# A Stochastic Approximation Algorithm for Making Pricing Decisions in Network Revenue Management Problems

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

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

April 29, 2010

**Abstract**

In this paper, we develop a stochastic approximation algorithm for making pricing decisions in network revenue management problems. In the setting we consider, the probability of observing a request for an itinerary depends on the price for the itinerary. We are interested in finding a set of prices that maximize the total expected revenue. Our approach is based on visualizing the total expected revenue as a function of the prices and using the stochastic gradients of the total revenue to search for a good set of prices. To compute the stochastic gradients of the total revenue, we use a construction that decouples the prices for the itineraries from the probability distributions of the itinerary requests. This construction ensures that the probability distributions of the underlying random variables do not change when we change the prices for the itineraries. We establish the convergence of our stochastic approximation algorithm. Computational experiments indicate that the prices obtained by our stochastic approximation algorithm perform significantly better than those obtained by standard benchmark strategies, especially when the leg capacities are tight and there are large differences between the price sensitivities of the different market segments.

Pricing and capacity allocation constitute two fundamental control mechanisms in network revenue management. Pricing assumes that the demands for different itineraries are price sensitive and deals with the question of what prices to charge for the itineraries. Capacity allocation, on the other hand, assumes that the prices for the itineraries are fixed and deals with the question of which itineraries to leave open for sale and which itineraries to close. It has been traditionally argued that airlines are more suitable for capacity allocation than for pricing. The first part of the argument is that advertising and administrative needs require airlines to fix the prices in advance of sales, ruling out the possibility of adjusting the prices dynamically as sales take place. The second part of the argument is that airlines are able to impose restrictions, such as Saturday night stays, to ensure that if an itinerary that appeals to a particular customer segment is closed for sale, then this customer segment does not switch to another itinerary. As a result, airlines can adjust the demand from a customer segment simply by opening or closing the itinerary that appeals to the customer segment. Although the argument above was certainly valid for the early days of network revenue management, online sales channels nowadays do not require fixing the prices in advance and competition makes it harder to impose restrictions. Due to these developments, pricing has become an increasingly appealing control mechanism for airlines.

In this paper, we develop a stochastic approximation algorithm to make pricing decisions in network revenue management problems. In the setting we consider, the probability of observing a request for an itinerary depends on the price for the itinerary. Pricing serves as the main control mechanism and we do not have the option of rejecting an itinerary request when we have enough capacity to serve the itinerary request. We are interested in finding a set of prices that maximize the total expected revenue from the accepted itinerary requests. Our approach visualizes the total revenue as a function of the prices and uses the stochastic gradients of the total revenue to search for a good set of prices.

We overcome two difficulties to compute the stochastic gradients of the total revenue. First, the prices for the itineraries affect the probability distributions of the itinerary requests. Therefore, every time we change the prices, we end up changing the probability distributions of the underlying random variables. We address this difficulty by using a standard construction that decouples the prices for the itineraries from the probability distributions of the underlying random variables. In this construction, each customer is interested in an itinerary and associates a reservation price with this itinerary. The customer purchases the itinerary if its reservation price exceeds the price for the itinerary. In this case, the itineraries that are of interest to the customers and the reservation prices serve as the underlying random variables and samples of these random variables can be generated without using the prices at all. Second, if we perturb the price for an itinerary, then the number of accepted itinerary requests either does not change or changes by a discrete amount. Therefore, the stochastic gradient of the total revenue is either zero or does not exist. We address this difficulty by using a smoothed version of the problem, which assumes that the leg capacities are continuous and we can serve fractional numbers of itinerary requests. This modification ensures that the total expected revenue is differentiable respect to the prices and we can use the stochastic gradients of the total revenue to find a good set of prices.

There are a number of papers that use stochastic approximation algorithms to find good control mechanisms for network revenue management problems. The distinguishing aspect of our work is that our control mechanism is pricing, whereas the control mechanism in the earlier work is exclusively capacity control. Bertsimas and de Boer (2005) were the first to use stochastic approximation ideas to find good booking limits in network revenue management problems. Since booking limits are restricted to be integers, the authors work with the finite differences of the total revenue, rather than the stochastic gradients. As a result, although their approach performs quite well, it does not have a convergence guarantee. This shortcoming is addressed by van Ryzin and Vulcano (2008*b*), where the problem is smoothed by assuming that the leg capacities are continuous and it is possible to serve fractional numbers of itinerary requests. In particular, these authors observe that if there are multiple flight legs whose capacities are simultaneously binding, then the total revenue is not necessarily differentiable. To avoid such cases, they perturb the remaining leg capacities by small random amounts. The smoothed problem that we use in our paper utilizes similar capacity perturbations. The work by van Ryzin and Vulcano (2008*b*) is subsequently extended by van Ryzin and Vulcano (2008*a*) to incorporate the customer choice process, where each customer, rather than having a fixed itinerary of interest, observes the set of itineraries that are available for sale and makes a purchase within this set. A customer has a ranked list of itineraries that are of interest and it purchases the first itinerary in the ranked list that is available for sale. As a result, the underlying random variables in van Ryzin and Vulcano (2008*a*) are significantly more complicated than the ones considered in the earlier literature.

The papers above focus on computing a good set of booking limits. A booking limit policy essentially restricts how many seats can be used by different virtual classes. On the other hand, Topaloglu (2008) and Chaneton and Vulcano (2009) focus on computing a good set of bid prices by using stochastic approximation algorithms. In a bid price policy, each flight leg in the airline network has a bid price associated with it and an itinerary request is accepted only if the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs that are in the itinerary. As such, the bid prices serve as revenue barriers. Topaloglu (2008) considers a setting where each customer has a fixed itinerary of interest, whereas Chaneton and Vulcano (2009) explicitly incorporate the customer choice process. Both papers end up perturbing the remaining leg capacities by small random amounts to ensure that the total expected revenue is differentiable. Kunnumkal and Topaloglu (2009) consider the problem of finding a good set of bid prices when overbooking is possible. They show that if overbooking is possible, then it is not necessary to perturb the remaining leg capacities by small random amounts to ensure that the total expected revenue is differentiable.

There is rich literature on dynamic pricing, but majority of the papers focus on pricing a single product, whereas network revenue management requires pricing multiple itineraries jointly. Gallego and van Ryzin (1994) focus on adjusting the price of one product with price sensitive demand. They establish how the optimal price changes as a function of the length of the selling horizon and the remaining inventory. In addition, they show that charging a single price over the entire selling horizon is asymptotically optimal as the length of the selling horizon and the initial inventory increases with the same rate. Feng and Gallego (1995) focus on only one price change, which can be either from low to high or from high to low and they characterize the timing of the price change. Feng and Gallego (2000), Feng and Xiao (2000) and Zhao and Zheng (2000) extend the work of Gallego and van Ryzin (1994) to more complicated demand dynamics and price constraints. Maglaras and Meissner (2006) show that certain pricing problems can be cast as capacity allocation problems and this result allows them to extend the structural properties for capacity allocation problems to pricing problems. Gallego and van Ryzin (1997) propose a deterministic optimization problem to price multiple products that interact with each other. This optimization problem computes a single price for each product to charge over the entire selling horizon. The authors show that charging a single price over the entire selling horizon is asymptotically optimal in the same sense as in Gallego and van Ryzin (1994). We use the optimization problem proposed by Gallego and van Ryzin (1997) as a benchmark strategy. Kleywegt (2001) develops a joint pricing and overbooking model over an airline network. He assumes that the demands for the itineraries are deterministic functions of the prices and solves the model by using Lagrangian duality ideas. McGill and van Ryzin (1999), Bitran and Caldentey (2003), Elmaghraby and Keskinocak (2003) and Talluri and van Ryzin (2005) give comprehensive reviews of dynamic pricing problems.

Stochastic approximation algorithms have many applications. Fu (1994) focuses on the problem of computing base stock levels in single echelon inventory systems, whereas Glasserman and Tayur (1995) consider the same problem in multiple echelon setting. Bashyam and Fu (1998) show how to deal with random lead times and service level constraints. The paper by van Ryzin and McGill (2000) computes booking limits for a revenue management problem with a single flight leg. Mahajan and van Ryzin (2001) consider the problem of setting stock levels for multiple products, where the products can serve as substitutes of each other. Karaesmen and van Ryzin (2004) find good overbooking policies for parallel flight legs that operate between the same origin destination pair. Kunnumkal and Topaloglu (2008) focus on various inventory control problems, where the optimal policy is known to be a base stock policy. Kushner and Clark (1978), Benveniste, Metivier and Priouret (1991), Bertsekas and Tsitsiklis (1996) and Kushner and Yin (2003) cover the theory of stochastic approximation algorithms.

In this paper, we make the following research contributions. 1) We build a stochastic approximation algorithm to find a good set of prices in network revenue management problems. It appears that there are few practical algorithms for jointly pricing multiple products and our paper fills this gap. 2) We use a construction that decouples the prices for the itineraries from the probability distributions of the underlying random variables. In this construction, each customer is interested in a particular itinerary and associates a reservation price with this itinerary. The itineraries that are of interest to the customers and the reservation prices serve as the underlying random variables and samples of these random variables can be generated without using the prices. It is widely known that using a probability distribution for the reservation price is equivalent to using a function that relates the expected demand to the price, but using this equivalence in a stochastic approximation algorithm is particularly helpful and this equivalence may be useful to build algorithms for other pricing problems. 3) We establish the convergence of our stochastic approximation algorithm. 4) Computational experiments indicate that the prices obtained by our approach perform significantly better than those obtained by standard benchmark strategies. The performance gaps are especially significant when the leg capacities are tight and there are large differences between the price sensitivities of the different market segments.

The rest of the paper is organized as follows. In Section 1, we give a description of the pricing problem that is of interest to us. In Section 2, we show how to smooth the problem so that the total expected revenue is differentiable with respect to the prices and we formulate a basic optimization problem that chooses the prices to maximize the total expected revenue. In Section 3, we show how to compute the stochastic gradients of the total revenue with respect to the prices. In Section 4, we describe our stochastic approximation algorithm and establish its convergence. In Section 5, we present our computational experiments. In Section 6, we provide concluding remarks.

## 1 Problem Formulation

We have a set of flight legs that can be used to serve the itinerary requests that arrive randomly over time. The probability of observing a request for an itinerary at a particular time period depends on the price that we charge for the itinerary. Whenever we observe an itinerary request, if there is enough capacity on the flight legs, then we serve the itinerary request and generate a revenue reflecting the price for the itinerary. Otherwise, the itinerary request simply leaves the system. Therefore, pricing serves as the main control mechanism and we do not have the option of rejecting an itinerary request when we have enough capacity to serve the itinerary request. We are interested in finding a set of prices for the itineraries that maximize the total expected revenue.

The itinerary requests arrive over the finite planning horizon $\mathcal{T} = \{1, \ldots, \tau\}$. The flight legs depart at time period $\tau + 1$. We assume that a time period corresponds to a small enough interval of time that we observe at most one itinerary request at each time period. This is a standard modeling approach in the literature. The set of flight legs is $\mathcal{L}$ and the set of itineraries is $\mathcal{J}$. We use $\mathcal{L}_j$ to denote the set of flight legs that are used by itinerary $j$. In other words, if we serve one request for itinerary $j$, then we consume one unit of capacity on each flight leg that is in set $\mathcal{L}_j$. For notational brevity, we assume that there are no group bookings, but it is indeed possible to extend our approach to incorporate group bookings. We use $x_{it}$ to denote the remaining capacity on flight leg $i$ at the beginning of time period $t$ and $x_t = \{x_{it} : i \in \mathcal{L}\}$ to denote the vector of remaining leg capacities. Naturally, $x_1$ is a part of the problem data capturing the initial leg capacities. We use $p_j$ to denote the price for itinerary $j$ and $p = \{p_j : j \in \mathcal{J}\}$ to denote the vector of prices. If we charge price $p_j$ for itinerary $j$, then the probability of observing a request for itinerary $j$ at a time period is $\lambda_j(p_j)$. We assume that $\lambda_j(\cdot)$ is a strictly decreasing function that satisfies $\lim_{p_j \to \infty} \lambda_j(p_j) = 0$. Since we observe at most one itinerary request at a time period, we have $\sum_{j \in \mathcal{J}} \lambda_j(p_j) \leq 1$ for all $p \in \Re_+^{|\mathcal{J}|}$ and we do not observe an itinerary request at a time period with probability $1 - \sum_{j \in \mathcal{J}} \lambda_j(p_j)$. As evident from our notation, the functions $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ do not depend on the time period, but this assumption is only for notational brevity and it is straightforward to make these functions dependent on the time period. We also note that the probability of observing a request for itinerary $j$ at a particular time period depends on the price for itinerary $j$, but not on the prices for the other itineraries. This assumption is reasonable when the itineraries do not serve as substitutes of each other.

The first difficulty in developing a stochastic approximation algorithm arises from the fact that the prices for the itineraries affect the probability distributions of the itinerary requests. Therefore, every time we change the prices for the itineraries, we end up changing the probability distributions of the underlying random variables. We begin with a standard construction that decouples the prices from the probability distributions of the underlying random variables. For this purpose, we let $\pi_j = \lambda_j(0)$ and $\bar{F}_j(p_j) = \lambda_j(p_j)/\lambda_j(0)$ so that we have $\lambda_j(p_j) = \pi_j \, \bar{F}_j(p_j)$. Since $\sum_{j \in \mathcal{J}} \lambda_j(p_j) \leq 1$ for all $p \in \Re_+^{|\mathcal{J}|}$, we also have $\sum_{j \in \mathcal{J}} \pi_j \leq 1$. Furthermore, since $\lambda_j(\cdot)$ is a strictly decreasing function, $\bar{F}_j(\cdot)$ is also a strictly decreasing function and $\bar{F}_j(p_j) \in [0, 1]$ for all $p_j \in \Re_+$. We visualize $\pi_j$ as the probability that there is a customer arrival at a time period that is interested in itinerary $j$ and $1 - \bar{F}_j(\cdot)$ as the cumulative distribution function for the reservation price of a customer that is interested in itinerary $j$. In our construction, we observe a request for itinerary $j$ if there is a customer arrival that is interested in itinerary $j$ and the reservation price of the customer exceeds the price for itinerary $j$. This implies that if the price for itinerary $j$ is $p_j$, then the probability that we observe a request for itinerary $j$ at a time period is given by $\pi_j \, \bar{F}_j(p_j)$. Therefore, noting that $\lambda_j(p_j) = \pi_j \, \bar{F}_j(p_j)$, using $\{(\pi_j, \bar{F}_j(\cdot)) : j \in \mathcal{J}\}$ is conceptually equivalent to using $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$, but the advantage of using $\{(\pi_j, \bar{F}_j(\cdot)) : j \in \mathcal{J}\}$ is that we can capture all of the underlying random variables by $\omega = \{(j_t, q_t) : t \in \mathcal{T}\}$, where $j_t$ is the itinerary that is of interest to the customer arriving at time period $t$ and $q_t$ is the reservation price of the customer arriving at time period $t$. We can generate a sample of the random variable $j_t$ by using the probabilities $\{\pi_j : j \in \mathcal{J}\}$ and a sample of the random variable $q_t$ by using the cumulative distribution functions $\{1 - \bar{F}_j(\cdot) : j \in \mathcal{J}\}$. Since we have $\sum_{j \in \mathcal{J}} \pi_j \leq 1$, we may not observe a customer arrival at time period $t$, in which case, we simply set $j_t = \varnothing$ and $q_t = \infty$. We emphasize that the prices do not play any role when generating a sample of $\omega$. As a result, this construction allows us to decouple the prices for the itineraries from the probability distributions of the underlying random variables.

The second difficulty in developing a stochastic approximation algorithm arises from the fact that the number of itinerary requests that we serve at a time period is not necessarily a differentiable function of the prices and remaining leg capacities. To be more specific, we recall that the customer arriving at time period $t$ is interested in itinerary $j_t$ and we have a request for itinerary $j_t$ whenever the reservation price of the customer exceeds the price for itinerary $j_t$. On the other hand, our ability to serve an itinerary request is limited by the remaining leg capacities. In particular, noting the definition of $\mathcal{L}_{j_t}$, we can serve at most $\min_{i \in \mathcal{L}_{j_t}} \{x_{it}\}$ requests for itinerary $j_t$. Therefore, letting $\mathbf{1}(\cdot)$ be the indicator function, as a function of the remaining leg capacities, prices and realizations of the underlying random variables, we define the decision function at time period $t$ as

$$u_t(x_t, p, \omega) = \min \Big\{ \mathbf{1}(q_t \geq p_{j_t}), \min_{i \in \mathcal{L}_{j_t}} \{x_{it}\} \Big\}, \tag{1}$$

where $u_t(x_t, p, \omega)$ takes value one if we serve an itinerary request at time period $t$ and $u_t(x_t, p, \omega)$ takes value zero otherwise. Noting the indicator function and the min operator inside the curly brackets above, $u_t(x_t, p, \omega)$ is not necessarily a differentiable function of $p$ and $x_t$. This lack of differentiability ultimately prevents us from searching for a good set of prices by using the stochastic gradients of the total revenue with respect to the prices. In the next section, we develop a smoothed version of the decision function that is differentiable with respect to the prices and remaining leg capacities with probability one. The smoothed version of the decision function ensures that the total revenue is a differentiable function of the prices with probability one, which, in turn, allows us to develop a stochastic approximation algorithm that can be used to find a good set of prices.

## 2 Smoothing the Decision Function

There are two potential sources of nondifferentiability in the decision function in (1). The first source is related to the indicator function, whereas the second source is related to the min operator inside the curly brackets. We deal with the first source of nondifferentiability by assuming that we can serve fractional numbers of itinerary requests at any time period. For this purpose, we let $\theta(\cdot)$ be a differentiable and increasing function that satisfies $\lim_{q\to-\infty}\theta(q) = 0$ and $\lim_{q\to\infty}\theta(q) = 1$ and replace $\mathbf{1}(q_t \geq p_{j_t})$ in (1) with $\theta(q_t - p_{j_t})$. As $\theta(\cdot)$ approaches the step function $\mathbf{1}(\cdot \geq 0)$, we begin recovering the decision function in (1), but we note that since $\theta(\cdot)$ is differentiable, it can never be exactly equal to the step function. In addition to assuming the differentiability of $\theta(\cdot)$, we assume that $\theta(\cdot)$ and its derivative $\dot{\theta}(\cdot)$ are Lipschitz in the sense that there exist finite scalars $L_\theta$ and $L_{\dot{\theta}}$ that satisfy $|\theta(q) - \theta(s)| \leq L_\theta\,|q - s|$ and $|\dot{\theta}(q) - \dot{\theta}(s)| \leq L_{\dot{\theta}}\,|q - s|$ for all $q, s \in \Re$.

We deal with the second source of nondifferentiability in the decision function by using random perturbations of the remaining leg capacities. In particular, we let $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ be uniformly distributed random variables over the small interval $[0, \epsilon]$ and perturb the remaining leg capacities at the beginning of time period $t$ by using $\alpha_t = \{\alpha_{it} : i \in \mathcal{L}\}$. Therefore, using perturbations of the remaining leg capacities is equivalent to assuming that random but small amounts of capacity become available at the beginning of each time period. If the interval $[0, \epsilon]$ is not too large, then perturbing the remaining leg capacities in this fashion should not cause too much error. Incorporating $\theta(\cdot)$ and $\alpha_t$ into (1), we redefine the decision function as

$$u_t(x_t, p, \omega) = \min\Big\{\theta(q_t - p_{j_t}), \min_{i \in \mathcal{L}_{j_t}}\{x_{it} + \alpha_{it}\}\Big\}, \tag{2}$$

where we also redefine the underlying random variables as $\omega = \{(j_t, q_t, \alpha_t) : t \in \mathcal{T}\}$ to incorporate the capacity perturbations. We emphasize that the decision function in (2) returns fractional quantities and the implicit assumption in using this decision function is that we can serve a fraction of an itinerary request. If we assume that the capacity perturbations $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ are independent of each other and $\{(j_t, q_t) : t \in \mathcal{T}\}$, then the arguments of all of the min operators in (2) are distinct from each other with probability one. This immediately implies that the decision function in (2) is differentiable with respect to the prices and remaining leg capacities with probability one.

To define the total revenue function, we let $a_t$ be the $|\mathcal{L}|$ dimensional vector with a one in the element corresponding to flight leg $i$ if itinerary $j_t$ uses flight leg $i$ and a zero otherwise. Therefore, $a_t$ captures the leg capacities that we consume whenever we serve one request for itinerary $j_t$. In this case, as a function of the remaining leg capacities, prices and realizations of the underlying random variables, we can recursively define the total revenue function as

$$R_t(x_t, p, \omega) = p_{j_t}\, u_t(x_t, p, \omega) + R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega) \tag{3}$$

with the boundary condition that $R_{\tau+1}(x_{\tau+1}, p, \omega) = 0$. In other words, if the remaining leg capacities at the beginning of time period $t$ are given by $x_t$, the prices are given by $p$ and the realizations of the underlying random variables are given by $\omega$, then $R_t(x_t, p, \omega)$ computes the total revenue that we obtain over time periods $\{t, \ldots, \tau\}$. Of course, the computation of the total revenue is under the assumption that we can serve fractional numbers of itinerary requests and random amounts of capacity become available at the beginning of each time period. Noting that $x_1$ is a part of the problem data capturing the initial capacities on the flight legs, if the prices for the itineraries are given by $p$ and the realizations of the underlying random variables are given by $\omega$, then $R_1(x_1, p, \omega)$ computes the total revenue that we obtain over the planning horizon. In this case, using $\bar{p} = \{\bar{p}_j : j \in \mathcal{J}\}$ to denote the finite upper bounds that we impose on the feasible prices, we can solve the problem

$$\max_{p \in [0, \bar{p}]} \mathbb{E}\big\{R_1(x_1, p, \omega)\big\} \tag{4}$$

to find a set of prices that maximize the total expected revenue. In the problem above, the constraint $p \in [0, \bar{p}]$ should be understood as $p_j \in [0, \bar{p}_j]$ for all $j \in \mathcal{J}$. Imposing upper bounds on the prices should not be a huge concern from a practical perspective since the upper bounds can be arbitrarily large. We also note that problem (4) assumes that the prices are static over the whole planning horizon, but one can adopt a rolling horizon framework and periodically solve problems of the form (4) to recompute the prices. Using the fact that $u_t(x_t, p, \omega)$ is differentiable with respect to $p$ and $x_t$ with probability one, we can check by backward induction on (3) that $R_t(x_t, p, \omega)$ is also differentiable with respect to $p$ and $x_t$ with probability one, in which case, it becomes possible to solve problem (4) by using the stochastic gradients of the total revenue function with respect to the prices. In the next section, we show how to compute the stochastic gradients of the total revenue function in a tractable fashion.

## 3 Stochastic Gradients of the Total Revenue Function

In this section, we develop a recursion that can be used to compute the stochastic gradients of the total revenue function with respect to the prices. We begin by defining some notation. We let $\partial_j^P R_t(x_t, p, \omega)$ be the derivative of $R_t(\cdot, \cdot, \omega)$ with respect to the price for itinerary $j$ evaluated at remaining leg capacities $x_t$ and prices $p$. Similarly, we let $\partial_i^X R_t(x_t, p, \omega)$ be the derivative of $R_t(\cdot, \cdot, \omega)$ with respect to the remaining capacity on flight leg $i$ evaluated at remaining leg capacities $x_t$ and prices $p$. In other words, $\partial_j^P R_t(x_t, p, \omega)$ and $\partial_i^X R_t(x_t, p, \omega)$ are respectively given by

$$\partial_j^P R_t(x_t, p, \omega) = \frac{\partial R_t(z_t, r, \omega)}{\partial r_j}\bigg|_{(z_t, r) = (x_t, p)}$$
$$\partial_i^X R_t(x_t, p, \omega) = \frac{\partial R_t(z_t, r, \omega)}{\partial z_{it}}\bigg|_{(z_t, r) = (x_t, p)}.$$

We use the notation $\partial_j^P u_t(x_t, p, \omega)$ and $\partial_i^X u_t(x_t, p, \omega)$ with similar interpretations.

We proceed to derive a number of recursions that can be used to compute $\partial_j^P R_t(x_t, p, \omega)$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$. In this case, $\partial_j^P R_1(x_1, p, \omega)$ gives the stochastic gradient of the total revenue function with respect to the price for itinerary $j$. To compute $\partial_j^P R_t(x_t, p, \omega)$, we differentiate (3) with respect to the price for itinerary $j$ and use the chain rule to obtain

$$\begin{aligned}\partial_j^P R_t(x_t, p, \omega) = p_{j_t}\, \partial_j^P u_t(x_t, p, \omega) + \mathbf{1}(j = j_t)\, u_t(x_t, p, \omega) \\ + \partial_j^P R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega) \\ - \sum_{i \in \mathcal{L}_{j_t}} \partial_j^P u_t(x_t, p, \omega)\ \partial_i^X R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega),\end{aligned} \tag{5}$$

where we use the fact that $a_t$ has a one in component $i$ whenever $i \in \mathcal{L}_{j_t}$ and a zero otherwise. To compute the terms on the right side of (5), we differentiate (3) with respect to the remaining capacity on flight leg $i$ to obtain

$$\partial_i^X R_t(x_t, p, \omega) = p_{j_t}\, \partial_i^X u_t(x_t, p, \omega) + \partial_i^X R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega) \\ - \sum_{l \in \mathcal{L}_{j_t}} \partial_i^X u_t(x_t, p, \omega)\ \partial_l^X R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega). \quad (6)$$

On the other hand, differentiating (2) with respect to the price for itinerary $j$ yields

$$\partial_j^P u_t(x_t, p, \omega) = \begin{cases} -\,\dot{\theta}(q_t - p_{j_t}) & \text{if } j = j_t \text{ and } \theta(q_t - p_{j_t}) \le \min_{i \in \mathcal{L}_{j_t}} \{x_{it} + \alpha_{it}\} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

We emphasize that the expression above is accurate only in with probability one sense. In particular, if we have $j = j_t$ and $\theta(q_t - p_{j_t}) = \min_{i \in \mathcal{L}_{j_t}} \{x_{it} + \alpha_{it}\}$, then $u_t(x_t, p, \omega)$ in (2) is not necessarily differentiable with respect to the price for itinerary $j$. However, since $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ are uniformly distributed random variables that are independent of each other and $\{(j_t, q_t) : t \in \mathcal{T}\}$, the event that $j = j_t$ and $\theta(q_t - p_{j_t}) = \min_{i \in \mathcal{L}_{j_t}} \{x_{it} + \alpha_{it}\}$ occurs with probability zero. We arbitrarily set $\partial_j^P u_t(x_t, p, \omega) = -\,\dot{\theta}(q_t - p_{j_t})$ under this probability zero event. Finally, differentiating (2) with respect to the remaining capacity on flight leg $i$ yields

$$\partial_i^X u_t(x_t, p, \omega) = \begin{cases} 1 & \text{if } i \in \mathcal{L}_{j_t} \text{ and } x_{it} + \alpha_{it} \le \min \Big\{ \theta(q_t - p_{j_t}), \min_{l \in \mathcal{L}_{j_t} \setminus \{i\}} \{x_{lt} + \alpha_{lt}\} \Big\} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Similar to (7), the expression in (8) is accurate only in with probability one sense.

We can compute $\partial_j^P R_1(x_1, p, \omega)$ for all $j \in \mathcal{J}$ by simulating the evolution of the system under prices $p$ and realizations of the underlying random variables $\omega$. In this case, the remaining leg capacities at time period $t+1$ are recursively given by $X_{t+1}(p, \omega) = X_t(p, \omega) + \alpha_t - a_t\, u_t(X_t(p, \omega), p, \omega)$ with the initial condition that $X_1(p, \omega) = x_1$. We can use (7) and (8) to compute $\partial_j^P u_t(X_t(p, \omega), p, \omega)$ and $\partial_i^X u_t(X_t(p, \omega), p, \omega)$ for all $j \in \mathcal{J}$, $i \in \mathcal{L}$, $t \in \mathcal{T}$. This allows us to compute $\partial_i^X R_t(X_t(p, \omega), p, \omega)$ for all $i \in \mathcal{L}$, $t \in \mathcal{T}$ by using (6) and moving backwards over the time periods. Finally, we can compute $\partial_j^P R_t(X_t(p, \omega), p, \omega)$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$ by using (5) and moving backwards over the time periods one more time. Since $u_t(x_t, p, \omega)$ is differentiable with respect to $p$ and $x_t$ with probability one, all of the stochastic gradients exist only in with probability one sense.

## 4 Stochastic Approximation Algorithm

In this section, we analyze the following stochastic approximation algorithm to solve problem (4).

**Algorithm 1**
**Step 1.** Choose the initial prices $p^1 = \{p_j^1 : j \in \mathcal{J}\}$ such that $p^1 \in [0, \bar{p}]$. Initialize the iteration counter by setting $k = 1$.
**Step 2.** Letting $\omega^k$ be the realizations of the underlying random variables at iteration $k$, compute $\partial_j^P R_1(x_1, p^k, \omega^k)$ for all $j \in \mathcal{J}$ by using (5)-(8).

**Step 3.** Letting $\sigma^k$ be the step size parameter at iteration $k$ and using $[\cdot]^+$ to denote $\max\{\cdot, 0\}$, compute the prices $p^{k+1} = \{p_j^{k+1} : j \in \mathcal{J}\}$ at the next iteration as

$$p_j^{k+1} = \min\Big\{[p_j^k + \sigma^k\, \partial_j^P R_1(x_1, p^k, \omega^k)]^+, \bar{p}_j\Big\}.$$

**Step 4.** Increase $k$ by one and go to Step 2.

The algorithm above is a standard stochastic approximation algorithm that uses the stochastic gradient $\{\partial_j^P R_1(x_1, p^k, \omega^k) : j \in \mathcal{J}\}$ at iteration $k$. In Step 1, we choose the initial prices so that they satisfy $p_j^1 \in [0, \bar{p}_j]$ for all $j \in \mathcal{J}$. In Step 2, we sample a realization of the underlying random variables and compute the stochastic gradient $\{\partial_j^P R_1(x_1, p^k, \omega^k) : j \in \mathcal{J}\}$. Letting $\mathcal{F}^k = \{p^1, \omega^1, \ldots, \omega^{k-1}\}$ be the history of the algorithm up to iteration $k$, we assume that the distribution of $\omega^k$ conditional on $\mathcal{F}^k$ is the same as the distribution of $\omega$. In Step 3, we update the prices $p^k$ by using the step size parameter $\sigma^k$ and the stochastic gradient $\{\partial_j^P R_1(x_1, p^k, \omega^k) : j \in \mathcal{J}\}$. The role of the operator $\min\{[\cdot]^+, \bar{p}_j\}$ in this step is to ensure that $p_j^{k+1} \in [0, \bar{p}_j]$ and we can visualize this operator simply as a projection on to the interval $[0, \bar{p}_j]$. We have the next convergence result for Algorithm 1.

**Proposition 1** *Assume that the sequence of prices $\{p^k\}$ is generated by Algorithm* 1 *and the sequence of step size parameters $\{\sigma^k\}$ satisfies $\sigma^k \geq 0$ for all $k \geq 1$, $\sum_{k=1}^{\infty} \sigma^k = \infty$ and $\sum_{k=1}^{\infty} [\sigma^k]^2 < \infty$. If we use $\Theta$ to denote the set of Kuhn Tucker points for problem* (4) *and $\Theta$ is connected, then we have $p^k \to \Theta$ in probability as $k \to \infty$.*

**Proof** Propositions 2, 3 and 4 in Appendices A, B and C show that the following statements hold.

**(A.1)** We have $|\partial_j^P R_1(x_1, p, \omega)| \leq B_R^P$ with probability one for all $j \in \mathcal{J}$, $p \in [0, \bar{p}]$ for a finite scalar $B_R^P$.
**(A.2)** Using $\partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\}$ to denote the derivative of $\mathbb{E}\{R_1(x_1, \cdot, \omega)\}$ with respect to the price for itinerary $j$ evaluated at prices $p$, we have $\partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\} = \mathbb{E}\{\partial_j^P R_1(x_1, p, \omega)\}$ for all $j \in \mathcal{J}$, $p \in [0, \bar{p}]$.
**(A.3)** Using $\|\cdot\|$ to denote the Euclidean norm on $\Re^{|\mathcal{J}|}$, we have $\mathbb{E}\{|\partial_j^P R_1(x_1, p, \omega) - \partial_j^P R_1(x_1, r, \omega)|\} \leq L_R^P \|p - r\|$ for all $j \in \mathcal{J}$, $p, r \in [0, \bar{p}]$ for a finite scalar $L_R^P$.

In this case, the desired result can be shown by referring to Theorem 6.3.1 in Kushner and Clark (1978), which we briefly state in Appendix D. In particular, (A.2) and (A.3) imply that

$$\begin{aligned}\big|\partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\} - \partial_j^P \mathbb{E}\{R_1(x_1, r, \omega)\}\big| &= \big|\mathbb{E}\{\partial_j^P R_1(x_1, p, \omega) - \partial_j^P R_1(x_1, r, \omega)\}\big| \\ &\leq \mathbb{E}\{\big|\partial_j^P R_1(x_1, p, \omega) - \partial_j^P R_1(x_1, r, \omega)\big|\} \leq L_R^P \|p - r\|\end{aligned}$$

for all $p, r \in [0, \bar{p}]$. Therefore, (B.1) in Appendix D holds. Using the fact that the distribution of $\omega^k$ conditional on $\mathcal{F}^k$ is the same as the distribution of $\omega$, (A.2) implies that $\mathbb{E}\{\partial_j^P R_1(x_1, p, \omega^k) \,|\, \mathcal{F}^k\} = \mathbb{E}\{\partial_j^P R_1(x_1, p, \omega)\} = \partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\}$. Therefore, (B.2) in Appendix D holds. (A.1) and (A.2) imply that $\big|\partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\}\big| = \big|\mathbb{E}\{\partial_j^P R_1(x_1, p, \omega)\}\big| \leq B_R^P$ for all $j \in \mathcal{J}$, $p \in [0, \bar{p}]$ so that we obtain $\big|\partial_j^P R_1(x_1, p, \omega) - \partial_j^P \mathbb{E}\{R_1(x_1, p, \omega)\}\big| \leq 2\, B_R^P$ by (A.1). Therefore, (B.3) in Appendix D holds. The constraints in problem (4) are of the form $0 \leq p_j \leq \bar{p}_j$ for all $j \in \mathcal{J}$. This implies that for a given

$j \in \mathcal{J}$, the constraints $0 \leq p_j$ and $p_j \leq \bar{p}_j$ cannot be simultaneously active. Thus, the gradients of the active constraints are always linearly independent. Therefore, (B.4) in Appendix D holds. $\Box$

In the proof of Proposition 1, (A.1) implies that the stochastic gradient of the total revenue function is uniformly bounded. (A.2) implies that the expectation of the stochastic gradient of the total revenue function is equal to the gradient of the objective function of problem (4). Therefore, the expectation of the stochastic gradient is an ascent direction for problem (4). Finally, (A.3) implies that the expectation of the stochastic gradient of the total revenue function is Lipschitz when viewed as a function of the prices. It is worthwhile to note that Proposition 1 does not make an assumption about the correlation structure between the customer arrivals and reservation prices at different time periods. In particular, Proposition 1 continues to hold even when there are correlations among $\{(j_t, q_t) : t \in \mathcal{T}\}$.

An undesirable aspect of Proposition 1 is that it is difficult to verify the connectedness of the set of Kuhn Tucker points for problem (4). Theorem 6.3.1 in Kushner and Clark (1978) also provides a weaker but more technical convergence result for Algorithm 1 without assuming that the set of Kuhn Tucker points is connected. Roughly speaking, this convergence result focuses on the averages of the prices generated by Algorithm 1 over a certain number of iterations and shows that the probability that these averages lie away from the set of Kuhn Tucker points for an extended period of time diminishes as the number of iterations in Algorithm 1 increases. We do not go into the details of this technical convergence result, but emphasize that it is possible to provide a convergence result for Algorithm 1 without assuming that the set of Kuhn Tucker points for problem (4) is connected.

## 5 Computational Experiments

In this section, we compare the performances of the prices obtained by Algorithm 1 with those obtained by other benchmark strategies. We begin by describing our benchmark strategies and experimental setup. Following this, we present our computational results.

### 5.1 Benchmark Strategies

We test the performances of the following three benchmark strategies.

**Stochastic Approximation Algorithm (SAA)** SAA corresponds to the solution approach that we propose in this paper, but our practical implementation differs from the earlier development in two important aspects. First, our practical implementation of SAA divides the planning horizon into $S$ equal segments and recomputes the prices at time periods $\{1 + (s-1)\,\tau/S : s = 1, \ldots, S\}$. In particular, if the remaining leg capacities at the beginning of segment $s$ are given by $x_{1+(s-1)\tau/S}$, then we solve the problem $\max_{p \in [0,\bar{p}]} \mathbb{E}\big\{R_{1+(s-1)\tau/S}(x_{1+(s-1)\tau/S}, p, \omega)\big\}$ by using Algorithm 1. Letting $p^* = \{p^*_j : j \in \mathcal{J}\}$ be the prices that we obtain in this fashion, we use the prices $p^*$ until we reach the beginning of the next segment and recompute the prices. Second, although the earlier development in the paper is under the assumption that we can serve fractional numbers of itinerary requests, we drop this assumption when testing the performances of the prices. In particular, if the prices obtained by Algorithm 1 are given by $p^*$, then we observe a request for itinerary $j$ at a time period with probability $\lambda_j(p^*_j)$. If there is enough capacity, then we serve the itinerary request. Otherwise, the itinerary request leaves the system. In other words, we do not use $\theta(\cdot)$ and $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ when testing the performances of the prices.

In our computational experiments, we use $S = 12$. A few setup runs indicated that increasing $S$ further does not improve the performance of SAA noticeably. Our choice of $\theta(\cdot)$ is given by $\theta(q) = \frac{1}{2} - \frac{1}{2}\, e^{-\zeta\,[q]^+} + \frac{1}{2}\, e^{-\zeta\,[-q]^+}$ with $\zeta > 0$. Figure 1 plots $\theta(\cdot)$ for different values of $\zeta$ and indicates that $\theta(\cdot)$ approaches the step function as $\zeta$ increases. We use $\zeta = 0.075$ in our computational experiments. In Step 1 of Algorithm 1, we choose the initial prices as $\{\bar{p}_j/2 : j \in \mathcal{J}\}$. In Step 3 of Algorithm 1, we use the step size parameter $\sigma^k = 200/(400 + k)$. We terminate Algorithm 1 after 1,000 iterations.

**Deterministic Linear Program (DLP)** DLP is based on a deterministic linear program that is formulated under the assumption that the numbers of itinerary requests always take on their expected values. Noting that the price for itinerary $j$ lies in the interval $[0, \bar{p}_j]$, we discretize the interval $[0, \bar{p}_j]$ into $N$ pieces to obtain the price levels $\{\hat{p}_j^n : n = 1, \ldots, N\}$, where $\hat{p}_j^n = (n-1)\,\bar{p}_j/(N-1)$. At any time period in the planning horizon, if we charge the price level $\hat{p}_j^n$ for itinerary $j$ and there is enough capacity on the flight legs to serve a request for itinerary $j$, then the expected number of requests that we serve for itinerary $j$ is given by $\hat{\Lambda}_j^n = \lambda_j(\hat{p}_j^n)$ and the expected revenue that we generate from itinerary $j$ is given by $\hat{\rho}_j^n = \hat{p}_j^n\, \lambda_j(\hat{p}_j^n)$. In this case, letting $y_j^n$ be the number of time periods at which we plan to charge the price level $\hat{p}_j^n$ for itinerary $j$, we can solve the problem

$$\max \quad \sum_{j \in \mathcal{J}} \sum_{n=1}^{N} \hat{\rho}_j^n\, y_j^n \tag{9}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} \sum_{n=1}^{N} \mathbf{1}(i \in \mathcal{L}_j)\, \hat{\Lambda}_j^n\, y_j^n \leq x_{i1} \qquad \text{for all } i \in \mathcal{L} \tag{10}$$

$$\sum_{n=1}^{N} y_j^n \leq \tau \qquad \text{for all } j \in \mathcal{J} \tag{11}$$

$$y_j^n \geq 0 \qquad \text{for all } j \in \mathcal{J},\ n = 1, \ldots, N \tag{12}$$

to estimate the total expected revenue over the planning horizon. In the problem above, the first set of constraints ensure that the pricing decisions do not violate the capacities on the fight legs, whereas the second set of constraints ensure that the total number of time periods at which we use the different price levels does not exceed the number of time periods in the planning horizon. Letting $\{y_j^{n*} : n = 1, \ldots, N,\ j \in \mathcal{J}\}$ be the optimal solution to problem (9)-(12), at any time period in the planning horizon, we charge the price level $\hat{p}_j^n$ for itinerary $j$ with probability $y_j^{n*}/\tau$. With probability $1 - \sum_{n=1}^{N} y_j^{n*}/\tau$, we do not make itinerary $j$ available for purchase. Not making an itinerary available for purchase can be visualized as charging a prohibitively large price for the itinerary. DLP is based on the work of Gallego and van Ryzin (1997).

In addition to making the pricing decisions, another useful aspect of problem (9)-(12) is that it can be used to obtain an upper bound on the optimal total expected revenue and such an upper bound becomes useful when assessing the optimality gap of a suboptimal benchmark strategy, such as SAA and DLP. In particular, we can follow the same approach in Gallego and van Ryzin (1997) to show that the optimal objective value of problem (9)-(12) provides an upper bound on the total expected revenue obtained by the optimal pricing policy under the assumption that we are only allowed to use the price levels $\{\hat{p}_j^n : n = 1, \ldots, N,\ j \in \mathcal{J}\}$. In our computational experiments, we solve problem (9)-(12) with a large value for $N$ so that the price levels $\{\hat{p}_j^n : n = 1, \ldots, N\}$ virtually lie on a continuum over the interval $[0, \bar{p}_j]$. In this case, we expect that the optimal objective value of problem (9)-(12) would be an approximate upper bound on the optimal total expected revenue even when we are allowed to use any set of prices $p$ that satisfy $p_j \in [0, \bar{p}_j]$ for all $j \in \mathcal{J}$.

Similar to SAA, our practical implementation of DLP divides the planning horizon into $S$ equal segments and recomputes the prices at time periods $\{1 + (s-1)\tau/S : s = 1, \ldots, S\}$. In particular, if the remaining leg capacities at the beginning of segment $s$ are given by $x_{1+(s-1)\tau/S}$, then we replace the right side of constraints (10) with $\{x_{i,1+(s-1)\tau/S} : i \in \mathcal{L}\}$ and the right side of constraints (11) with $\tau - (s-1)\tau/S$, and solve problem (9)-(12). Letting $\{y_j^{n*} : n = 1, \ldots, N,\ j \in \mathcal{J}\}$ be the optimal solution to problem (9)-(12), we charge the price level $\hat{p}_j^n$ for itinerary $j$ with probability $y_j^{n*}/[\tau - (s-1)\tau/S]$ until we reach the beginning of the next segment and recompute the prices. We use $N = 40$ and $S = 12$ in our computational experiments. A few setup runs indicated that increasing either $N$ or $S$ further does not improve the performance of DLP significantly. When obtaining an approximate upper bound on the optimal total expected revenue, we solve problem (9)-(12) with $N = 400$.

**Capacity Allocation with Single Price (CSP)** CSP is a restricted version of DLP, where we only use the prices that maximize the immediate expected revenue from the itineraries. In particular, letting $\hat{p}_j = \operatorname{argmax}_{p_j \in [0, \bar{p}_j]} p_j \, \lambda_j(p_j)$, $\hat{\Lambda}_j = \lambda_j(\hat{p}_j)$ and $\hat{\rho}_j = \hat{p}_j \, \lambda_j(\hat{p}_j)$, we solve the problem

$$\max \quad \sum_{j \in \mathcal{J}} \hat{\rho}_j \, y_j \tag{13}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} \mathbf{1}(i \in \mathcal{L}_j) \, \hat{\Lambda}_j \, y_j \leq x_{i1} \qquad \text{for all } i \in \mathcal{L} \tag{14}$$

$$0 \leq y_j \leq \tau \qquad \text{for all } j \in \mathcal{J} \tag{15}$$

to obtain the optimal solution $\{y_j^* : j \in \mathcal{J}\}$. At any time period in the planning horizon, we use the price $\hat{p}_j$ for itinerary $j$ with probability $y_j^*/\tau$. With probability $1 - y_j^*/\tau$, we do not make itinerary $j$ available for purchase. Since CSP is a restricted version of DLP, we do not expect it to perform as well as DLP, but the goal of CSP is to show how well we can perform by choosing the prices in a myopic fashion and adjusting only the availability of the itineraries. Similar to DLP, CSP divides the planning horizon into $S$ equal segments and resolves problem (13)-(15) at the beginning of each segment. Similar to SAA and DLP, we use $S = 12$ for CSP.

## 5.2 Experimental Setup

In our computational experiments, we consider two function types that capture the relationship between the price and the probability of observing an itinerary request. In the first function type, we assume that $\lambda_j(\cdot)$ is a linear function of the form $\lambda_j(p_j) = \pi_j \, (1 - \kappa_j \, p_j)$, where $\pi_j$ can be interpreted as the probability that we observe a request for itinerary $j$ when we charge nothing for this itinerary and $\kappa_j$

as the price sensitivity. The upper bound on the price for itinerary $j$ is given by $\bar{p}_j = 1/\kappa_j$, in which case, we have $0 \leq \lambda_j(p_j) \leq \pi_j$ for all $p_j \in [0, \bar{p}_j]$. We assume that $\sum_{j \in \mathcal{J}} \pi_j \leq 1$ so that we also have $\sum_{j \in \mathcal{J}} \lambda_j(p_j) \leq 1$ for all $p \in [0, \bar{p}]$. In the second function type, we assume that $\lambda_j(\cdot)$ is an exponential function of the form $\lambda_j(p_j) = \pi_j \, e^{-\kappa_j p_j}$, where the interpretations of $\pi_j$ and $\kappa_j$ are the same as those for the linear case. The upper bound on the price for itinerary $j$ is given by $\bar{p}_j = \ln 10/\kappa_j$, in which case, we have $\pi_j/10 \leq \lambda_j(p_j) \leq \pi_j$ for all $p_j \in [0, \bar{p}_j]$. Similar to the linear case, we assume that $\sum_{j \in \mathcal{J}} \pi_j \leq 1$ so that we also have $\sum_{j \in \mathcal{J}} \lambda_j(p_j) \leq 1$ for all $p \in [0, \bar{p}]$.

In all of our test problems, we consider an airline network with one hub and $K$ spokes. There are two flight legs associated with each spoke, one of which is from the spoke to the hub and the other one is from the hub to the spoke. There is a highly price sensitive and a moderately price sensitive itinerary that connects every possible origin destination pair in the airline network. Therefore, we have $2K$ flight legs and $2K(K+1)$ itineraries, $4K$ of which use one flight leg and $2K(K-1)$ use two flight legs. Figure 2 shows the structure of the airline network with $K = 8$. The price sensitivity associated with a highly price sensitive itinerary is $\delta$ times larger than the price sensitivity associated with the corresponding moderately price sensitive itinerary. To measure the tightness of the leg capacities, we let $\hat{p}_j = \text{argmax}_{p_j \in [0, \bar{p}_j]} \, p_j \, \lambda_j(p_j)$ for all $j \in \mathcal{J}$ so that $\sum_{j \in \mathcal{J}} \mathbf{1}(i \in \mathcal{L}_j) \, \tau \, \lambda_j(\hat{p}_j)$ gives the total expected demand for the capacity on flight leg $i$ when we use the prices that maximize the immediate expected revenue from the itineraries. In this case, we measure the tightness of the leg capacities by

$$\gamma = \frac{\sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \mathbf{1}(i \in \mathcal{L}_j) \, \tau \, \lambda_j(\hat{p}_j)}{\sum_{i \in \mathcal{L}} x_{i1}}.$$

We vary $K$, $\gamma$ and $\delta$ in our test problems and label them by $(T, K, \gamma, \delta) \in \{\text{L}, \text{E}\} \times \{4, 8\} \times \{1.2, 1.6, 2.0\} \times \{2, 4, 8\}$, where the first component describes whether $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are linear or exponential functions and the other three components are as described above.

## 5.3 Computational Results

Our main computational results are shown in Tables 1 and 2. In particular, Tables 1 and 2 respectively consider the cases where $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are linear and exponential functions. The first column in these tables shows the problem characteristics. The second, third and fourth columns respectively show the total expected revenues obtained by SAA, DLP and CSP. We estimate these total expected revenues by simulating the performances of the three benchmark strategies for 100 sample paths. We use common random numbers when simulating the performances of the three benchmark strategies. The fifth and sixth columns show the percent gaps between the total expected revenues obtained by SAA and the other two benchmark strategies. The seventh column shows the upper bound on the optimal total expected revenue provided by the optimal objective value of problem (9)-(12).

The results in Table 1 indicate that SAA performs significantly better than DLP. The average performance gap between SAA and DLP is 4.3% and there are test problems where the performance gap between the two benchmark strategies can be as high as 8.9%. As expected CSP performs worse than DLP and the average performance gap between SAA and CSP is 13.6%. For all of the test problems, the performance gaps between all of the benchmark strategies turn out to be statistically significant at 95% level. Table 2 provides similar observations. The average performance gap between SAA and DLP is 8.7% and the average performance gap between SAA and CSP is 14.8%. There are test problems where the performance gap between SAA and DLP is as high as 16.4%. Overall, SAA yields significant improvements over DLP and CSP.

To illustrate the critical problem characteristics that affect the relative performances of the three benchmark strategies, Figure 3 shows the performance gaps between SAA and the other two benchmark strategies for the case where $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are exponential functions. In this figure, the thin data series plot the performance gaps between SAA and DLP, whereas the thick data series plot the performance gaps between SAA and CSP. The test problems in the horizontal axis are arranged in such a fashion that the first nine test problems involve four spokes, whereas the last nine test problems involve eight spokes. If we move from left to right within a block of nine test problems, then the leg capacities get tighter, whereas if we move from left to right within a block of three test problems, then the differences in the price sensitivities of the highly and moderately price sensitive itineraries get larger. The results in Figure 3 indicate that the performance gaps between SAA and the other two benchmark strategies tend to get larger as the leg capacities get tighter and the differences in price sensitivities of the highly and moderately price sensitive itineraries get larger. We note that if the leg capacities were unlimited, then the pricing decisions at different time periods would not interact and it would be trivially optimal to use the prices $\{\hat{p}_j : j \in \mathcal{J}\}$ with $\hat{p}_j = \text{argmax}_{p_j \in [0,\bar{p}_j]} \, p_j \, \lambda_j(p_j)$. Therefore, we intuitively expect the test problems with tighter leg capacities to be more difficult. On the other hand, as the differences in the price sensitivities of the highly and moderately price sensitive itineraries get larger, the itineraries become more diverse in terms of their responses to the pricing decisions and we need to use a richer set of prices to obtain good performance. Therefore, we also intuitively expect the test problems with larger differences in the price sensitivities to be more difficult. These observations indicate that for the test problems that we expect to be more difficult, SAA provides especially good performance when compared with the other two benchmark strategies.

In all of our computational experiments, we choose the initial prices in Step 1 of Algorithm 1 as $\{\bar{p}_j/2 : j \in \mathcal{J}\}$. This is essentially an arbitrary choice, but the performance of Algorithm 1 turns out to be relatively insensitive to the choice of the initial prices. To illustrate this behavior, Figure 4 plots $\mathbb{E}\big\{R_1(x_1, p^k, \omega)\big\}$ for two test problems as a function of the iteration counter $k$ in Algorithm 1. The charts on the left and right sides of Figure 4 respectively correspond to test problems $(\text{L}, 8, 1.6, 8)$ and $(\text{E}, 4, 1.2, 4)$. The three data series in these charts correspond to three different choices of the initial prices. In the first choice, the initial prices are $\{\bar{p}_j/2 : j \in \mathcal{J}\}$. In the second choice, the initial prices are sampled uniformly over $\{p \in \Re_+^{|\mathcal{J}|} : p \in [0, \bar{p}]\}$. In the third choice, letting $\{y_j^{n*} : n = 1, \ldots, N, \; j \in \mathcal{J}\}$ be the optimal solution to problem (9)-(12), the initial prices are $\{\sum_{n=1}^N y_j^{n*} \, \hat{p}_j^n / \tau : j \in \mathcal{J}\}$. We note that the third choice of initial prices can roughly be interpreted as the weighted average of the prices used by DLP, but this interpretation is not always correct since we do not necessarily have $\sum_{n=1}^N y_j^{n*} = \tau$ for all $j \in \mathcal{J}$. The results in Figure 4 indicate that the performance of Algorithm 1 is relatively insensitive to the choice of the initial prices. In particular, the performances of the final prices that are obtained by starting from the different initial prices are always within 0.3% of each other. These results are encouraging, but we still caution the reader that the objective function of problem (4) is not necessarily concave and the performance of Algorithm 1 can indeed depend on the choice of the initial prices. Another interesting observation from Figure 4 is that the performances of the prices obtained by Algorithm 1 stabilize after about 500 iterations. Nevertheless, due to the lack of good stopping criteria for stochastic approximation algorithms, we prefer to err on the conservative side and terminate Algorithm 1 after 1,000 iterations.

Our computational results indicate that the prices used by SAA perform significantly better than those used by DLP. In this case, a natural question is whether the prices used by SAA and DLP are indeed very different or they are simply minor adjustments to each other. Letting $\{p_j^1 : j \in \mathcal{J}\}$ be the prices obtained by Algorithm 1, $\{y_j^{n*} : n = 1, \ldots, N,\ j \in \mathcal{J}\}$ be the optimal solution to problem (9)-(12) and $p_j^2 = \sum_{n=1}^N y_j^{n*}\, \hat{p}_j^n / \tau$, Figure 5 compares the prices used by SAA and DLP by providing scatter plots of $\{(p_j^1, p_j^2) : j \in \mathcal{J}\}$ for two test problems. The charts on the left and right sides of Figure 5 respectively correspond to test problems $(\mathrm{L}, 8, 1.6, 8)$ and $(\mathrm{E}, 4, 1.2, 4)$. The results in Figure 5 indicate that the prices obtained by Algorithm 1 and problem (9)-(12) show the same general trends and they may essentially be minor adjustments to each other, but the prices obtained by Algorithm 1 still provide significant improvements over those obtained by problem (9)-(12).

The computational results in Tables 1 and 2 correspond to the case where we recompute the prices 12 times over the planning horizon, but it turns out that the performance of SAA is relatively insensitive to the number of times that we recompute the prices. Figure 6 plots the total expected revenues obtained by SAA, DLP and CSP for two test problems as a function of the number of times that we recompute the prices. The charts of the left and right sides of Figure 6 respectively correspond to test problems $(\mathrm{L}, 8, 1.6, 8)$ and $(\mathrm{E}, 4, 1.2, 4)$. In the horizontal axis, we vary the number of times that we recompute the prices over the planning horizon. The thin, thick and dashed data series respectively show the total expected revenues obtained by SAA, DLP and CSP. The results in Figure 6 indicate that the performances of DLP and CSP quickly deteriorate when we recompute the prices fewer than six times over the planning horizon, whereas the performance of SAA remains stable. Furthermore, the performance of SAA even when we compute the prices only once at the beginning of the planning horizon can be uniformly better than the performances of DLP and CSP.

For different numbers of spokes in the airline network and for different numbers of time periods in the planning horizon, Table 3 shows the CPU seconds required for Algorithm 1 to compute one set of prices. All of our computational experiments are carried out on a Pentium PC running Windows XP with Intel Xeon 2.8 GHz CPU and 4 GB RAM. The CPU seconds in Table 3 correspond to the case where we terminate Algorithm 1 after 1,000 iterations. The results in Table 3 indicate that increasing the number of spokes by a factor of two increases the CPU seconds by about a factor of three. This type of scaling is reasonable as increasing the number of spokes by a factor of two increases the number of flight legs by a factor of two and increases the number of itineraries by about a factor of four. On the other hand, increasing the number of time periods in the planning horizon by a factor of two increases the CPU seconds by about a factor of two. For the largest test problems, the CPU seconds are on the order of five to ten minutes. It turns out that the CPU seconds for DLP and CSP are consistently less than a fraction of a second and we do not give detailed CPU seconds for DLP and CSP. The CPU seconds for DLP and CSP are shorter than those for SAA by an order of magnitude, but the CPU seconds for SAA are still reasonable for practical implementation. Furthermore, noting that SAA can provide significant improvements over DLP and CSP in terms of total expected revenue, SAA appears to be a very good candidate for solving practical pricing problems.

## 6 Conclusions

In this paper, we proposed a stochastic approximation algorithm that can be used to find a good set of prices in network revenue management problems. To develop our stochastic approximation algorithm, we used a construction that decouples the prices for the itineraries from the probability distributions of the underlying random variables. To facilitate our convergence proof, we used a smoothed version of the problem, which assumes that the leg capacities are continuous and we can serve fractional numbers of itinerary requests. These modifications ensured that the total revenue is differentiable with respect to the prices with probability one and we can use the stochastic gradients of the total revenue to find a good set of prices. Our computational experiments indicated that the prices obtained by our approach perform significantly better than those obtained by a deterministic linear program and the performance gaps become especially large when the leg capacities are tight and there are large differences between the price sensitivities of the different market segments.

## References


Bashyam, S. and Fu, M. C. (1998), 'Optimizaton of $(s, S)$ inventory systems with random lead times and a service level constraint', *Management Science* **44**(12), 243–256.

Benveniste, A., Metivier, M. and Priouret, P. (1991), *Adaptive Algorithms and Stochastic Approximations*, Springer.

Bertsekas, D. P. and Tsitsiklis, J. N. (1996), *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.

Bertsimas, D. and de Boer, S. (2005), 'Simulation-based booking limits for airline revenue management', *Operations Research* **53**(1), 90–106.

Bitran, G. and Caldentey, R. (2003), 'An overview of pricing models for revenue management', *Manufacturing & Service Operations Management* **5**(3), 203–229.

Chaneton, J. and Vulcano, G. (2009), Computing bid-prices for revenue management under customer choice behavior, Technical report, New York University, Stern School of Business.

Elmaghraby, W. and Keskinocak, P. (2003), 'Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions', *Management Science* **49**(10), 1287–1309.

Feng, Y. and Gallego, G. (1995), 'Optimal starting times for end-of-season sales and optimal stopping times for promotional fares', *Management Science* **41**(8), 1371–1391.

Feng, Y. and Gallego, G. (2000), 'Perishable asset revenue management with markovian time dependent demand intensities', *Management Science* **46**(7), 641–656.

Feng, Y. and Xiao, B. (2000), 'A continuous-time yield management model with multiple prices and reversible price changes', *Management Science* **46**(5), 644–657.

Fu, M. (1994), 'Sample path derivatives for $(s, S)$ inventory systems', *Operations Research* **42**(2), 351–363.

Gallego, G. and van Ryzin, G. (1994), 'Optimal dynamic pricing of inventories with stochastic demand over finite horizons', *Management Science* **40**(8), 999–1020.

Gallego, G. and van Ryzin, G. (1997), 'A multiproduct dynamic pricing problem and its applications to yield management', *Operations Research* **45**(1), 24–41.

Glasserman, P. (1994), Perturbation analysis of production networks, *in* D. D. Yao, ed., '*Stochastic Modeling and Analysis of Manufacturing Systems*', Springer, New York, NY.

Glasserman, P. and Tayur, S. (1995), 'Sensitivity analysis for base-stock levels in multiechelon production-inventory systems', *Management Science* **41**(2), 263–281.

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

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

Kunnumkal, S. and Topaloglu, H. (2008), 'Using stochastic approximation algorithms to compute optimal base-stock levels in inventory control problems', *Operations Research* **56**(3), 646–664.

Kunnumkal, S. and Topaloglu, H. (2009), A stochastic approximation algorithm to compute bid prices for joint capacity allocation and overbooking over an airline network, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at `http://legacy.orie.cornell.edu/~huseyin/publications/publications.html`.

Kushner, H. J. and Clark, D. S. (1978), *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlang, Berlin.

Kushner, H. J. and Yin, G. G. (2003), *Stochastic Approximation and Recursive Algorithms and Applications*, Springer, New York.

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

Mahajan, S. and van Ryzin, G. (2001), 'Stocking retail assortments under dynamic customer substitution', *Operations Research* **49**(3), 334–351.

McGill, J. I. and van Ryzin, G. J. (1999), 'Revenue management: Research overview and prospects', *Transportation Science* **33**(2), 233–256.

Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer, New York, NY.

Topaloglu, H. (2008), 'A stochastic approximation method to compute bid prices in network revenue management problems', *INFORMS Journal on Computing* **20**(4), 596–610.

van Ryzin, G. and McGill, J. (2000), 'Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels', *Management Science* **46**(6), 760–775.

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

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

Zhao, W. and Zheng, Y.-S. (2000), 'Optimal dynamic pricing for perishable assets with nonhomogenous demand', *Management Science* **46**(3), 375–388.

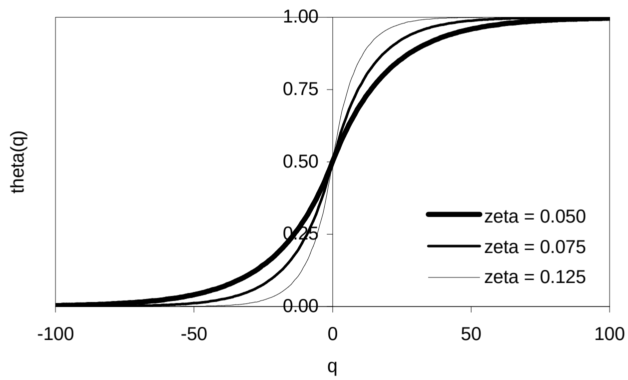


Figure 1: The function $\theta(q) = \frac{1}{2} - \frac{1}{2}e^{-\zeta\,[q]^+} + \frac{1}{2}e^{-\zeta\,[-q]^+}$ for different values of $\zeta$.

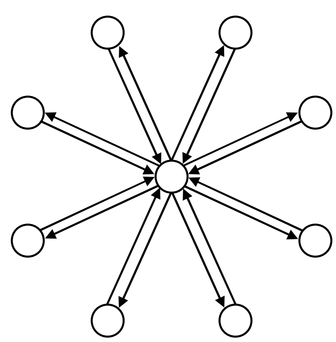

Figure 2: Structure of the airline network with eight spokes.

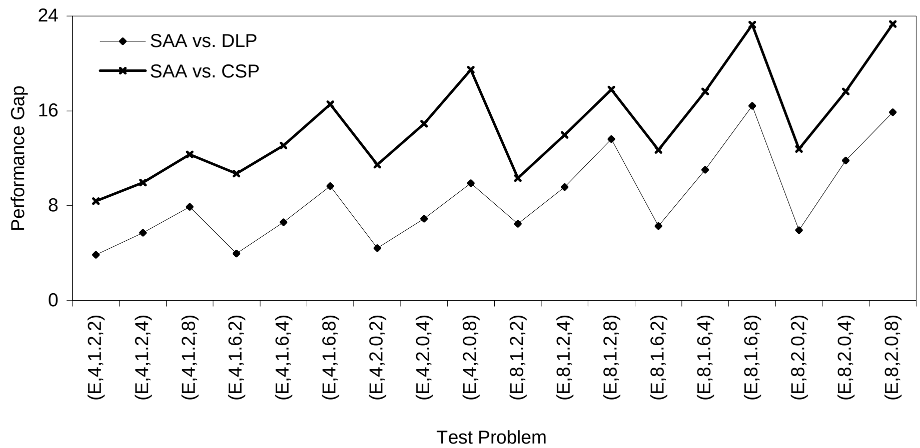


Figure 3: Performance gaps between SAA and the other two benchmark strategies for the case where $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are exponential functions.

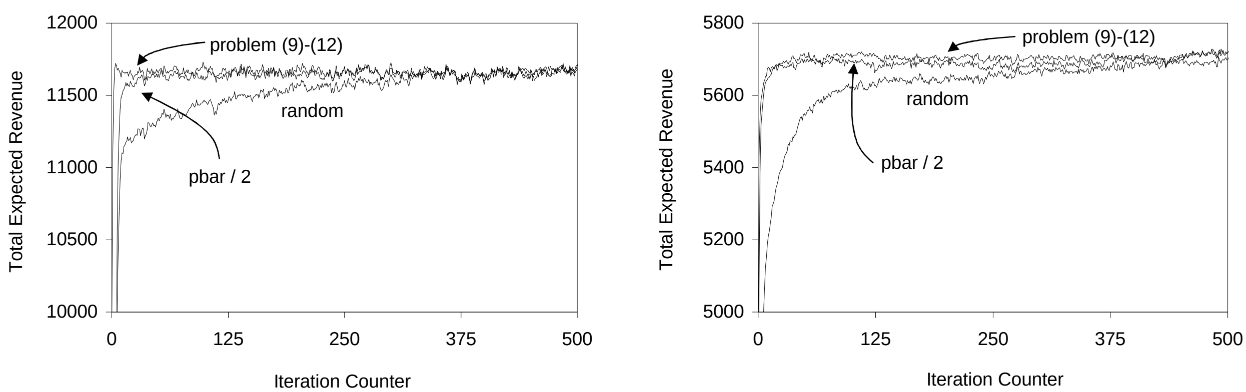


Figure 4: Total expected revenues obtained by SAA as a function of the iteration counter.

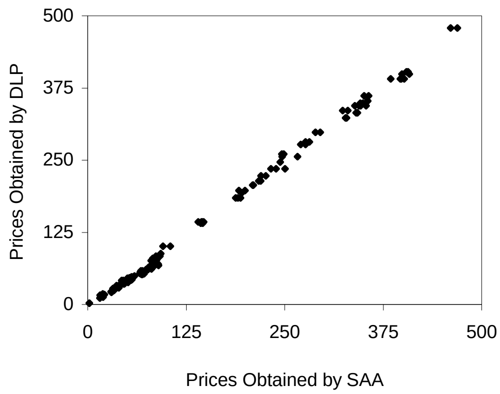


Figure 5: Comparison of the prices obtained by SAA and DLP.

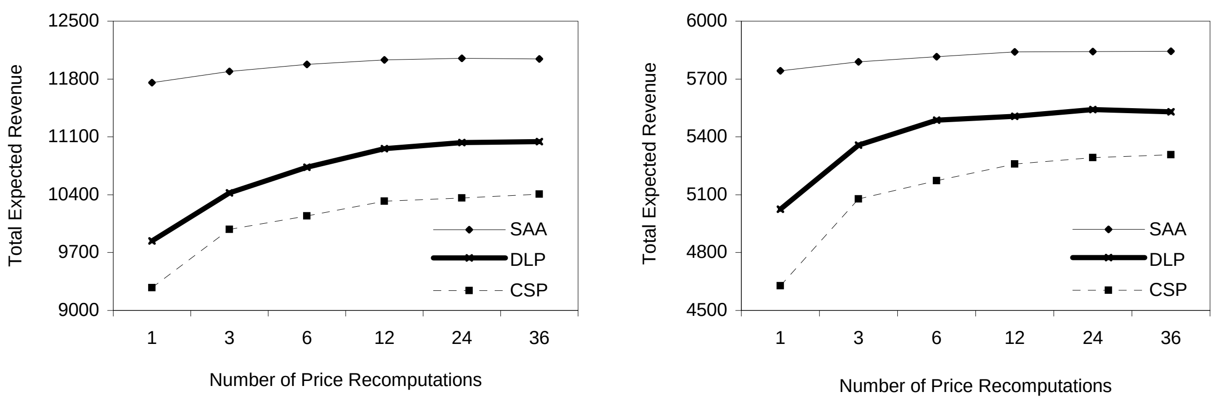


Figure 6: Total expected revenues obtained by SAA, DLP and CSP as a function of the number of times that we recompute the prices over the planning horizon.

| Problem $(T, K, \gamma, \delta)$ | Total Exp. Revenue Obtained by SAA | DLP | CSP | Perc. Gap with SAA DLP | CSP | Upper Bound |
|---|---|---|---|---|---|---|
| (L, 4, 1.2, 2) | 6,165 | 6,049 | 5,460 | 1.89 | 11.44 | 6,606 |
| (L, 4, 1.2, 4) | 8,977 | 8,654 | 8,043 | 3.60 | 10.41 | 9,538 |
| (L, 4, 1.2, 8) | 14,704 | 13,938 | 13,202 | 5.21 | 10.22 | 15,416 |
| (L, 4, 1.6, 2) | 5,355 | 5,304 | 4,456 | 0.95 | 16.79 | 5,922 |
| (L, 4, 1.6, 4) | 8,011 | 7,789 | 6,964 | 2.77 | 13.08 | 8,792 |
| (L, 4, 1.6, 8) | 13,637 | 12,963 | 11,972 | 4.94 | 12.21 | 14,640 |
| (L, 4, 2.0, 2) | 4,719 | 4,674 | 3,804 | 0.96 | 19.40 | 5,271 |
| (L, 4, 2.0, 4) | 7,240 | 7,067 | 6,273 | 2.39 | 13.35 | 8,084 |
| (L, 4, 2.0, 8) | 12,717 | 12,144 | 11,208 | 4.51 | 11.87 | 13,897 |
| (L, 8, 1.2, 2) | 5,632 | 5,454 | 5,041 | 3.17 | 10.49 | 6,273 |
| (L, 8, 1.2, 4) | 8,049 | 7,587 | 7,184 | 5.74 | 10.75 | 8,924 |
| (L, 8, 1.2, 8) | 13,118 | 11,951 | 11,482 | 8.90 | 12.47 | 14,239 |
| (L, 8, 1.6, 2) | 4,898 | 4,774 | 4,131 | 2.54 | 15.65 | 5,680 |
| (L, 8, 1.6, 4) | 7,159 | 6,800 | 6,163 | 5.02 | 13.92 | 8,281 |
| (L, 8, 1.6, 8) | 12,032 | 10,959 | 10,323 | 8.92 | 14.20 | 13,569 |
| (L, 8, 2.0, 2) | 4,354 | 4,261 | 3,545 | 2.15 | 18.58 | 5,152 |
| (L, 8, 2.0, 4) | 6,497 | 6,176 | 5,530 | 4.94 | 14.88 | 7,710 |
| (L, 8, 2.0, 8) | 11,239 | 10,241 | 9,566 | 8.88 | 14.89 | 12,972 |
| Average | | | | 4.30 | 13.59 | |

Table 1: Performances of SAA, DLP and CSP for the case where $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are linear functions.

| Problem $(T, K, \gamma, \delta)$ | Total Exp. Revenue Obtained by SAA | DLP | CSP | Perc. Gap with SAA DLP | CSP | Upper Bound |
|---|---|---|---|---|---|---|
| (E, 4, 1.2, 2) | 4,056 | 3,899 | 3,716 | 3.86 | 8.38 | 4,271 |
| (E, 4, 1.2, 4) | 5,840 | 5,506 | 5,259 | 5.72 | 9.95 | 6,149 |
| (E, 4, 1.2, 8) | 9,493 | 8,742 | 8,322 | 7.91 | 12.33 | 9,909 |
| (E, 4, 1.6, 2) | 3,697 | 3,551 | 3,301 | 3.95 | 10.69 | 4,018 |
| (E, 4, 1.6, 4) | 5,425 | 5,067 | 4,715 | 6.60 | 13.08 | 5,866 |
| (E, 4, 1.6, 8) | 9,008 | 8,138 | 7,515 | 9.66 | 16.57 | 9,608 |
| (E, 4, 2.0, 2) | 3,357 | 3,209 | 2,972 | 4.42 | 11.46 | 3,723 |
| (E, 4, 2.0, 4) | 4,956 | 4,614 | 4,217 | 6.90 | 14.92 | 5,531 |
| (E, 4, 2.0, 8) | 8,297 | 7,476 | 6,681 | 9.90 | 19.48 | 9,242 |
| (E, 8, 1.2, 2) | 3,695 | 3,456 | 3,314 | 6.47 | 10.32 | 4,050 |
| (E, 8, 1.2, 4) | 5,255 | 4,751 | 4,521 | 9.58 | 13.98 | 5,748 |
| (E, 8, 1.2, 8) | 8,457 | 7,305 | 6,952 | 13.62 | 17.79 | 9,147 |
| (E, 8, 1.6, 2) | 3,275 | 3,070 | 2,859 | 6.28 | 12.70 | 3,827 |
| (E, 8, 1.6, 4) | 4,712 | 4,192 | 3,881 | 11.04 | 17.65 | 5,501 |
| (E, 8, 1.6, 8) | 7,737 | 6,466 | 5,935 | 16.42 | 23.28 | 8,884 |
| (E, 8, 2.0, 2) | 2,890 | 2,719 | 2,520 | 5.92 | 12.79 | 3,558 |
| (E, 8, 2.0, 4) | 4,179 | 3,686 | 3,442 | 11.81 | 17.65 | 5,194 |
| (E, 8, 2.0, 8) | 6,832 | 5,747 | 5,239 | 15.89 | 23.31 | 8,547 |
| Average | | | | 8.66 | 14.80 | |

Table 2: Performances of SAA, DLP and CSP for the case where $\{\lambda_j(\cdot) : j \in \mathcal{J}\}$ are exponential functions.

| No. of. Spokes | CPU secs. for SAA |
|---|---|
| 4 | 31 |
| 8 | 75 |
| 16 | 227 |
| 32 | 722 |

| No. of. Time Per. | CPU secs. for SAA |
|---|---|
| 180 | 33 |
| 360 | 75 |
| 720 | 165 |
| 1,440 | 331 |

Table 3: CPU seconds required for Algorithm 1 to compute one set of prices.

# A Appendix: Proposition 2

**Proposition 2** *We have $|\partial_j^P R_1(x_1, p, \omega)| \le B_R^P$ with probability one for all $j \in \mathcal{J}$, $p \in [0, \bar{p}]$ for a finite scalar $B_R^P$.*

**Proof** All of the statements in the proof should be understood in with probability one sense and the proof follows from an argument that is similar to the one in Topaloglu (2008). Letting $\bar{P} = \max_{j \in \mathcal{J}} \bar{p}_j$, we begin by using induction over the time periods to show that $|\partial_i^X R_t(x_t, p, \omega)| \le \bar{P}\, 4^{\tau - t}\, |\mathcal{L}|^{\tau - t}$ for all $x_t \in \Re_+^{|\mathcal{L}|}$, $p \in [0, \bar{p}]$, $i \in \mathcal{L}$, $t \in \mathcal{T}$. Since we have $|\partial_i^X u_t(x_t, p, \omega)| \le 1$ by (8), (6) implies that

$$\big|\partial_i^X R_t(x_t, p, \omega)\big| \le \bar{P} + 2 \sum_{l \in \mathcal{L}} \big|\partial_l^X R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega)\big|.$$

Assuming that the induction hypothesis holds at time period $t+1$, the inequality above implies that $\big|\partial_i^X R_t(x_t, p, \omega)\big| \le \bar{P} + 2\, |\mathcal{L}| \big[\bar{P}\, 4^{\tau-t-1}\, |\mathcal{L}|^{\tau-t-1}\big] \le \bar{P}\, 2\, |\mathcal{L}| \big[\, 4^{\tau-t-1}\, |\mathcal{L}|^{\tau-t-1}\big] + 2\, |\mathcal{L}| \big[\bar{P}\, 4^{\tau-t-1}\, |\mathcal{L}|^{\tau-t-1}\big] = \bar{P}\, 4^{\tau-t}\, |\mathcal{L}|^{\tau-t}$ and the result holds at time period $t$. This completes the induction argument and letting $B_R^X = \bar{P}\, 4^{\tau-1}\, |\mathcal{L}|^{\tau-1}$, we have $|\partial_i^X R_t(x_t, p, \omega)| \le B_R^X$ for all $x_t \in \Re_+^{|\mathcal{L}|}$, $p \in [0, \bar{p}]$, $i \in \mathcal{L}$, $t \in \mathcal{T}$.

Since $\theta(\cdot)$ is Lipschitz with modulus $L_\theta$, its derivative is bounded by $L_\theta$. Therefore, by (7), we have $|\partial_j^P u_t(x_t, p, \omega)| \le L_\theta$. Furthermore, since $\theta(q) \in [0, 1]$ for all $q \in \Re$, we have $|u_t(x_t, p, \omega)| \le 1$ by (2). If we use the last two inequalities in (5) and note that $|\partial_i^X R_t(x_t, p, \omega)| \le B_R^X$, then we obtain $|\partial_j^P R_t(x_t, p, \omega)| \le \bar{P}\, L_\theta + 1 + |\partial_j^P R_{t+1}(x_t + \alpha_t - a_t\, u_t(x_t, p, \omega), p, \omega)| + |\mathcal{L}|\, L_\theta\, B_R^X$ for all $x_t \in \Re_+^{|\mathcal{L}|}$, $p \in [0, \bar{p}]$, $j \in \mathcal{J}$, $t \in \mathcal{T}$. Using this inequality and moving backwards over the time periods, it is straightforward to see that we have $|\partial_j^P R_t(x_t, p, \omega)| \le [\tau - t + 1]\big[1 + \bar{P}\, L_\theta + |\mathcal{L}|\, L_\theta\, B_R^X\big]$ and the result follows by letting $B_R^P = \tau\, [1 + \bar{P}\, L_\theta + |\mathcal{L}|\, L_\theta\, B_R^X]$. □

# B Appendix: Proposition 3

**Proposition 3** *Using $\partial_j^P \mathbb{E}\big\{R_1(x_1, p, \omega)\big\}$ to denote the derivative of $\mathbb{E}\big\{R_1(x_1, \cdot, \omega)\big\}$ with respect to the price for itinerary $j$ evaluated at prices $p$, we have $\partial_j^P \mathbb{E}\big\{R_1(x_1, p, \omega)\big\} = \mathbb{E}\big\{\partial_j^P R_1(x_1, p, \omega)\big\}$ for all $j \in \mathcal{J}$, $p \in [0, \bar{p}]$.*

**Proof** We have $\big|\min_{i \in \mathcal{L}_{j_t}}\{x_{it}\} - \min_{i \in \mathcal{L}_{j_t}}\{z_{it}\}\big| \le \|x_t - z_t\|$ for all $x_t, z_t \in \Re_+^{|\mathcal{L}|}$, where with slight notational abuse, we use $\|\cdot\|$ to denote the Euclidean norm on $\Re^{|\mathcal{L}|}$ as well as the Euclidean norm on $\Re^{|\mathcal{J}|}$. Therefore, $\min_{i \in \mathcal{L}_{j_t}}\{\cdot\} : \Re^{|\mathcal{L}|} \to \Re$ is Lipschitz. We assume that $\theta(\cdot) : \Re \to \Re$ is Lipschitz. Noting that $\big|\min\{q_1, q_2\} - \min\{s_1, s_2\}\big| \le |q_1 - s_1| + |q_2 - s_2|$ for all $q_1, q_2, s_1$ and $s_2 \in \Re$, $\min\{\cdot, \cdot\} : \Re^2 \to \Re$ is also Lipschitz. By Lemma 6.3.3 in Glasserman (1994), the composition of Lipschitz functions is Lipschitz. Therefore, the decision function in (2) is Lipschitz when viewed as a function of the remaining leg capacities and prices. In this case, by using backward induction on (3) and using the fact that the composition of Lipschitz functions is Lipschitz, we can check that the total revenue function is Lipschitz when viewed as a function of the remaining leg capacities and prices. This implies that $R_1(x_1, \cdot, \omega)$ is Lipschitz. Noting the discussion at the end of Section 2, $R_1(x_1, \cdot, \omega)$ is differentiable with respect to the prices with probability one. Finally, $R_1(x_1, p, \omega)$ is bounded by $\tau\, \bar{P}$ for all $p \in [0, \bar{p}]$, where $\bar{P}$ is as in the proof of Proposition 2. In this case, the result follows from Lemma 6.3.1 in Glasserman (1994), which we briefly state in Appendix E. $\square$

## C Appendix: Proposition 4

**Proposition 4** *We have* $\mathbb{E}\{|\partial^P_j R_1(x_1,p,\omega) - \partial^P_j R_1(x_1,r,\omega)|\} \leq L^P_R \, \|p-r\|$ *for all* $j \in \mathcal{J}$, $p, r \in [0, \bar{p}]$ *for a finite scalar* $L^P_R$.

We need a number of intermediate results to show Proposition 4. The next lemma shows that if the remaining leg capacities and prices change by a small amount, then the outcome of the decision function also changes by a small amount. Throughout this section, all of the statements in the proofs should be understood in with probability one sense.

**Lemma 5** *We have* $|u_t(x_t,p,\omega) - u_t(z_t,r,\omega)| \leq \|x_t - z_t\| + L_\theta \, \|p-r\|$ *with probability one for all* $t \in \mathcal{T}$, $x_t, z_t \in \Re^{|\mathcal{L}|}_+$, $p, r \in [0, \bar{p}]$.

**Proof** We note that $|\min\{q_1,q_2\} - \min\{s_1,s_2\}| \leq |q_1 - s_1| + |q_2 - s_2|$ for all $q_1, q_2, s_1$ and $s_2 \in \Re$ and $|\min_{i \in \mathcal{L}_{j_t}}\{x_{it}\} - \min_{i \in \mathcal{L}_{j_t}}\{z_{it}\}| \leq \|x_t - z_t\|$ for all $x_t, z_t \in \Re^{|\mathcal{L}|}_+$, in which case, (2) implies that $|u_t(x_t,p,\omega) - u_t(z_t,r,\omega)| \leq |\theta(q_t - p_{j_t}) - \theta(q_t - r_{j_t})| + |\min_{i \in \mathcal{L}_{j_t}}\{x_{it} + \alpha_{it}\} - \min_{i \in \mathcal{L}_{j_t}}\{z_{it} + \alpha_{it}\}| \leq L_\theta \, |p_{j_t} - r_{j_t}| + \|x_t - z_t\| \leq L_\theta \, \|p - r\| + \|x_t - z_t\|$. $\square$

In the next lemma and throughout the rest of this section, we let $x^p_t$ be the remaining leg capacities at time period $t$ when the prices are given by $p$. In other words, the random variables $\{x^p_t : t \in \mathcal{T}\}$ are recursively given by $x^p_{t+1} = x^p_t + \alpha_t - a_t \, u_t(x^p_t, p, \omega)$ with the boundary condition that $x^p_1 = x_1$. In the next lemma, we show that if the prices change by a small amount, then the remaining leg capacities also change by a small amount.

**Lemma 6** *We have* $\|x^p_t - x^r_t\| \leq L_X \, \|p - r\|$ *with probability one for all* $t \in \mathcal{T}$, $p, r \in [0, \bar{p}]$ *for a finite scalar* $L_X$.

**Proof** We use induction over the time periods to show that $\|x^p_t - x^r_t\| \leq 4^{t-1} \, (1 + |\mathcal{L}|)^{t-1} \, L_\theta \, \|p - r\|$ for all $t \in \mathcal{T}$. Assuming that the induction hypothesis holds at time period $t$, we have

$$\begin{aligned}
\|x^p_{t+1} - x^r_{t+1}\| &\leq \|x^p_t - x^r_t\| + \|a_t\| \, |u_t(x^p_t,p,\omega) - u_t(x^r_t,r,\omega)| \\
&\leq \|x^p_t - x^r_t\| + |\mathcal{L}| \big[ \, \|x^p_t - x^r_t\| + L_\theta \, \|p-r\| \, \big] \\
&\leq (1 + |\mathcal{L}|) \, 4^{t-1} \, (1 + |\mathcal{L}|)^{t-1} \, L_\theta \, \|p - r\| + |\mathcal{L}| \, L_\theta \, \|p-r\| \leq 4^t \, (1 + |\mathcal{L}|)^t \, L_\theta \, \|p-r\|,
\end{aligned}$$

where the first inequality follows from the definition of $x^p_{t+1}$, the second inequality follows from Lemma 5 and the third inequality follows from the induction hypothesis. This completes the induction argument and the result follows by letting $L_X = 4^{\tau-1} \, (1 + |\mathcal{L}|)^{\tau-1} \, L_\theta$. $\square$

In the next lemma, we show that the expectation of the derivative of the decision function with respect to the prices is Lipschitz when viewed as a function of the prices.

**Lemma 7** *We have* $\mathbb{E}\{|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)|\} \leq L_u^P \|p - r\|$ *for all* $j \in \mathcal{J}$, $t \in \mathcal{T}$, $p, r \in [0, \bar{p}]$ *for a finite scalar* $L_u^P$.

**Proof** If we have $j \neq j_t$, then (7) implies that $|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)| = 0$. Assuming that $j = j_t$, we consider the following four cases.

**Case 1.** If we have $\theta(q_t - p_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\}$ and $\theta(q_t - r_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\}$, then (7) implies that $|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)| = |\dot{\theta}(q_t - p_{j_t}) - \dot{\theta}(q_t - r_{j_t})| \leq L_{\dot{\theta}} \|p - r\|$.
**Case 2.** If we have $\theta(q_t - p_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\}$ and $\theta(q_t - r_{j_t}) > \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\}$, then noting that the derivative of $\theta(\cdot)$ is bounded by its Lipschitz modulus, we obtain $|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)| = |\dot{\theta}(q_t - p_{j_t})| \leq L_\theta$.
**Case 3.** Similar to Case 2, if we have $\theta(q_t - p_{j_t}) > \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\}$ and $\theta(q_t - r_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\}$, then we obtain $|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)| = |\dot{\theta}(q_t - r_{j_t})| \leq L_\theta$.
**Case 4.** Finally, if we have $\theta(q_t - p_{j_t}) > \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\}$ and $\theta(q_t - r_{j_t}) > \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\}$, then we obtain $|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)| = 0$.

We proceed to bound the probability of Case 2. We have $|[\theta(q_t - p_{j_t}) - x_{it}^p] - [\theta(q_t - r_{j_t}) - x_{it}^r]| \leq |\theta(q_t - p_{j_t}) - \theta(q_t - r_{j_t})| + |x_{it}^p - x_{it}^r| \leq [L_\theta + L_X] \|p - r\|$ by Lemma 6. In this case, noting that $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ are uniformly distributed over the interval $[0, \epsilon]$ and they are independent of each other and $\{(j_t, q_t) : t \in \mathcal{T}\}$, we obtain

$$\begin{aligned}
&\mathbb{P}\Big\{\theta(q_t - p_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\},\ \theta(q_t - r_{j_t}) > \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\}\Big\} \\
&\quad = \sum_{l \in \mathcal{L}} \mathbb{P}\Big\{l = \operatorname*{argmin}_{l \in \mathcal{L}_{j_t}} \{x_{it}^r + \alpha_{it}\},\ \theta(q_t - p_{j_t}) \leq \min_{i \in \mathcal{L}_{j_t}} \{x_{it}^p + \alpha_{it}\},\ \theta(q_t - r_{j_t}) > x_{lt}^r + \alpha_{lt}\Big\} \\
&\qquad \leq \sum_{l \in \mathcal{L}} \mathbb{P}\Big\{\theta(q_t - p_{j_t}) \leq x_{lt}^p + \alpha_{lt},\ \theta(q_t - r_{j_t}) > x_{lt}^r + \alpha_{lt}\Big\} \\
&\qquad\quad = \sum_{l \in \mathcal{L}} \mathbb{P}\Big\{\theta(q_t - p_{j_t}) - x_{lt}^p \ \leq \alpha_{lt} < \theta(q_t - r_{j_t}) - x_{lt}^r\Big\} \leq |\mathcal{L}|\, [L_\theta + L_X]\, \|p - r\| / \epsilon.
\end{aligned}$$

By symmetry, the same bound applies to the probability of Case 3. Combining the four cases and using the trivial bound of one on the probability of Case 1, we have

$$\mathbb{E}\{|\partial_j^P u_t(x_t^p, p, \omega) - \partial_j^P u_t(x_t^r, r, \omega)|\} \leq L_{\dot{\theta}} \|p - r\| + 2\, |\mathcal{L}|\, [L_\theta + L_X]\, L_\theta\, \|p - r\| / \epsilon$$

and the result follows by letting $L_u^P = L_{\dot{\theta}} + 2\, |\mathcal{L}|\, [L_\theta + L_X]\, L_\theta / \epsilon$. □

In the next lemma, we show that the expectation of the derivative of the decision function with respect to the remaining leg capacities is Lipschitz when viewed as a function of the prices.

**Lemma 8** *We have* $\mathbb{E}\{|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)|\} \leq L_u^X \|p - r\|$ *for all* $i \in \mathcal{L}$, $t \in \mathcal{T}$, $p, r \in [0, \bar{p}]$ *for a finite scalar* $L_u^X$.

**Proof** We follow an argument similar to the one in the proof of Lemma 7. If we have $i \notin \mathcal{L}_{j_t}$, then (8) implies that $|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)| = 0$. Assuming that $i \in \mathcal{L}_{j_t}$, there are only two cases under which $\left|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)\right|$ is not equal to zero. The first case corresponds to $x_{it}^p + \alpha_{it} \leq \min\left\{\theta(q_t - p_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^p + \alpha_{lt}\}\right\}$ and $x_{it}^r + \alpha_{it} > \min\left\{\theta(q_t - r_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^r + \alpha_{lt}\}\right\}$, whereas the second case corresponds to $x_{it}^p + \alpha_{it} > \min\left\{\theta(q_t - p_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^p + \alpha_{lt}\}\right\}$ and $x_{it}^r + \alpha_{it} \leq \min\left\{\theta(q_t - r_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^r + \alpha_{lt}\}\right\}$. If one of these two cases holds, then we have $\left|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)\right| = 1$. We proceed to bound the probability of the first case. Using the fact that $\left|\min\{q_1, q_2\} - \min\{s_1, s_2\}\right| \leq |q_1 - s_1| + |q_2 - s_2|$ for all $q_1, q_2, s_1$ and $s_2 \in \Re$ and $\left|\min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}\} - \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{z_{lt}\}\right| \leq \|x_t - z_t\|$ for all $x_t, z_t \in \Re_+^{|\mathcal{L}|}$, we have

$$\left|\left[\min\left\{\theta(q_t - p_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^p + \alpha_{lt}\}\right\} - x_{it}^p\right] - \left[\min\left\{\theta(q_t - r_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^r + \alpha_{lt}\}\right\} - x_{it}^r\right]\right|$$
$$\leq \left|\theta(q_t - p_{j_t}) - \theta(q_t - r_{j_t})\right| + 2\,\|x_t^p - x_t^r\| \leq L_\theta\,\|p - r\| + 2\,\|x_t^p - x_t^r\| \leq [L_\theta + 2\,L_X]\,\|p - r\|, \quad (16)$$

where the last inequality follows from Lemma 6. In this case, noting that $\{\alpha_{it} : i \in \mathcal{L},\ t \in \mathcal{T}\}$ are uniformly distributed over the interval $[0, \epsilon]$ and they are independent of each other and $\{(j_t, q_t) : t \in \mathcal{T}\}$, we bound the probability of the first case by

$$\mathbb{P}\Big\{x_{it}^p + \alpha_{it} \leq \min\left\{\theta(q_t - p_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^p + \alpha_{lt}\}\right\},\ x_{it}^r + \alpha_{it} > \min\left\{\theta(q_t - r_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^r + \alpha_{lt}\}\right\}\Big\}$$
$$= \mathbb{P}\Big\{\min\left\{\theta(q_t - r_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^r + \alpha_{lt}\}\right\} - x_{it}^r < \alpha_{it} \leq \min\left\{\theta(q_t - p_{j_t}), \min_{l\in\mathcal{L}_{j_t}\setminus\{i\}}\{x_{lt}^p + \alpha_{lt}\}\right\} - x_{it}^p\Big\}$$
$$\leq [L_\theta + 2\,L_X]\,\|p - r\|/\epsilon,$$

where the inequality follows from (16). By symmetry, the same bound applies to the probability of the second case. Since we have $\left|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)\right| = 1$ whenever one of the two cases holds and $\left|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)\right| = 0$ otherwise, combining the two cases yields

$$\mathbb{E}\left\{\left|\partial_i^X u_t(x_t^p, p, \omega) - \partial_i^X u_t(x_t^r, r, \omega)\right|\right\} \leq 2\,[L_\theta + 2\,L_X]\,\|p - r\|/\epsilon$$

and the result follows by letting $L_u^X = 2\,[L_\theta + 2\,L_X]/\epsilon$. □

In the next lemma, we show that the expectation of the derivative of the total revenue function with respect to the remaining leg capacities is Lipschitz when viewed as a function of the prices.

**Lemma 9** *We have* $\mathbb{E}\left\{\left|\partial_i^X R_t(x_t^p, p, \omega) - \partial_i^X R_t(x_t^r, r, \omega)\right|\right\} \leq L_R^X\,\|p - r\|$ *for all* $i \in \mathcal{L}$, $t \in \mathcal{T}$, $p, r \in [0, \bar{p}]$ *for a finite scalar* $L_R^X$.

**Proof** For notational brevity, we let $\partial_i^X R_t^p = \partial_i^X R_t(x_t^p, p, \omega)$ and $\partial_i^X u_t^p = \partial_i^X u_t(x_t^p, p, \omega)$, in which case, (6) can be written as $\partial_i^X R_t^p = p_{j_t}\ \partial_i^X u_t^p + \partial_i^X R_{t+1}^p - \sum_{l\in\mathcal{L}_{j_t}} \partial_i^X u_t^p \times \partial_l^X R_{t+1}^p$ and we obtain

$$\left|\partial_i^X R_t^p - \partial_i^X R_t^r\right| \leq \left|p_{j_t}\ \partial_i^X u_t^p - r_{j_t}\ \partial_i^X u_t^r\right| + \left|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\right|$$
$$+ \sum_{l\in\mathcal{L}_{j_t}} \left|\partial_i^X u_t^p \times \partial_l^X R_{t+1}^p - \partial_i^X u_t^r \times \partial_l^X R_{t+1}^r\right|. \quad (17)$$

Since we have $|q_1\, s_1 - q_2\, s_2| \leq |q_1|\,|s_1 - s_2| + |q_1 - q_2|\,|s_2|$ for all $q_1, q_2, s_1$ and $s_2 \in \Re$ and $\left|\partial_i^X u_t^p\right| \leq 1$ by (8), we bound the first term on the right side of (17) by

$$\left|p_{j_t}\ \partial_i^X u_t^p - r_{j_t}\ \partial_i^X u_t^r\right| \leq |p_{j_t}|\,\left|\partial_i^X u_t^p - \partial_i^X u_t^r\right| + |p_{j_t} - r_{j_t}|\,\left|\partial_i^X u_t^r\right| \leq \bar{P}\,\left|\partial_i^X u_t^p - \partial_i^X u_t^r\right| + \|p - r\|,$$

where $\bar{P}$ is as in the proof of Proposition 2. Using a similar argument, we bound the third term on the right side of (17) by

$$\begin{aligned}\big|\partial_i^X u_t^p \times \partial_l^X R_{t+1}^p - \partial_i^X u_t^r \times \partial_l^X R_{t+1}^r\big| &\leq \big|\partial_i^X u_t^p\big|\,\big|\partial_l^X R_{t+1}^p - \partial_l^X R_{t+1}^r\big| + \big|\partial_i^X u_t^p - \partial_i^X u_t^r\big|\,\big|\partial_l^X R_{t+1}^r\big| \\ &\leq \big|\partial_l^X R_{t+1}^p - \partial_l^X R_{t+1}^r\big| + B_R^X\,\big|\partial_i^X u_t^p - \partial_i^X u_t^r\big|,\end{aligned}$$

where $B_R^X$ is as in the proof of Proposition 2. Using these two bounds in (17) and taking expectations, we obtain

$$\begin{aligned}\mathbb{E}\big\{\big|\partial_i^X R_t^p - \partial_i^X R_t^r\big|\big\} \leq \bar{P}\,\mathbb{E}\big\{\big|\partial_i^X u_t^p - \partial_i^X u_t^r\big|\big\} + \|p-r\| + \mathbb{E}\big\{\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big|\big\} \\ + \sum_{l\in\mathcal{L}_{j_t}} \mathbb{E}\big\{\big|\partial_l^X R_{t+1}^p - \partial_l^X R_{t+1}^r\big|\big\} + B_R^X \sum_{l\in\mathcal{L}_{j_t}} \mathbb{E}\big\{\big|\partial_i^X u_t^p - \partial_i^X u_t^r\big|\big\},\end{aligned}$$

in which case, since Lemma 8 gives that $\mathbb{E}\big\{\big|\partial_i^X u_t^p - \partial_i^X u_t^r\big|\big\} \leq L_u^X\,\|p-r\|$, we have

$$\begin{aligned}\mathbb{E}\big\{\big|\partial_i^X R_t^p - \partial_i^X R_t^r\big|\big\} \leq \big[\bar{P}\,L_u^X + 1 + B_R^X\,|\mathcal{L}|\,L_u^X\big]\,\|p-r\| \\ + \mathbb{E}\big\{\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big|\big\} + \sum_{l\in\mathcal{L}_{j_t}} \mathbb{E}\big\{\big|\partial_l^X R_{t+1}^p - \partial_l^X R_{t+1}^r\big|\big\}. \quad (18)\end{aligned}$$

Letting $\Psi = \bar{P}\,L_u^X + 1 + B_R^X\,|\mathcal{L}|\,L_u^X$ for notational brevity, we can use induction over the time periods to show that $\mathbb{E}\big\{\big|\partial_i^X R_t^p - \partial_i^X R_t^r\big|\big\} \leq 2^{\tau-t}\,(1+|\mathcal{L}|)^{\tau-t}\,\Psi\,\|p-r\|$ for all $t\in\mathcal{T}$. In particular, noting (18), the induction hypothesis holds at time period $\tau$. Assuming that the induction hypothesis holds at time period $t+1$, (18) implies that

$$\begin{aligned}\mathbb{E}\big\{\big|\partial_i^X R_t^p - \partial_i^X R_t^r\big|\big\} &\leq \Psi\,\|p-r\| + 2^{\tau-t-1}\,(1+|\mathcal{L}|)^{\tau-t-1}\,\Psi\,\|p-r\| \\ &\qquad + |\mathcal{L}|\,2^{\tau-t-1}\,(1+|\mathcal{L}|)^{\tau-t-1}\,\Psi\,\|p-r\| \\ &\leq 2^{\tau-t-1}\,(1+|\mathcal{L}|)^{\tau-t}\,\Psi\,\|p-r\| + 2^{\tau-t-1}\,(1+|\mathcal{L}|)^{\tau-t}\,\Psi\,\|p-r\|.\end{aligned}$$

The right side of the chain of inequalities above is equal to $2^{\tau-t}\,(1+|\mathcal{L}|)^{\tau-t}\,\Psi\,\|p-r\|$ and this completes the induction argument. In this case, the result follows by letting $L_R^X = 2^{\tau-1}\,(1+|\mathcal{L}|)^{\tau-1}\,\Psi$. $\square$

We are now ready to show Proposition 4. We follow an argument similar to the one in the proof of Lemma 9. For notational brevity, we let $\partial_j^P R_t^p = \partial_j^P R_t(x_t^p,p,\omega)$ and $\partial_j^P u_t^p = \partial_j^P u_t(x_t^p,p,\omega)$, in which case, (5) can be written as $\partial_j^P R_t^p = p_{j_t}\;\partial_j^P u_t^p + \mathbf{1}(j=j_t)\,u_t(x_t^p,p,\omega) + \partial_j^P R_{t+1}^p - \sum_{i\in\mathcal{L}_{j_t}} \partial_j^P u_t^p \times \partial_i^X R_{t+1}^p$ and we obtain

$$\begin{aligned}\big|\partial_j^P R_t^p - \partial_j^P R_t^r\big| \leq \big|p_{j_t}\;\partial_j^P u_t^p - r_{j_t}\;\partial_j^P u_t^r\big| + \big|u_t(x_t^p,p,\omega) - u_t(x_t^r,r,\omega)\big| \\ + \big|\partial_j^P R_{t+1}^p - \partial_j^P R_{t+1}^r\big| + \sum_{i\in\mathcal{L}_{j_t}} \big|\partial_j^P u_t^p \times \partial_i^X R_{t+1}^p - \partial_j^P u_t^r \times \partial_i^X R_{t+1}^r\big|. \quad (19)\end{aligned}$$

Since the derivative of $\theta(\cdot)$ is bounded by its Lipschitz modulus $L_\theta$, we have $\big|\partial_j^P u_t^p\big| \leq L_\theta$ by (7). In this case, we bound the first term on the right side of (19) by

$$\big|p_{j_t}\;\partial_j^P u_t^p - r_{j_t}\;\partial_j^P u_t^r\big| \leq |p_{j_t}|\,\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big| + |p_{j_t} - r_{j_t}|\,\big|\partial_j^P u_t^r\big| \leq \bar{P}\,\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big| + L_\theta\,\|p-r\|,$$

where we use the fact that $|q_1\,s_1 - q_2\,s_2| \leq |q_1|\,|s_1-s_2| + |q_1-q_2|\,|s_2|$ for all $q_1, q_2, s_1$ and $s_2\in\Re$ and $\bar{P}$ is as in the proof of Proposition 2. Using Lemmas 5 and 6, we bound the second term on the right side of (19) by $\big|u_t(x_t^p,p,\omega)-u_t(x_t^r,r,\omega)\big| \leq [L_X+L_\theta]\,\|p-r\|$. Finally, noting $B_R^X$ in the proof of Proposition 2, we bound the fourth term on the right side of (19) by

$$\begin{aligned}\big|\partial_j^P u_t^p \times \partial_i^X R_{t+1}^p - \partial_j^P u_t^r \times \partial_i^X R_{t+1}^r\big| &\leq \big|\partial_j^P u_t^p\big|\,\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big| + \big|\partial_j^P u_t^p - \partial_j^P u_t^r\big|\,\big|\partial_i^X R_{t+1}^r\big| \\ &\leq L_\theta\,\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big| + B_R^X\,\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big|.\end{aligned}$$

Using these three bounds in (19) and taking expectations, we obtain

$$\begin{aligned}\mathbb{E}\big\{\big|\partial_j^P R_t^p - \partial_j^P R_t^r\big|\big\} \leq \bar P\,\mathbb{E}\big\{\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big|\big\} + L_\theta\,\|p-r\| \\ + [L_X+L_\theta]\,\|p-r\| + \mathbb{E}\big\{\big|\partial_j^P R_{t+1}^p - \partial_j^P R_{t+1}^r\big|\big\} \\ + L_\theta \sum_{i\in\mathcal{L}_{j_t}} \mathbb{E}\big\{\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big|\big\} + |\mathcal{L}|\,B_R^X\,\mathbb{E}\big\{\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big|\big\},\end{aligned}$$

in which case, since Lemmas 7 and 9 respectively give that $\mathbb{E}\big\{\big|\partial_j^P u_t^p - \partial_j^P u_t^r\big|\big\} \leq L_u^P\,\|p-r\|$ and $\mathbb{E}\big\{\big|\partial_i^X R_{t+1}^p - \partial_i^X R_{t+1}^r\big|\big\} \leq L_R^X\,\|p-r\|$, we have

$$\begin{aligned}\mathbb{E}\big\{\big|\partial_j^P R_t^p - \partial_j^P R_t^r\big|\big\} \leq \big[\bar P\,L_u^P + 2\,L_\theta + L_X + |\mathcal{L}|\,L_\theta\,L_R^X + |\mathcal{L}|\,L_u^P\,B_R^X\big]\,\|p-r\| \\ + \mathbb{E}\big\{\big|\partial_j^P R_{t+1}^p - \partial_j^P R_{t+1}^r\big|\big\}.\end{aligned}$$

Letting $\Phi = \bar P\,L_u^P + 2\,L_\theta + L_X + |\mathcal{L}|\,L_\theta\,L_R^X + |\mathcal{L}|\,L_u^P\,B_R^X$, it is straightforward to use the inequality above and induction over the time periods to show that $\mathbb{E}\big\{\big|\partial_j^P R_t^p - \partial_j^P R_t^r\big|\big\} \leq [\tau - t + 1]\,\Phi\,\|p-r\|$ for all $t \in \mathcal{T}$. In this case, the result follows by letting $L_R^P = \tau\,\Phi$.

## D Appendix: Theorem 6.3.1 in Kushner and Clark (1978)

For a function $g(\cdot) : \Re^n \to \Re$, a convex set $\mathcal{Z} = \{p \in \Re^n : z_i(p) \leq 0 \text{ for all } i = 1,\ldots,I\}$ and a sequence of random step directions $\{s^k\}$, we consider solving the problem $\max_{p\in\mathcal{Z}} g(p)$ by using the algorithm

$$p^{k+1} = \operatorname*{argmin}_{z\in\mathcal{Z}} \ \big\|\big[p^k + \sigma^k\,s^k\big] - z\big\|,$$

where $\{\sigma^k\}$ is a sequence of step size parameters and the operator $\operatorname{argmin}_{z\in\mathcal{Z}}\|\cdot - z\|$ projects its argument on to $\mathcal{Z}$. We assume that the following statements hold.

**(B.1)** The objective function $g(\cdot)$ is continuously differentiable.
**(B.2)** We have $\mathbb{E}\big\{s^k \,|\, p^1, s^1, \ldots, s^{k-1}\big\} = \nabla g(p^k)$ with probability one for all $k \geq 1$.
**(B.3)** We have $\|s^k - \nabla g(p^k)\| \leq M$ with probability one for all $k \geq 1$ for a finite scalar $M$.
**(B.4)** The feasible set $\mathcal{Z}$ is closed and bounded, $z_i(\cdot)$ is continuously differentiable for all $i = 1,\ldots,I$ and the gradients of the active constraints at any point in $\mathcal{Z}$ are linearly independent.

In this case, the next proposition is a somewhat specialized version of Theorem 6.3.1 in Kushner and Clark (1978). This specialized version is shown in Theorem 3 in van Ryzin and Vulcano (2008*b*).

**Proposition 10** *Assume that the sequence of points $\{p^k\}$ is generated by the algorithm above,* (B.1)-(B.4) *hold and the sequence of step size parameters $\{\sigma^k\}$ satisfies $\sigma^k \geq 0$ for all $k \geq 1$, $\sum_{k=1}^\infty \sigma^k = \infty$ and $\sum_{k=1}^\infty [\sigma^k]^2 < \infty$. If we use $\Theta$ to denote the set of Kuhn Tucker points for the problem* $\max_{p\in\mathcal{Z}} g(p)$ *and $\Theta$ is connected, then we have $p^k \to \Theta$ in probability as $k \to \infty$.*

# E Appendix: Lemma 6.3.1 in Glasserman (1994)

For a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a set $\mathcal{D} \subseteq \Re^n$ and a function $f(\cdot,\cdot) : \mathcal{D} \times \Omega \to \Re$, we assume that the following statements hold for all $p, r \in \mathcal{D}$.

**(C.1)** The function $f(\cdot, \omega)$ is differentiable at $p$ for $\mathbb{P}$-almost all values of $\omega$.
**(C.2)** There exists a finite scalar $L_f$ such that we have $\|f(p,\omega) - f(r,\omega)\| \leq L_f \|p - r\|$ for $\mathbb{P}$-almost all values of $\omega$.

In this case, the next result is from Lemma 6.3.1 in Glasserman (1994).

**Lemma 11** *Assume that there exists a finite scalar $B_f$ that satisfies $\mathbb{E}\big\{|f(p,\omega)|\big\} \leq B_f$ for all $p \in \mathcal{D}$. In this case, if* (C.1) *and* (C.2) *hold, then $\nabla \mathbb{E}\big\{f(p,\omega)\big\}$ exists and we have $\nabla \mathbb{E}\big\{f(p,\omega)\big\} = \mathbb{E}\big\{\nabla f(p,\omega)\big\}$ for all $p \in \mathcal{D}$.*