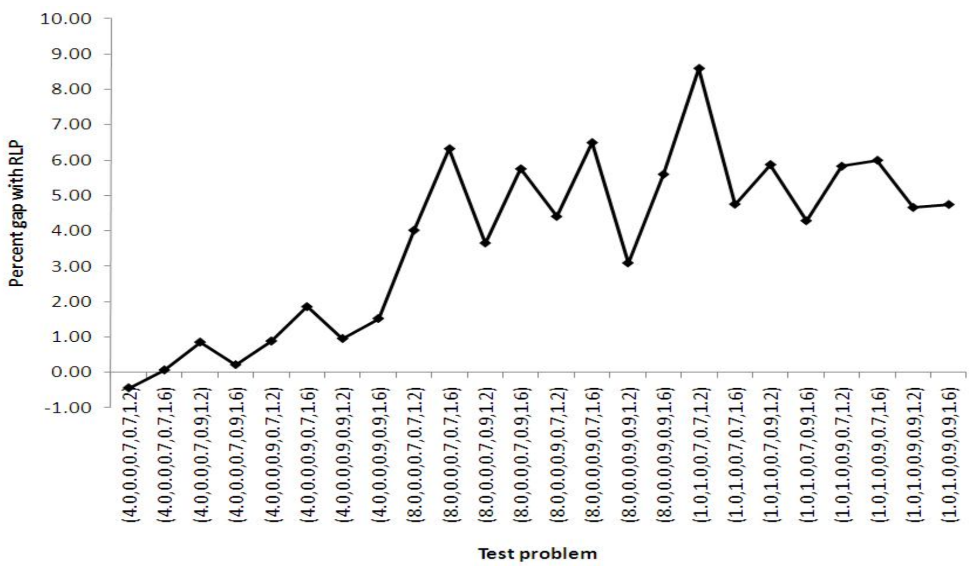

Figure 2: Performance gap between DLP and RLP for the network with a single hub.

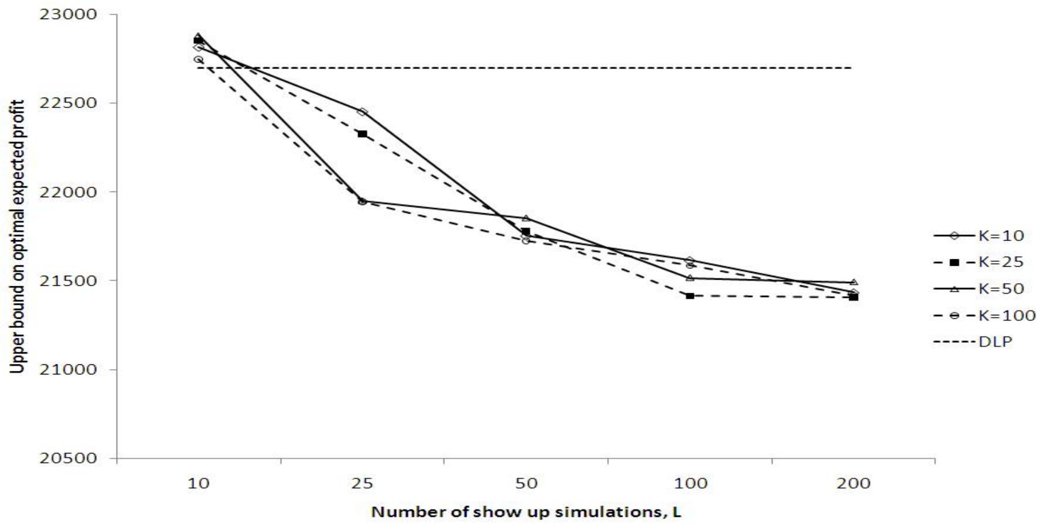

Figure 3: Sensitivity of the upper bound obtained by RLP to the number of demand and show-up simulations. The plot corresponds to the test problem on a network with a single hub with parameters (4.0, 0.0, 0.7, 0.7, 1.2).

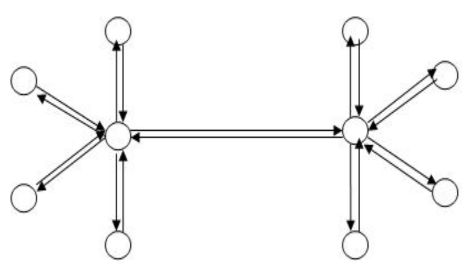

Figure 4: Structure of the network with two hubs for the case where $N = 8$.

Problem $(\gamma, \sigma, q^l, q^h, \alpha)$	Upper Bound obtained by DLP	RLP	PRLP	Percentage Gap with RLP DLP	PRLP
(4.0,0.0,0.7,0.7,1.2)	22702	21406	22108	6.06 ✓	3.28 ✓
(4.0,0.0,0.7,0.7,1.6)	21623	20202	21188	7.04 ✓	4.88 ✓
(4.0,0.0,0.7,0.9,1.2)	38538	36286	37894	6.21 ✓	4.43 ✓
(4.0,0.0,0.7,0.9,1.6)	24886	23598	24505	5.46 ✓	3.84 ✓
(4.0,0.0,0.9,0.7,1.2)	27591	26201	26998	5.30 ✓	3.04 ✓
(4.0,0.0,0.9,0.7,1.6)	31987	30478	31735	4.95 ✓	4.12 ✓
(4.0,0.0,0.9,0.9,1.2)	30735	30033	30547	2.34 ✓	1.71 ✓
(4.0,0.0,0.9,0.9,1.6)	20518	20069	20351	2.24 ✓	1.41 ✓
(8.0,0.0,0.7,0.7,1.2)	29435	27539	28832	6.89 ✓	4.69 ✓
(8.0,0.0,0.7,0.7,1.6)	32649	30259	32390	7.90 ✓	7.04 ✓
(8.0,0.0,0.7,0.9,1.2)	28493	26547	27814	7.33 ✓	4.77 ✓
(8.0,0.0,0.7,0.9,1.6)	26712	25548	26273	4.56 ✓	2.84 ✓
(8.0,0.0,0.9,0.7,1.2)	28944	27599	28423	4.87 ✓	2.99 ✓
(8.0,0.0,0.9,0.7,1.6)	29991	28448	29876	5.43 ✓	5.02 ✓
(8.0,0.0,0.9,0.9,1.2)	29684	28691	29008	3.46 ✓	1.10 ✓
(8.0,0.0,0.9,0.9,1.6)	29124	27981	28915	4.08 ✓	3.34 ✓
(1.0,1.0,0.7,0.7,1.2)	33411	31376	32509	6.48 ✓	3.61 ✓
(1.0,1.0,0.7,0.7,1.6)	30265	28269	29972	7.06 ✓	6.03 ✓
(1.0,1.0,0.7,0.9,1.2)	36912	35084	36179	5.21 ✓	3.12 ✓
(1.0,1.0,0.7,0.9,1.6)	33660	31707	33300	6.16 ✓	5.02 ✓
(1.0,1.0,0.9,0.7,1.2)	37749	36359	36984	3.82 ✓	1.72 ✓
(1.0,1.0,0.9,0.7,1.6)	27103	25376	26746	6.81 ✓	5.40 ✓
(1.0,1.0,0.9,0.9,1.2)	30249	28945	29710	4.50 ✓	2.64 ✓
(1.0,1.0,0.9,0.9,1.6)	23820	23156	23775	2.87 ✓	2.67 ✓

Table 1: Comparison of upper bounds on the network with one hub.

Problem $(\gamma, \sigma, q^l, q^h, \alpha)$	Exp. Profit obtained by DLP	RLP	PRLP	VCE	SAA	Percentage Gap with RLP DLP	PRLP	VCE	SAA
(4.0,0.0,0.7,0.7,1.2)	20777	20685	19384	20178	20788	-0.44 ×	6.29 ✓	2.45 ✓	-0.50 ⊙
(4.0,0.0,0.7,0.7,1.6)	19310	19322	18646	18433	19452	0.06 ⊙	3.50 ✓	4.60 ✓	-0.67 ×
(4.0,0.0,0.7,0.9,1.2)	34797	35096	32614	33297	35189	0.85 ✓	7.07 ✓	5.13 ✓	-0.27 ⊙
(4.0,0.0,0.7,0.9,1.6)	21815	21862	21397	20673	21935	0.21 ⊙	2.13 ✓	5.44 ✓	-0.33 ⊙
(4.0,0.0,0.9,0.7,1.2)	25895	26125	23808	25231	26094	0.88 ✓	8.87 ✓	3.42 ✓	0.12 ⊙
(4.0,0.0,0.9,0.7,1.6)	28888	29433	27589	27840	29380	1.85 ✓	6.27 ✓	5.41 ✓	0.18 ⊙
(4.0,0.0,0.9,0.9,1.2)	28454	28728	25906	27811	28750	0.95 ✓	9.82 ✓	3.19 ✓	-0.08 ⊙
(4.0,0.0,0.9,0.9,1.6)	18252	18533	17390	17758	18430	1.52 ✓	6.17 ✓	4.18 ✓	0.55 ✓
(8.0,0.0,0.7,0.7,1.2)	25207	26261	24644	25351	26172	4.01 ✓	6.16 ✓	3.47 ✓	0.34 ⊙
(8.0,0.0,0.7,0.7,1.6)	26120	27883	27312	26435	28011	6.32 ✓	2.05 ✓	5.19 ✓	-0.46 ⊙
(8.0,0.0,0.7,0.9,1.2)	24611	25544	23880	24383	25410	3.65 ✓	6.52 ✓	4.55 ✓	0.52 ⊙
(8.0,0.0,0.7,0.9,1.6)	21823	23155	22433	21849	23089	5.75 ✓	3.12 ✓	5.64 ✓	0.28 ⊙
(8.0,0.0,0.9,0.7,1.2)	25400	26571	24373	25608	26704	4.41 ✓	8.27 ✓	3.62 ✓	-0.50 ⊙
(8.0,0.0,0.9,0.7,1.6)	24650	26362	25149	25270	26397	6.49 ✓	4.60 ✓	4.14 ✓	-0.13 ⊙
(8.0,0.0,0.9,0.9,1.2)	25968	26796	24427	26155	27009	3.09 ✓	8.84 ✓	2.39 ✓	-0.79 ×
(8.0,0.0,0.9,0.9,1.6)	23977	25399	24168	24493	25382	5.60 ✓	4.85 ✓	3.57 ✓	0.07 ⊙
(1.0,1.0,0.7,0.7,1.2)	27002	29541	27595	27342	29500	8.59 ✓	6.59 ✓	7.44 ✓	0.14 ⊙
(1.0,1.0,0.7,0.7,1.6)	24378	25594	25098	23561	25912	4.75 ✓	1.94 ✓	7.94 ✓	-1.24 ×
(1.0,1.0,0.7,0.9,1.2)	30816	32739	30380	31412	32831	5.87 ✓	7.20 ✓	4.05 ✓	-0.28 ⊙
(1.0,1.0,0.7,0.9,1.6)	27751	28993	27500	27669	29060	4.28 ✓	5.15 ✓	4.57 ✓	-0.23 ⊙
(1.0,1.0,0.9,0.7,1.2)	32469	34479	31596	31458	34497	5.83 ✓	8.36 ✓	8.76 ✓	-0.05 ⊙
(1.0,1.0,0.9,0.7,1.6)	22731	24181	22549	22570	24093	6.00 ✓	6.75 ✓	6.66 ✓	0.36 ⊙
(1.0,1.0,0.9,0.9,1.2)	25650	26904	24836	25866	27091	4.66 ✓	7.69 ✓	3.86 ✓	-0.70 ⊙
(1.0,1.0,0.9,0.9,1.6)	20036	21034	19978	19931	21206	4.74 ✓	5.02 ✓	5.24 ✓	-0.82 ⊙

Table 2: Comparison of expected total profits on the network with one hub.

No. of spokes	CPU secs. RLP	SAA
4	63	83
8	209	528
12	317	1737
16	658	4143

No. of time periods	CPU secs. RLP	SAA
180	54	500
360	209	528
540	404	539
720	683	539

Table 3: CPU seconds for RLP and SAA as a function of the number of spokes in the airline network and the number of time periods in the booking horizon.

Problem	Exp. Service Level					Exp. Occupancy				
$(\gamma, \sigma, q^l, q^h, \alpha)$	DLP	RLP	PRLP	VCE	SAA	DLP	RLP	PRLP	VCE	SAA
(4.0,0.0,0.7,0.7,1.2)	0.95	0.98	0.99	0.96	0.95	0.90	0.84	0.73	0.90	0.89
(4.0,0.0,0.7,0.7,1.6)	0.92	0.97	0.98	0.94	0.95	0.90	0.81	0.76	0.89	0.85
(4.0,0.0,0.7,0.9,1.2)	0.94	0.98	0.99	0.98	0.97	0.92	0.85	0.75	0.87	0.87
(4.0,0.0,0.7,0.9,1.6)	0.93	0.96	0.98	0.97	0.94	0.91	0.83	0.79	0.88	0.86
(4.0,0.0,0.9,0.7,1.2)	0.96	0.98	1.00	0.96	0.98	0.92	0.89	0.71	0.93	0.88
(4.0,0.0,0.9,0.7,1.6)	0.94	0.97	0.99	0.95	0.98	0.93	0.88	0.73	0.94	0.85
(4.0,0.0,0.9,0.9,1.2)	0.95	0.97	1.00	0.98	0.98	0.92	0.89	0.70	0.91	0.88
(4.0,0.0,0.9,0.9,1.6)	0.95	0.97	0.99	0.97	0.97	0.91	0.89	0.74	0.91	0.86
(8.0,0.0,0.7,0.7,1.2)	0.96	0.99	1.00	0.98	0.99	0.90	0.82	0.71	0.86	0.81
(8.0,0.0,0.7,0.7,1.6)	0.95	0.98	0.99	0.96	0.99	0.90	0.80	0.75	0.87	0.76
(8.0,0.0,0.7,0.9,1.2)	0.96	0.98	1.00	0.99	0.98	0.90	0.84	0.74	0.85	0.82
(8.0,0.0,0.7,0.9,1.6)	0.95	0.98	0.99	0.98	0.99	0.90	0.83	0.77	0.86	0.78
(8.0,0.0,0.9,0.7,1.2)	0.97	0.99	1.00	0.97	0.99	0.91	0.88	0.69	0.92	0.84
(8.0,0.0,0.9,0.7,1.6)	0.96	0.98	1.00	0.97	0.99	0.92	0.87	0.73	0.91	0.79
(8.0,0.0,0.9,0.9,1.2)	0.97	0.98	1.00	0.99	0.99	0.92	0.89	0.69	0.90	0.83
(8.0,0.0,0.9,0.9,1.6)	0.97	0.98	1.00	0.98	0.99	0.92	0.88	0.72	0.90	0.79
(1.0,1.0,0.7,0.7,1.2)	0.96	0.99	1.00	0.96	0.99	0.90	0.81	0.71	0.90	0.79
(1.0,1.0,0.7,0.7,1.6)	0.96	0.98	0.99	0.95	0.99	0.90	0.78	0.75	0.90	0.76
(1.0,1.0,0.7,0.9,1.2)	0.96	0.98	1.00	0.98	0.99	0.91	0.83	0.72	0.87	0.81
(1.0,1.0,0.7,0.9,1.6)	0.96	0.98	0.99	0.97	0.98	0.90	0.83	0.74	0.89	0.80
(1.0,1.0,0.9,0.7,1.2)	0.97	0.99	1.00	0.96	0.99	0.92	0.86	0.68	0.94	0.81
(1.0,1.0,0.9,0.7,1.6)	0.97	0.98	0.99	0.96	0.99	0.92	0.86	0.74	0.93	0.79
(1.0,1.0,0.9,0.9,1.2)	0.98	0.99	1.00	0.99	1.00	0.91	0.88	0.69	0.91	0.80
(1.0,1.0,0.9,0.9,1.6)	0.97	0.98	1.00	0.98	0.99	0.92	0.87	0.72	0.92	0.79

Table 4: Comparison of additional performance measures on the network with one hub.

Problem $(\gamma,\sigma,q^l,q^h,\alpha)$	Upper Bound obtained by DLP	RLP	PRLP	Percentage Gap with RLP DLP	PRLP
(4.0,0.0,0.7,0.7,1.2)	26014	24375	25235	6.73 ✓	3.53 ✓
(4.0,0.0,0.7,0.7,1.6)	21688	20038	21404	8.24 ✓	6.81 ✓
(4.0,0.0,0.7,0.9,1.2)	27751	26184	27513	5.98 ✓	5.08 ✓
(4.0,0.0,0.7,0.9,1.6)	29798	28464	29892	4.69 ✓	5.02 ✓
(4.0,0.0,0.9,0.7,1.2)	23977	23110	23766	3.75 ✓	2.84 ✓
(4.0,0.0,0.9,0.7,1.6)	25418	24496	24957	3.76 ✓	1.88 ✓
(4.0,0.0,0.9,0.9,1.2)	26027	25199	25401	3.29 ✓	0.80 ⊙
(4.0,0.0,0.9,0.9,1.6)	24554	23599	24188	4.05 ✓	2.49 ⊙
(8.0,0.0,0.7,0.7,1.2)	26332	24497	25830	7.49 ✓	5.44 ⊙
(8.0,0.0,0.7,0.7,1.6)	21766	20208	21721	7.71 ✓	7.48 ✓
(8.0,0.0,0.7,0.9,1.2)	23213	21551	22409	7.71 ✓	3.98 ✓
(8.0,0.0,0.7,0.9,1.6)	31292	29466	30649	6.19 ✓	4.01 ✓
(8.0,0.0,0.9,0.7,1.2)	25396	24364	25263	4.24 ✓	3.69 ✓
(8.0,0.0,0.9,0.7,1.6)	29750	27979	29276	6.33 ✓	4.63 ✓
(8.0,0.0,0.9,0.9,1.2)	31200	29691	30431	5.08 ✓	2.49 ✓
(8.0,0.0,0.9,0.9,1.6)	23116	22132	22423	4.45 ✓	1.31 ⊙
(1.0,1.0,0.7,0.7,1.2)	32805	30344	32204	8.11 ✓	6.13 ✓
(1.0,1.0,0.7,0.7,1.6)	23403	21980	23362	6.48 ✓	6.29 ✓
(1.0,1.0,0.7,0.9,1.2)	31309	28986	30521	8.01 ✓	5.30 ✓
(1.0,1.0,0.7,0.9,1.6)	22469	20807	22111	7.99 ✓	6.27 ✓
(1.0,1.0,0.9,0.7,1.2)	25073	24005	24733	4.45 ✓	3.04 ✓
(1.0,1.0,0.9,0.7,1.6)	25202	23861	25134	5.62 ✓	5.34 ✓
(1.0,1.0,0.9,0.9,1.2)	25124	23999	24579	4.69 ✓	2.42 ✓
(1.0,1.0,0.9,0.9,1.6)	21245	20340	20826	4.45 ✓	2.39 ⊙

Table 5: Comparison of upper bounds on the network with two hubs.

Problem $(\gamma,\sigma,q^l,q^h,\alpha)$	Exp. Profit obtained by DLP	RLP	PRLP	VCE	SAA	Percentage Gap with RLP DLP	PRLP	VCE	SAA
(4.0,0.0,0.7,0.7,1.2)	23247	23203	21430	22562	23402	-0.19 ⊙	7.64 ✓	2.76 ✓	-0.86 ×
(4.0,0.0,0.7,0.7,1.6)	19285	19341	18234	18416	19540	0.29 ⊙	5.72 ✓	4.78 ✓	-1.03 ×
(4.0,0.0,0.7,0.9,1.2)	25440	25386	23580	23914	25550	-0.21 ⊙	7.11 ✓	5.80 ✓	-0.65 ×
(4.0,0.0,0.7,0.9,1.6)	26125	26320	25121	24643	26498	0.74 ⊙	4.56 ✓	6.37 ✓	-0.68 ×
(4.0,0.0,0.9,0.7,1.2)	22299	22483	19976	21640	22440	0.82 ✓	11.15 ✓	3.75 ✓	0.19 ⊙
(4.0,0.0,0.9,0.7,1.6)	22631	23094	21118	21697	23089	2.01 ✓	8.56 ✓	6.05 ✓	0.02 ⊙
(4.0,0.0,0.9,0.9,1.2)	23922	24341	21548	23576	24341	1.72 ✓	11.47 ✓	3.14 ✓	0.00 ⊙
(4.0,0.0,0.9,0.9,1.6)	21868	22199	20368	21223	22194	1.49 ✓	8.25 ✓	4.40 ✓	0.02 ⊙
(8.0,0.0,0.7,0.7,1.2)	21444	22723	21031	21903	22804	5.63 ✓	7.45 ✓	3.61 ✓	-0.35 ⊙
(8.0,0.0,0.7,0.7,1.6)	17501	18652	17552	17621	18715	6.17 ✓	5.90 ✓	5.53 ✓	-0.34 ⊙
(8.0,0.0,0.7,0.9,1.2)	19137	20206	18524	19211	20090	5.29 ✓	8.32 ✓	4.92 ✓	0.57 ✓
(8.0,0.0,0.7,0.9,1.6)	25439	26928	24785	25688	26644	5.53 ✓	7.96 ✓	4.61 ✓	1.05 ✓
(8.0,0.0,0.9,0.7,1.2)	22542	23359	20832	22208	23399	3.50 ✓	10.82 ✓	4.93 ✓	-0.17 ⊙
(8.0,0.0,0.9,0.7,1.6)	24844	26324	23938	24477	26186	5.62 ✓	9.06 ✓	7.01 ✓	0.52 ⊙
(8.0,0.0,0.9,0.9,1.2)	27348	28117	24604	27444	27986	2.73 ✓	12.49 ✓	2.39 ✓	0.47 ⊙
(8.0,0.0,0.9,0.9,1.6)	18719	20171	18520	19008	20299	7.20 ✓	8.19 ✓	5.77 ✓	-0.63 ⊙
(1.0,1.0,0.7,0.7,1.2)	26944	28969	26211	27116	28982	6.99 ✓	9.52 ✓	6.39 ✓	-0.04 ⊙
(1.0,1.0,0.7,0.7,1.6)	18945	20248	18565	18229	20114	6.43 ✓	8.31 ✓	9.97 ✓	0.66 ⊙
(1.0,1.0,0.7,0.9,1.2)	25481	27418	24887	26033	27312	7.06 ✓	9.23 ✓	5.05 ✓	0.39 ⊙
(1.0,1.0,0.7,0.9,1.6)	17862	19527	18073	18014	19650	8.52 ✓	7.44 ✓	7.75 ✓	-0.63 ⊙
(1.0,1.0,0.9,0.7,1.2)	21032	22479	19793	19183	22590	6.44 ✓	11.95 ✓	14.66 ✓	-0.49 ⊙
(1.0,1.0,0.9,0.7,1.6)	20415	22210	20377	19149	22289	8.08 ✓	8.25 ✓	13.78 ✓	-0.36 ⊙
(1.0,1.0,0.9,0.9,1.2)	20714	22072	19775	20802	22248	6.15 ✓	10.41 ✓	5.75 ✓	-0.80 ⊙
(1.0,1.0,0.9,0.9,1.6)	17323	18652	16720	17178	18748	7.13 ✓	10.36 ✓	7.90 ✓	-0.52 ⊙

Table 6: Comparison of expected total profits on the network with two hubs.

Problem	Exp. Service Level					Exp. Occupancy				
$(\gamma, \sigma, q^l, q^h, \alpha)$	DLP	RLP	PRLP	VCE	SAA	DLP	RLP	PRLP	VCE	SAA
(4.0,0.0,0.7,0.7,1.2)	0.95	0.98	1.00	0.96	0.97	0.90	0.84	0.73	0.90	0.87
(4.0,0.0,0.7,0.7,1.6)	0.93	0.98	0.99	0.95	0.96	0.90	0.82	0.76	0.90	0.85
(4.0,0.0,0.7,0.9,1.2)	0.95	0.98	0.99	0.99	0.97	0.89	0.82	0.74	0.84	0.84
(4.0,0.0,0.7,0.9,1.6)	0.94	0.98	0.99	0.98	0.97	0.88	0.80	0.75	0.84	0.82
(4.0,0.0,0.9,0.7,1.2)	0.97	0.98	1.00	0.96	0.98	0.90	0.87	0.70	0.93	0.87
(4.0,0.0,0.9,0.7,1.6)	0.95	0.98	1.00	0.94	0.98	0.91	0.87	0.73	0.94	0.85
(4.0,0.0,0.9,0.9,1.2)	0.96	0.98	1.00	0.98	0.98	0.92	0.90	0.70	0.91	0.87
(4.0,0.0,0.9,0.9,1.6)	0.96	0.97	1.00	0.98	0.98	0.92	0.89	0.72	0.91	0.86
(8.0,0.0,0.7,0.7,1.2)	0.96	0.98	1.00	0.98	0.98	0.89	0.82	0.72	0.86	0.80
(8.0,0.0,0.7,0.7,1.6)	0.95	0.98	0.99	0.97	0.98	0.89	0.80	0.74	0.87	0.78
(8.0,0.0,0.7,0.9,1.2)	0.95	0.98	1.00	0.99	0.98	0.89	0.82	0.73	0.81	0.80
(8.0,0.0,0.7,0.9,1.6)	0.95	0.98	0.99	0.99	0.99	0.89	0.81	0.74	0.82	0.77
(8.0,0.0,0.9,0.7,1.2)	0.97	0.99	1.00	0.97	0.99	0.90	0.87	0.69	0.92	0.83
(8.0,0.0,0.9,0.7,1.6)	0.97	0.98	1.00	0.97	0.99	0.91	0.86	0.72	0.92	0.79
(8.0,0.0,0.9,0.9,1.2)	0.98	0.99	1.00	0.99	0.99	0.90	0.88	0.68	0.89	0.83
(8.0,0.0,0.9,0.9,1.6)	0.97	0.98	1.00	0.98	0.99	0.91	0.87	0.71	0.90	0.80
(1.0,1.0,0.7,0.7,1.2)	0.97	0.99	1.00	0.97	0.99	0.89	0.80	0.71	0.89	0.78
(1.0,1.0,0.7,0.7,1.6)	0.96	0.99	0.99	0.96	0.99	0.87	0.76	0.72	0.89	0.73
(1.0,1.0,0.7,0.9,1.2)	0.97	0.99	1.00	0.99	0.99	0.90	0.82	0.72	0.85	0.78
(1.0,1.0,0.7,0.9,1.6)	0.96	0.99	1.00	0.98	0.99	0.90	0.80	0.73	0.86	0.77
(1.0,1.0,0.9,0.7,1.2)	0.98	0.99	1.00	0.96	1.00	0.90	0.86	0.68	0.94	0.80
(1.0,1.0,0.9,0.7,1.6)	0.97	0.98	1.00	0.95	0.99	0.92	0.85	0.72	0.94	0.77
(1.0,1.0,0.9,0.9,1.2)	0.98	0.99	1.00	0.98	0.99	0.91	0.87	0.68	0.91	0.80
(1.0,1.0,0.9,0.9,1.6)	0.97	0.99	1.00	0.97	1.00	0.91	0.86	0.70	0.91	0.78

Table 7: Comparison of additional performance measures on the network with two hubs.