

A Stochastic Approximation Algorithm to Compute Bid Prices for Joint Capacity Allocation and Overbooking over an Airline Network

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

February 2, 2010

Abstract

In this paper, we develop a stochastic approximation algorithm to find good bid price policies for the joint capacity allocation and overbooking problem over an airline network. Our approach is based on visualizing the total expected profit as a function of the bid prices and searching for a good set of bid prices by using the stochastic gradients of the total expected profit function. We show that the total expected profit function that we use is differentiable with respect to the bid prices and derive a simple expression that can be used to compute its stochastic gradients. We show that the iterates of our stochastic approximation algorithm converge to a stationary point of the total expected profit function with probability 1. Our computational experiments indicate that the bid prices computed by our approach perform significantly better than those computed by standard benchmark strategies and the performance of our approach is relatively insensitive to the frequency with which we recompute the bid prices over the planning horizon.

The notion of bid prices forms a powerful tool for finding good policies for network revenue management problems. The fundamental idea is to associate a bid price with each flight leg that captures the opportunity cost of a seat. In this case, 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 requested itinerary; see Williamson (1992) and Talluri and van Ryzin (1998). It is known that the optimal policies are not necessarily characterized by bid prices, but the intuitive appeal and ease of implementation of bid price policies make them a popular choice in practice.

Bid prices are traditionally computed by solving a deterministic linear program. This deterministic linear program can be visualized as an approximation to the network revenue management problem that is formulated under the assumption that all random quantities are known in advance and they take on their expected values. In the deterministic linear program, there exists one capacity availability constraint for each flight leg and the right side of this constraint is the total capacity available on the flight leg. Therefore, the optimal values of the dual variables associated with the capacity availability constraints are used as bid prices. This approach for computing bid prices has seen acceptance from both academic community and industry, but it is inherently a deterministic approximation to a problem that actually takes places under uncertainty and it does not capture the temporal dynamics of the network revenue management problem accurately.

In this paper, we consider the problem of finding good bid price policies for making the capacity allocation and overbooking decisions over an airline network. We have a set of flight legs that can be used to satisfy the itinerary requests that arrive randomly over time. Whenever an itinerary request arrives, we need to decide whether to accept or reject this itinerary request. An accepted itinerary request generates a revenue and becomes a reservation. At the departure time of the flight legs, a portion of the reservations show up and we need to decide which reservations should be allowed boarding. The objective is to maximize the total expected profit, which is the difference between the total expected revenue obtained by accepting the itinerary requests and the total expected penalty cost incurred by denying boarding to the reservations. Our approach in this paper is based on visualizing the total expected profit as a function of the bid prices and searching for a good set of bid prices by using the stochastic gradients of the total expected profit function in a stochastic approximation algorithm. Since the stochastic gradients of the total expected profit function depend on the realizations of the itinerary requests and the show up decisions of the reservations, the hope is that our approach captures the stochastic aspects of the problem more accurately than the deterministic linear program.

Focusing on a class of policies that are characterized by a small number of parameters and using stochastic approximation algorithms to search for a good set of values for the parameters is a common approach in stochastic optimization. However, the nondifferentiable nature of the bid price policies creates problems when we use this idea in the network revenue management setting. In particular, if we perturb the bid price associated with a flight leg by an infinitesimal amount, then the cardinality of the subset of the itineraries for which are willing to accept a request either does not change at all or changes by a discrete amount. Therefore, an infinitesimal change in the bid prices either does not change the total expected profit at all or changes the total expected profit by a discrete amount. This implies that

the total expected profit function is not differentiable with respect to the bid prices and it becomes impossible to search for a good set of bid prices by using stochastic approximation algorithms. We resolve this difficulty by using randomized policies. Specifically, we assume that an itinerary request is accepted with a probability that depends on how much the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs that are in the requested itinerary. This modification ensures that the total expected profit function is differentiable with respect to the bid prices and we can develop a stochastic approximation algorithm that converges to a stationary point of the total expected profit function with probability 1 (w.p.1).

Our work builds on several papers that use stochastic approximation algorithms to find good policies for network revenue management problems. Bertsimas and de Boer (2005) were the first to use stochastic approximation ideas to compute good protection level policies in a setting where overbooking is not allowed and all reservations show up at the departure time. Since the protection levels are restricted to be integers and the total expected revenue is not a concave function of the protection levels, the authors work with finite differences rather than stochastic gradients to guide their search. Their approach yields good protection levels, but it does not have a convergence guarantee. The paper by van Ryzin and Vulcano (2008*b*) builds on Bertsimas and de Boer (2005), but it recognizes the need to smooth the total expected revenue function to obtain a convergent algorithm. In particular, van Ryzin and Vulcano (2008*b*) smooth the total expected revenue function by perturbing the leg capacities by small random amounts and by assuming that it is possible to serve a fraction of an itinerary request. The authors follow up on this work by van Ryzin and Vulcano (2008*a*) by showing that similar ideas can be used in a setting where the customers choose among the itineraries that are available for purchase. The papers mentioned above work with protection level policies. In contrast, Topaloglu (2008) searches for good bid price policies by using stochastic approximation algorithms. He also smoothes the total expected revenue function to obtain a convergent algorithm, but his smoothing mechanism is based on assuming that it is possible to accept a fraction of an itinerary request depending on how much the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs that are in the requested itinerary. Similar to Topaloglu (2008), Chaneton and Vulcano (2009) compute bid price policies by using stochastic approximation algorithms, but they consider the setting where the customers choose among the itineraries that are available for purchase.

Our paper complements the previous work in two important ways. First, the previous work assumes that overbooking is not allowed and all reservations show up at the departure time, whereas we explicitly address the overbooking issue. Our extension to overbooking is nontrivial and has important practical implications as overbooking plays a major role in airline operations. Second, due to the possibility of overbooking, the capacities on the flight legs do not physically limit the numbers of itinerary requests that we can accept. The leg capacities come into play only at the departure time. As a result, we can fully characterize the total numbers of accepted itinerary requests under a particular bid price policy and this allows us to construct a simple expression for the total expected profit. In contrast, if overbooking is not allowed and all reservations show up at the departure time, then it is not possible to obtain a simple expression for the total expected revenue. Instead, we need to solve a recursive equation to compute the total expected revenue. This recursive equation keeps track of the total expected revenue from the

current time period until the end of the planning horizon as a function of the remaining capacities. The computational effort required to solve the recursive equation grows exponentially with the number of flight legs. If overbooking is not allowed, then one alternative is to use the stochastic gradients of the total expected revenue on one sample path, but van Ryzin and Vulcano (2008*b*) show that the stochastic gradients of the total expected revenue on one sample path may not exist. As mentioned above, the authors deal with this difficulty by perturbing the leg capacities by small random amounts.

In this paper, we make the following research contributions. 1) We propose a new method to compute bid prices. Our approach allows overbooking and it uses the stochastic gradients of the total expected profit function in a stochastic approximation algorithm to search for a good set of bid prices. 2) We derive a simple expression for the stochastic gradient of the total expected profit function. 3) We show that the iterates of our algorithm converge to a stationary point of the total expected profit function w.p.1. 4) Computational experiments indicate that the bid prices computed by our approach perform significantly better than those computed by standard benchmark strategies. The performance gaps become especially more noticeable as the penalty cost of denying boarding to a reservation gets larger. Furthermore, the performance of our approach is relatively insensitive to the frequency with which we recompute the bid prices over the planning horizon.

The rest of the paper is organized as follows. In Section 1, we review the other related literature and compare our approach with the previous work. In Section 2, we formulate a basic optimization problem that maximizes the total expected profit by adjusting the bid prices. In Section 3, we show that the total expected profit function is differentiable with respect to the bid prices and derive a simple expression for the stochastic gradient of the total expected profit function. In Section 4, we give a stochastic approximation algorithm that can be used to find a good set of bid prices and show that the iterates of this algorithm converge to a stationary point of the total expected profit function w.p.1. In Section 5, we present our computational experiments. In Section 6, we conclude.

1 REVIEW OF RELATED LITERATURE

Despite the fact that there is substantial literature on the capacity allocation decisions under the assumption that overbooking is not allowed, the literature is surprisingly thin when one considers overbooking in the network revenue management setting. There are a few models that can be used to compute bid prices when overbooking is allowed. Bertsimas and Popescu (2003) propose computing bid prices by using a deterministic linear program. As mentioned above, an important shortcoming of the deterministic linear program is that it is formulated under the assumption that all random quantities take on their expected values. In addition to this shortcoming, the penalty cost of denying boarding to the reservations in the deterministic linear program implicitly assumes that one can oversee the whole airline network at the departure time and jointly decide which of the reservations in the whole system should be denied boarding to minimize the total penalty cost. This tends to be an optimistic assumption as it is difficult to oversee the whole airline network when deciding which reservations should be denied boarding. Karaesmen and van Ryzin (2004*a*) address these shortcomings by using a variant of the deterministic linear program that uses an approximation to the penalty cost of denying

boarding to a reservation over each flight leg and this approximation is computed by considering the probability distributions of the show up decisions. Their approach yields a bid price policy and we use their approach as a benchmark strategy in our computational experiments. Gallego and van Ryzin (1997) provide theoretical support for the deterministic linear program by showing that the control policy obtained from a variant of the deterministic linear program is asymptotically optimal as the leg capacities and the expected numbers of itinerary requests increase linearly with the same rate. Kleywegt (2001) develops a pricing and overbooking model in continuous time. The demand process in this paper is deterministic and the author solves the model by using Lagrangian duality ideas.

There are models that make the overbooking decisions over an airline network by considering the probabilistic aspects of the problem more carefully than the deterministic linear program. Karaesmen and van Ryzin (2004*b*) construct a model that is useful when making the overbooking decisions for flight legs that can serve as substitutes of each other, which is the case for multiple flights in a day that connect the same origin destination pair. This model is similar to a stochastic program, where the revenue from the accepted itinerary requests is obtained in the first stage and the penalty cost of denying boarding to the reservations is incurred in the second stage. Erdelyi and Topaloglu (2009*b*) and Kunnumkal and Topaloglu (2008*a*) develop two related overbooking models and both of these models are based on the dynamic programming formulation of the joint capacity allocation and overbooking problem over an airline network. The fundamental observation behind these two papers is that if the penalty cost of denying boarding to the reservations were given by a separable function, then the dynamic programming formulation of the problem would decompose by the itineraries. To exploit this observation, the authors develop separable approximations to the penalty cost of denying boarding to the reservations. Erdelyi and Topaloglu (2009*b*) use simulation based methods to construct the separable approximations, whereas Kunnumkal and Topaloglu (2008*a*) use the policies provided by the deterministic linear program. The policies ultimately obtained from these two models are not bid price policies. Erdelyi and Topaloglu (2009*a*) show how to decompose the dynamic programming formulation of the joint capacity allocation and overbooking problem by the flight legs. After decomposing the problem by the flight legs, the authors observe that the overbooking problem over a single flight leg is still a challenging problem and they resort to approximations to solve the single leg overbooking problem. The policy they obtain can be interpreted as a bid price policy, but the bid price associated with a flight leg is capacity dependent in the sense that it depends on how much capacity has been committed on the flight leg. There is also some work on the overbooking decisions over a single flight leg. Since our focus is on airline networks, we do not go into the details of this work and refer the reader to Section 4.4 in Talluri and van Ryzin (2005). To give a representative example, Subramanian, Stidham and Lautenbacher (1999) present an overbooking model over a single flight leg to maximize the total expected profit and their work represents the more traditional approach. More recently, Lan, Ball and Karaesmen (2007) and Ball and Queyranne (2009) study online algorithms with regret criteria for making the overbooking decision over a single flight leg.

Our approach has several distinguishing features when compared with the earlier work. Our ultimate goal is to find a good bid price policy in the traditional sense. In particular, our bid prices do not depend on the committed capacities on the flight legs and the policies we obtain are similar to the bid price

policies that are currently in use in airline operations. This is likely to make our approach more appealing to the practitioners. Furthermore, our approach is completely independent of the form of the total penalty cost of denying boarding to the reservations. As mentioned above, the form of the total penalty cost that is used by the deterministic linear program is somewhat optimistic as it is based on the assumption that one can oversee the whole airline network at the departure time and jointly decide which of the reservations in the whole system should be denied boarding to minimize the total penalty cost. As far as we are aware, different airlines use different methods to decide which reservations should be denied boarding and the fact that our approach does not depend on the form of the total penalty cost is quite beneficial from a practical perspective.

Finally, our work has strong connections with the stochastic approximation literature. Kushner and Clark (1978), Benveniste, Metivier and Priouret (1991) and Bertsekas and Tsitsiklis (1996) cover the theory of stochastic approximation in detail. In addition to the network revenue management applications in Bertsimas and de Boer (2005), van Ryzin and Vulcano (2008a), van Ryzin and Vulcano (2008b), Topaloglu (2008) and Chaneton and Vulcano (2009), L'Ecuyer and Glynn (1994) focus on queueing, Fu (1994), Glasserman and Tayur (1995), Bashyam and Fu (1998) and Kunnumkal and Topaloglu (2008b) focus on inventory control and Mahajan and van Ryzin (2001) focus on inventory assortment planning applications.

2 PROBLEM FORMULATION

We have a set of flight legs that can be used to satisfy the itinerary requests that arrive randomly over time. Whenever an itinerary request arrives, we need to decide whether to accept or reject this itinerary request. An accepted itinerary request generates a revenue and becomes a reservation, whereas a rejected itinerary request simply leaves the system. At the departure time of the flight legs, a portion of the reservations show up and we need to decide which reservations should be allowed boarding. The objective is to maximize the total expected profit over the planning horizon, which is the difference between the total expected revenue obtained by accepting the itinerary requests and the total expected penalty cost incurred by denying boarding to the reservations.

The problem takes place over the planning horizon $[0, \tau]$ and time τ corresponds to the departure time of the flight legs. The set of flight legs is \mathcal{L} and the set of itineraries is \mathcal{J} . The requests for itinerary j arrive according to a Poisson process with rate function $\{\Lambda_j(t) : t \in [0, \tau]\}$. This implies that the total number of requests for itinerary j is a Poisson random variable with mean $\lambda_j = \int_0^\tau \Lambda_j(t) dt$. We assume that the Poisson processes generating the arrivals of the requests for different itineraries are independent. If we accept a request for itinerary j , then we generate a revenue of f_j . The probability that a reservation for itinerary j shows up at the departure time is q_j . We assume that the show up decisions of different reservations are independent. If we allow boarding to a reservation for itinerary j , then we consume a_{ij} units of capacity on flight leg i . The total capacity on flight leg i is c_i . If we deny boarding to a reservation for itinerary j , then we incur a penalty cost of γ_j . For notational brevity, we assume that the reservations that do not show up at the departure time are fully refunded and the reservations are not canceled before the departure time. However, our results can easily be extended to

handle partial refunds and cancellations, as long as the cancellation decisions of different reservations are independent. Finally, we assume that $f_j \leq \gamma_j$ for all $j \in \mathcal{J}$. If there exists an itinerary j that satisfies $f_j > \gamma_j$, then it is trivially optimal to accept the requests for this itinerary, since we can always deny boarding to a reservation for itinerary j and still make a profit of $f_j - \gamma_j$.

Our goal is to use a stochastic approximation algorithm to find a good bid price policy. A bid price policy associates a bid price with each flight leg and accepts an itinerary request only if the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs that are in the requested itinerary. In other words, if we let x_i be the bid price associated with flight leg i and use x to denote the vector $\{x_i : i \in \mathcal{L}\}$, then the policy characterized by bid prices x accepts the requests for the itineraries in the set $\mathcal{J}_+(x) = \{j \in \mathcal{J} : f_j \geq \sum_{i \in \mathcal{L}} a_{ij} x_i\}$ and rejects the requests for the itineraries in the set $\mathcal{J}_-(x) = \{j \in \mathcal{J} : f_j < \sum_{i \in \mathcal{L}} a_{ij} x_i\}$. Unfortunately, such policies do not work well with stochastic approximation algorithms. In particular, for almost every $x \in \mathfrak{R}^{|\mathcal{L}|}$, if we perturb one component of x by an infinitesimal amount, then the elements of $\mathcal{J}_+(x)$ and $\mathcal{J}_-(x)$ do not change at all. This implies that the stochastic gradient of the total expected profit with respect to the bid prices is zero almost everywhere and this makes it impossible to search for a good set of bid prices by using the stochastic gradients of the total expected profit function.

We deal with this difficulty by using randomized bid price policies. In particular, the bid price policies that we consider accept an itinerary request with a probability that depends on how much the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs that are in the requested itinerary. For this purpose, we let $\theta(\cdot)$ be an increasing and differentiable function that satisfies $\lim_{h \rightarrow \infty} \theta(h) = 1$ and $\lim_{h \rightarrow -\infty} \theta(h) = 0$. In this case, the policy characterized by bid prices x accepts a request for itinerary j with probability $\theta(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$. Given that we use the bid price policy characterized by bid prices x , we let $s_j(x)$ be the number of reservations for itinerary j that show up at the departure time. Noting that the number of requests for itinerary j is a Poisson random variable with mean λ_j , we accept a request for itinerary j with probability $\theta(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$ and a reservation for itinerary j shows up with probability q_j , $s_j(x)$ is a Poisson random variable with mean $q_j \lambda_j \theta(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$. We assume that $\theta(\cdot)$ and its derivative $\dot{\theta}(\cdot)$ are Lipschitz in the sense that there exist finite scalars L_θ and $L_{\dot{\theta}}$ that satisfy $|\theta(h) - \theta(k)| \leq L_\theta |h - k|$ and $|\dot{\theta}(h) - \dot{\theta}(k)| \leq L_{\dot{\theta}} |h - k|$ for all $h, k \in \mathfrak{R}$. We use $s(x)$ to denote the vector $\{s_j(x) : j \in \mathcal{J}\}$.

Since we give full refunds to the reservations that do not show up at the departure time, the total revenue obtained by the policy characterized by bid prices x is $r(s(x)) = \sum_{j \in \mathcal{J}} f_j s_j(x)$. As a function of the reservations that show up at the departure time, we let $p(s(x))$ be the total penalty cost of denying boarding to the reservations. We leave the form of the total penalty cost unspecified for the time being. Depending on the application setting, there can be different ways to set up this cost component and our results do not depend on the form of the total penalty cost at all. We only assume that there exists a finite scalar \bar{p} such that $p(s(x)) \leq \bar{p}$ for all $x \in \mathfrak{R}^{|\mathcal{L}|}$ w.p.1. This assumption is not a huge practical concern since we can always choose \bar{p} quite large. In this case, the total profit obtained by the policy characterized by bid prices x can be written as $V(s(x)) = r(s(x)) - p(s(x))$. Therefore, we can solve the problem $\max_x \mathbb{E}\{V(s(x))\}$ to find a set of bid prices that maximize the total expected

profit. In the next section, we show that $\mathbb{E}\{V(s(x))\}$ is a differentiable function of x and derive a simple expression that can be used to compute its stochastic gradients. This result ultimately allows us to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$ by using a stochastic approximation algorithm.

3 STOCHASTIC GRADIENT OF THE TOTAL EXPECTED PROFIT

Letting e_i be the $|\mathcal{L}|$ -dimensional unit vector with a one in the component corresponding to flight leg i , our goal in this section is to show that the limit

$$\lim_{h \rightarrow 0} \frac{\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(x - h e_i))\}}{h} \quad (1)$$

exists and to derive a simple expression that can be used to compute this limit. The next lemma establishes a uniform bound on the total profit.

Lemma 1 *There exists a finite scalar \bar{V} such that $|V(s(x))| \leq \bar{V}$ for all $x \in \mathbb{R}^{|\mathcal{L}|}$ w.p.1.*

Proof Since $r(\cdot) \geq 0$, we have $V(s(x)) = r(s(x)) - p(s(x)) \geq -p(s(x)) \geq -\bar{p}$ w.p.1. If we let $\bar{c} = \max_{i \in \mathcal{L}} \{c_i\}$, then the number of reservations for itinerary j that we deny boarding always exceeds $s_j(x) - \bar{c}$ so that $p(s(x)) \geq \sum_{j \in \mathcal{J}} \gamma_j [s_j(x) - \bar{c}]$ w.p.1. Therefore, we obtain $V(s(x)) = r(s(x)) - p(s(x)) \leq \sum_{j \in \mathcal{J}} f_j s_j(x) - \sum_{j \in \mathcal{J}} \gamma_j [s_j(x) - \bar{c}] = \sum_{j \in \mathcal{J}} [f_j - \gamma_j] s_j(x) + \sum_{j \in \mathcal{J}} \gamma_j \bar{c} \leq \sum_{j \in \mathcal{J}} \gamma_j \bar{c}$ w.p.1, where the last inequality uses the fact that $f_j \leq \gamma_j$. The result follows by letting $\bar{V} = \max\{\bar{p}, \sum_{j \in \mathcal{J}} \gamma_j \bar{c}\}$. \square

For notational brevity, we let $\theta_j(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i) = q_j \lambda_j \theta(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$ so that $s_j(x)$ is a Poisson random variable with mean $\theta_j(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$. We use $\dot{\theta}_j(\cdot)$ to denote the derivative of $\theta_j(\cdot)$. We fix $x \in \mathbb{R}^{|\mathcal{L}|}$ and $i \in \mathcal{L}$, and begin with a construction that shows that $s_j(x - h e_i)$ can be visualized as the sum of two independent Poisson random variables as long as $h \geq 0$. In particular, the mean of the Poisson random variable $s_j(x - h e_i)$ is $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + a_{ij} h)$ and the Taylor series expansion of $\theta_j(\cdot)$ at the point $f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k$ yields

$$\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + a_{ij} h) = \theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) + a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h),$$

where $o(h)$ denotes any function $g(\cdot)$ that satisfies $\lim_{h \rightarrow 0} g(h)/h = 0$. Since $f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + a_{ij} h \geq f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k$ and $\theta_j(\cdot)$ is increasing, we have $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h) \geq 0$ in the expression above. Therefore, the mean of $s_j(x - h e_i)$ can be written as the sum of the two nonnegative terms $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$ and $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h)$. The first term $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$ is the mean of the Poisson random variable $s_j(x)$. We let $\Delta_j(h)$ be a Poisson random variable with mean $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h)$ that is independent of $s_j(x)$. Therefore, since the sum of two independent Poisson random variables is also a Poisson random variable, we can visualize $s_j(x - h e_i)$ as the sum of $s_j(x)$ and $\Delta_j(h)$. We assume that $\{\Delta_j(h) : j \in \mathcal{J}\}$ are independent of each other so that $\{s_j(x - h e_i) : j \in \mathcal{J}\}$ continue being independent of each other. A similar construction is also used by Karaesmen and van Ryzin (2004b). Throughout the paper, we use $\Delta(h)$ to denote the vector $\{\Delta_j(h) : j \in \mathcal{J}\}$. The next lemma follows from the fact that $\Delta_j(h)$ is a Poisson random variable.

Lemma 2 We have $\mathbb{P}\{\Delta(h) = 0\} = 1 - \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h)$ as long as $h \geq 0$.

Proof For notational brevity, we let $\alpha_{ij}(x) = a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$ for all $j \in \mathcal{J}$. Since $\Delta_j(h)$ is a Poisson random variable with mean $\alpha_{ij}(x) h + o(h)$, we have $\mathbb{P}\{\Delta_j(h) = 0\} = 1 - \alpha_{ij}(x) h + o(h)$. Expanding the product, it is easy to check that we have $\prod_{j \in \mathcal{J}} [1 - \alpha_{ij}(x) h + o(h)] = 1 - \sum_{j \in \mathcal{J}} \alpha_{ij}(x) h + o(h)$. Using the fact that $\{\Delta_j(h) : j \in \mathcal{J}\}$ are independent of each other, we obtain

$$\mathbb{P}\{\Delta(h) = 0\} = \mathbb{P}\{\Delta_j(h) = 0 \text{ for all } j \in \mathcal{J}\} = \prod_{j \in \mathcal{J}} [1 - \alpha_{ij}(x) h + o(h)] = 1 - \sum_{j \in \mathcal{J}} \alpha_{ij}(x) h + o(h)$$

and the result follows by noting the definition of $\alpha_{ij}(x)$. \square

Throughout the paper, we let ϵ_j be the $|\mathcal{J}|$ -dimensional unit vector with a one in the component corresponding to itinerary j . Similar to Lemma 2, the next lemma follows from the fact that $\Delta_j(h)$ is a Poisson random variable.

Lemma 3 We have $\mathbb{P}\{\Delta(h) = \epsilon_j\} = a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) h + o(h)$ for all $j \in \mathcal{J}$ as long as $h \geq 0$.

Proof We let $\alpha_{ij}(x)$ be as in the proof of Lemma 2. Since $\Delta_j(h)$ is a Poisson random variable with mean $\alpha_{ij}(x) h + o(h)$, we have $\mathbb{P}\{\Delta_j(h) = 1\} = \alpha_{ij}(x) h + o(h)$. Proceeding as in the proof of Lemma 2, we obtain

$$\begin{aligned} \mathbb{P}\{\Delta(h) = \epsilon_j\} &= \mathbb{P}\{\Delta_j(h) = 1 \text{ and } \Delta_l(h) = 0 \text{ for all } l \in \mathcal{J} \setminus \{j\}\} \\ &= [\alpha_{ij}(x) h + o(h)] \prod_{l \in \mathcal{J} \setminus \{j\}} [1 - \alpha_{il}(x) h + o(h)] = \alpha_{ij}(x) h + o(h). \quad \square \end{aligned}$$

The sum of the probability of the event in Lemma 2 and the probabilities of the events in Lemma 3 for all $j \in \mathcal{J}$ is $1 + o(h)$. Therefore, the total probability of the events that are not covered in Lemmas 2 and 3 is $o(h)$. The next proposition gives a simple expression for the limit in (1).

Proposition 4 We have

$$\lim_{h \rightarrow 0} \frac{\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(x - h e_i))\}}{h} = \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\} \quad (2)$$

and the expectation on the right side of the expression above is finite.

Proof We first consider the limit in (2) as h approaches to zero from right. Since we can visualize $s(x - h e_i)$ as the sum of the two independent random variables $s(x)$ and $\Delta(h)$, we have the conditional expectations $\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) \mid \Delta(h) = 0\} = 0$ and $\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) \mid \Delta(h) = \epsilon_j\} = \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}$. Noting that $V(s(x)) - V(s(x - h e_i)) \leq 2\bar{V}$ w.p.1 by Lemma 1 and the total probability of the events that are not covered in Lemmas 2 and 3 is $o(h)$, we have

$$\begin{aligned} \mathbb{E}\{V(s(x)) - V(s(x - h e_i))\} &= \mathbb{E}\{\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) \mid \Delta(h)\}\} \\ &\leq \sum_{j \in \mathcal{J}} \mathbb{P}\{\Delta(h) = \epsilon_j\} \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\} + 2\bar{V}o(h). \end{aligned}$$

Dividing both sides of the inequality above by h , using the probability in Lemma 3 and taking the limit as h approaches to zero from right, we obtain

$$\lim_{h \downarrow 0} \frac{\mathbb{E}\{V(s(x)) - V(s(x - h e_i))\}}{h} \leq \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}.$$

It is possible to see that the reverse inequality also holds by following the same argument, but using the fact that $V(s(x)) - V(s(x - h e_i)) \geq -2\bar{V}$ w.p.1. Therefore, the limit in (2) as h approaches to zero from right is given by the expression on the right side of (2). We consider the case where h approaches to zero from left in Appendix A. Finally, the expectation on the right side of (2) is finite since we have $|V(s(x)) - V(s(x - h e_i))| \leq 2\bar{V}$ w.p.1. \square

The expectation $\mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}$ can be difficult to compute. Therefore, the expression on the right side of (2) still does not provide a tractable method to differentiate $\mathbb{E}\{V(s(x))\}$ with respect to x . However, if we drop the expectation and let

$$\Psi_i(x, s(x)) = \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) \{V(s(x)) - V(s(x) + \epsilon_j)\}, \quad (3)$$

then the vector $\Psi(x, s(x)) = \{\Psi_i(x, s(x)) : i \in \mathcal{L}\}$ clearly satisfies $\mathbb{E}\{\Psi(x, s(x))\} = \nabla \mathbb{E}\{V(s(x))\}$. In other words, $\Psi(x, s(x))$ is the stochastic gradient of $\mathbb{E}\{V(s(x))\}$. Since $V(s(x)) = r(s(x)) - p(s(x)) = \sum_{j \in \mathcal{J}} f_j s_j(x) - p(s(x))$, we have $V(s(x)) - V(s(x) + \epsilon_j) = -f_j - p(s(x)) + p(s(x) + \epsilon_j)$. Therefore, we can compute $\Psi(x, s(x))$ simply by computing the total penalty cost $1 + |\mathcal{J}|$ times at $s(x)$ and $s(x) + \epsilon_j$ for all $j \in \mathcal{J}$. In the next section, we use the stochastic gradient provided by $\Psi(x, s(x))$ in a stochastic approximation algorithm to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$.

4 STOCHASTIC APPROXIMATION ALGORITHM

We propose the following algorithm to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$.

Algorithm 1

Step 1. Initialize the bid prices $x^1 = \{x_i^1 : i \in \mathcal{L}\}$ arbitrarily and initialize the iteration counter by letting $k = 1$.

Step 2. Letting $s^k(x^k) = \{s_j^k(x^k) : j \in \mathcal{J}\}$ be the numbers of reservations that show up at iteration k , compute $\Psi_i(x^k, s^k(x^k))$ for all $i \in \mathcal{L}$ by using (3).

Step 3. Letting σ^k be a step size parameter, compute the bid prices $x^{k+1} = \{x_i^{k+1} : i \in \mathcal{L}\}$ at the next iteration as $x_i^{k+1} = x_i^k + \sigma^k \Psi_i(x^k, s^k(x^k))$ for all $i \in \mathcal{L}$.

Step 4. Increase k by 1 and go to Step 2.

We let \mathcal{F}^k be the filtration generated by the random variables $\{x^1, s^1(x^1), \dots, s^{k-1}(x^{k-1})\}$ in this algorithm. With this definition of \mathcal{F}^k , x^k becomes \mathcal{F}^k -measurable. We assume that the conditional distribution of $s^k(x^k)$ given \mathcal{F}^k is the same as the conditional distribution of $s(x^k)$ given x^k . Therefore, $s^k(x^k)$ in Step 2 of Algorithm 1 can be visualized as the numbers of reservations that show up when we use bid prices x^k . In Step 3, we update the bid prices by using step size parameter σ^k and stochastic gradient $\{\Psi_i(x^k, s^k(x^k)) : i \in \mathcal{L}\}$. We have the next convergence result for Algorithm 1.

Proposition 5 Assume that the sequence of step size parameters $\{\sigma^k\}_k$ are \mathcal{F}^k -measurable and satisfy $\sigma^k \geq 0$ for all $k = 1, 2, \dots$, $\sum_{k=1}^{\infty} \sigma^k = \infty$ and $\sum_{k=1}^{\infty} [\sigma^k]^2 < \infty$ w.p.1. If the sequence of bid prices $\{x^k\}_k$ are generated by Algorithm 1, then we have $\lim_{k \rightarrow \infty} \nabla \mathbb{E}\{V(s(x^k))\} = 0$ w.p.1 and every limit point of the sequence of bid prices $\{x^k\}_k$ is a stationary point of the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$ w.p.1.

Proof We establish the result by verifying the assumptions of Proposition 4.1 from Bertsekas and Tsitsiklis (1996), which we briefly state in Appendix B for completeness. In particular, if we let $F(x) = \mathbb{E}\{V(s(x))\}$, then we have $|F(x)| \leq \bar{V}$ for all $x \in \mathfrak{R}^{|\mathcal{L}|}$ by Lemma 1. Therefore, (B.1) in Appendix B is satisfied. Noting that the conditional distribution of $s^k(x^k)$ given \mathcal{F}^k is the same as the conditional distribution of $s(x^k)$ given x^k , we have $\mathbb{E}\{\Psi(x^k, s^k(x^k)) | \mathcal{F}^k\} = \mathbb{E}\{\Psi(x^k, s(x^k)) | x^k\} = \nabla \mathbb{E}\{V(s(x^k)) | x^k\} = \nabla F(x^k)$, where the second equality follows from (2) and (3). Therefore, (B.2) in Appendix B is satisfied.

We have $|V(s(x) + \epsilon_j) - V(s(x))| \leq 2\bar{V}$ w.p.1. Since $\theta(\cdot)$ is Lipschitz with modulus L_θ , its derivative $\dot{\theta}(\cdot)$ is bounded by L_θ , which implies that $\dot{\theta}_j(\cdot)$ is bounded by $q_j \lambda_j L_\theta$. In this case, letting $\bar{a}_j = \max_{i \in \mathcal{L}} \{a_{ij}\}$, we obtain $|\Psi_i(x, s(x))| \leq |\mathcal{J}| \max_{j \in \mathcal{J}} \{\bar{a}_j q_j \lambda_j\} L_\theta 2\bar{V}$ by (3). If we use $\|\cdot\|$ to denote the Euclidean norm on $\mathfrak{R}^{|\mathcal{L}|}$, then we have $\|\Psi(x, s(x))\| \leq |\mathcal{L}| |\mathcal{J}| \max_{j \in \mathcal{J}} \{\bar{a}_j q_j \lambda_j\} L_\theta 2\bar{V}$. Therefore, (B.3) in Appendix B is satisfied.

Propositions 8 and 9 in Appendix C respectively show that $\mathbb{E}\{V(s(x))\}$ and $\mathbb{E}\{V(s(x) + \epsilon_j)\}$ for all $j \in \mathcal{J}$ are Lipschitz functions of x . By Lemma 6.3.3 in Glasserman (1994), the sum of Lipschitz functions is Lipschitz. Therefore, $\mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}$ is a Lipschitz function of x . Since $\dot{\theta}(\cdot)$ is Lipschitz, $\dot{\theta}_j(\cdot)$ is also Lipschitz. Lemma 6.3.3 in Glasserman (1994) shows that the product and composition of Lipschitz functions is Lipschitz. Therefore, $\sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}$ is a Lipschitz function of x . In this case, (2) implies that each component of $\nabla E\{V(s(x))\}$ is a Lipschitz function of x and (B.4) in Appendix B is satisfied. \square

We note that verifying (B.1) ensures that the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$ is bounded. Verifying (B.2) ensures the expectation of the step direction in Step 2 of Algorithm 1 is an ascent direction for the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$. Verifying (B.3) ensures that the step direction in Step 2 of Algorithm 1 is bounded. Finally, verifying (B.4) ensures that each component of the gradient of the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$ is Lipschitz.

It is easy to check that the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$ is not necessarily concave and the stationary point mentioned in Proposition 5 may be a local maximum, a saddle point, or even a local minimum. In particular, since we have $\mathbb{E}\{r(s(x))\} = \sum_{j \in \mathcal{J}} f_j \theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$, the total expected profit inherits the properties of $\{\theta_j(\cdot) : j \in \mathcal{J}\}$. Many possible choices for $\{\theta_j(\cdot) : j \in \mathcal{J}\}$, including the one that we use in our computational experiments, are not concave. In our computational experiments, we dwell on the question of how we choose $\{\theta_j(\cdot) : j \in \mathcal{J}\}$ and test the robustness of Algorithm 1 to different choices for $\{\theta_j(\cdot) : j \in \mathcal{J}\}$. It turns out that there are a wide range of choices for $\{\theta_j(\cdot) : j \in \mathcal{J}\}$ under which Algorithm 1 continues to perform quite well.

5 COMPUTATIONAL EXPERIMENTS

In this section, we compare the performances of the bid prices computed by Algorithm 1 with those of the bid prices computed by using other strategies.

5.1 EXPERIMENTAL SETUP

We consider an airline network that serves N spokes through a single hub. There are two flight legs associated with each spoke. One of these flights is from the hub to the spoke and the other one is from the spoke to the hub. There is a high fare and a low fare itinerary associated with each possible origin destination pair. Therefore, we have $2N$ flight legs and $2N(N + 1)$ itineraries, $4N$ of which include one flight leg and $2N(N - 1)$ of which include two flight legs. Figure 1 shows the structure of the airline network for the case where $N = 8$. To come up with the revenues associated with the itineraries, for each low fare itinerary j , we generate ξ_j from the exponential distribution with mean one. If itinerary j includes one flight leg, then we set the revenue associated with the itinerary as $f_j = 75 + 75 \xi_j$, whereas if itinerary j includes two flight legs, then we set the revenue associated with the itinerary as $f_j = 150 + 150 \xi_j$. Therefore, the average revenue associated with a low fare itinerary that includes one flight leg is 150 and the average revenue associated with a low fare itinerary that includes two flight legs is 300. In all of our test problems, the revenue associated with a high fare itinerary is three times larger than the revenue associated with the corresponding low fare itinerary. The penalty cost of denying boarding to a reservation for an itinerary is fixed at $\gamma_j = \gamma$ for all $j \in \mathcal{J}$, where γ is a parameter that we vary. Since the total expected demand for the capacity on flight leg i is $\sum_{j \in \mathcal{J}} a_{ij} q_j \lambda_j$, we measure the tightness of the leg capacities by using

$$\alpha = \frac{\sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} a_{ij} q_j \lambda_j}{\sum_{i \in \mathcal{L}} c_i}.$$

The probabilities that a reservation for a high fare and a low fare itinerary shows up at the departure time are respectively q^h and q^l . These probabilities do not depend on the origin and destination locations of the itinerary. We have eight spokes in all of our test problems. The total expected number of itinerary requests over the planning horizon is about 360. We label our test problems by the tuple $(\gamma, \alpha, q^h, q^l)$ and use $\gamma \in \{600, 900, 1200, 1500\}$, $\alpha \in \{1.05, 1.2, 1.6\}$, $q^h \in \{0.7, 0.9\}$ and $q^l \in \{0.7, 0.9\}$. This setup provides 48 test problems for our computational experiments.

If we let z_j be the number of reservations for itinerary j that we deny boarding and use z to denote the vector $\{z_j : j \in \mathcal{J}\}$, then we compute the total penalty cost by solving the problem

$$\min_{z \in \mathbb{Z}_+^{|\mathcal{J}|}} \sum_{j \in \mathcal{J}} \gamma_j z_j \tag{4}$$

$$\text{subject to } \sum_{j \in \mathcal{J}} a_{ij} [s_j(x) - z_j] \leq c_i \quad \text{for all } i \in \mathcal{L} \tag{5}$$

$$z_j \leq s_j(x) \quad \text{for all } j \in \mathcal{J} \tag{6}$$

and letting $p(s(x))$ be the optimal objective value of the problem above. Constraints (5) in problem (4)-(6) ensure that the reservations that we allow boarding do not violate the leg capacities, whereas

constraints (6) ensure that the reservations that we deny boarding do not exceed the reservations that show up. To make sure that there exists a finite scalar \bar{p} that satisfies $p(s(x)) \leq \bar{p}$ for all $x \in \mathfrak{R}^{|\mathcal{L}|}$ w.p.1, we round the optimal objective value of problem (4)-(6) down to \bar{p} whenever it exceeds \bar{p} . This is not a concern from a practical perspective since we can always choose \bar{p} quite large.

5.2 BENCHMARK STRATEGIES

We compare the performances of the following five benchmark strategies.

Stochastic Approximation with Randomized Policy (SAR) This is the approach developed in this paper, but our practical implementation divides the planning horizon into five equal segments and recomputes the bid prices at the beginning of each segment. In particular, if we let u_{jt} be the number of reservations for itinerary j at time t , then among these reservations, the portion that shows up at the departure time is given by a binomial random variable $b_{jt}(u_{jt})$ with parameters u_{jt} and q_j . On the other hand, if we use the policy characterized by bid prices x to accept or reject the requests for itinerary j over the time interval $[t, \tau]$, then among these reservations, the portion that shows up at the departure time is given by a Poisson random variable $s_{jt}(x)$ with mean $q_j \int_t^\tau \Lambda_j(h) dh \theta(f_j - \sum_{i \in \mathcal{L}} a_{ij} x_i)$. In this case, letting $b_t(u_t)$ and $s_t(x)$ respectively be the vectors $\{b_{jt}(u_{jt}) : j \in \mathcal{J}\}$ and $\{s_{jt}(x) : j \in \mathcal{J}\}$, the numbers of reservations that show up at the departure time are given by $b_t(u_t) + s_t(x)$ and we can solve the problem $\max_x \mathbb{E}\{V(b_t(u_t) + s_t(x))\}$ to find a good set of bid prices at time t . Therefore, SAR solves the problem $\max_x \mathbb{E}\{V(b_{(k-1)\tau/5}(u_{(k-1)\tau/5}) + s_{(k-1)\tau/5}(x))\}$ to recompute the bid prices at the beginning of segment k and uses these bid prices until it reaches the beginning of the next segment.

We use the step size parameter $\sigma^k = 20/(40+k)$ in Step 3 of Algorithm 1 and terminate Algorithm 1 after 2,500 iterations. Letting $[h]^+ = \max\{h, 0\}$, our choice of $\theta(\cdot)$ is given by $\theta(h) = \frac{1}{2} - \frac{1}{2} e^{-\zeta [h]^+} + \frac{1}{2} e^{-\zeta [-h]^+}$ with $\zeta > 0$. Figure 2 plots our choice of $\theta(\cdot)$ for different values of ζ and shows that this function looks like the step function as ζ approaches to infinity. Therefore, the distinction between using a randomized and a nonrandomized policy diminishes as ζ approaches to infinity. After a few setup runs, we settle on $\zeta = 0.075$ and use this value for ζ in all of our test problems. In our computational experiments, we test the robustness of Algorithm 1 to different choices for ζ . To give a feel for where $\theta(\cdot)$ in Figure 2 starts to flatten relative to the revenues associated with the itineraries, we note that the revenues in our test problems range roughly over the interval $[75, 1,200]$.

Stochastic Approximation with Deterministic Policy (SAD) Managers may not like the idea of flipping a coin to decide whether an itinerary request should be accepted or rejected. Therefore, the randomized policy used by SAR may not always be appropriate. SAD attempts to overcome this difficulty by using the bid prices computed by Algorithm 1 in a deterministic decision rule. In particular, letting x be the bid prices computed by Algorithm 1, SAD simply accepts the requests for itinerary j if and only if $f_j \geq \sum_{i \in \mathcal{L}} a_{ij} x_i$. We emphasize that the stochastic gradient in Step 2 of Algorithm 1 is computed under the assumption that we use a randomized policy. Therefore, SAD should be visualized only as a heuristic extension of SAR and its performance may be better or worse than that of SAR. Similar to SAR, SAD recomputes the bid prices five times over the planning horizon.

Deterministic Linear Program (DLP) A traditional approach for computing bid prices is to solve a linear program that is formulated under the assumption that the arrivals of the itinerary requests and the show up decisions of the reservations take on their expected values. In particular, if we let w_j be the number of requests for itinerary j that we plan to accept over the planning horizon and z_j be the number of reservations for itinerary j that we plan to deny boarding, and use w and z to respectively denote the vectors $\{w_j : j \in \mathcal{J}\}$ and $\{z_j : j \in \mathcal{J}\}$, then this linear program can be written as

$$\max_{(w,z) \in \mathbb{R}_+^{2|\mathcal{J}|}} \sum_{j \in \mathcal{J}} f_j q_j w_j - \sum_{j \in \mathcal{J}} \gamma_j z_j \quad (7)$$

$$\text{subject to } \sum_{j \in \mathcal{J}} a_{ij} [q_j w_j - z_j] \leq c_i \quad \text{for all } i \in \mathcal{L} \quad (8)$$

$$z_j \leq q_j w_j \quad \text{for all } j \in \mathcal{J} \quad (9)$$

$$w_j \leq \lambda_j \quad \text{for all } j \in \mathcal{J}. \quad (10)$$

The problem above assumes that if we accept w_j requests for itinerary j , then $q_j w_j$ reservations show up at the departure time. Constraints (8) and (9) in problem (7)-(10) are analogous to constraints (5) and (6) in problem (4)-(6). Constraints (10) ensure that the itinerary requests that we accept do not exceed the expected numbers of itinerary requests. The form of problem (7)-(10) is tightly linked to the fact that we compute the total penalty cost by solving problem (4)-(6). If the total penalty cost were computed in a different manner, then it may not be possible to formulate a linear program.

Letting $\{x_i : i \in \mathcal{L}\}$ be the optimal values of the dual variables associated with constraints (8) in problem (7)-(10), x_i captures the opportunity cost of a seat on flight leg i . Therefore, DLP uses x_i as the bid price associated with flight leg i and accepts the requests for itinerary j if and only if $f_j \geq \sum_{i \in \mathcal{L}} a_{ij} x_i$. Furthermore, Erdelyi and Topaloglu (2009b) show that the optimal objective value of problem (7)-(10) provides an upper bound on the optimal total expected profit. Such upper bounds become useful when assessing the optimality gap of a suboptimal control policy such as the ones used by SAR, SAD or DLP. Similar to SAR and SAD, our practical implementation of DLP recomputes the bid prices five times over the planning horizon. In particular, if the numbers of reservations at time t is given by $\{u_{jt} : j \in \mathcal{J}\}$, then DLP replaces the right side of constraints (8) with $\{c_i - \sum_{j \in \mathcal{J}} a_{ij} q_j u_{jt} : i \in \mathcal{L}\}$, the right side of constraints (9) with $\{q_j [w_j + u_{jt}] : j \in \mathcal{J}\}$ and the right side of constraints (10) with $\{\int_t^\tau \Lambda_j(h) dh : j \in \mathcal{J}\}$, and solves problem (7)-(10) to recompute the bid prices at time t .

Scaled Capacities with a Numerical Search (SCN) One natural criticism for DLP is that it makes its plans under the assumption that the show up decisions of the reservations take on their expected values. However, depending on the tradeoff between the penalty costs and the probabilities of showing up, one may want to be more or less aggressive than what the expected values of the show up decisions suggest. SCN attempts to remedy this shortcoming by scaling the leg capacities by a constant factor. In particular, letting β be a scaling factor, SCN replaces the right side of constraints (8) in problem (7)-(10) with $\{\beta c_i : i \in \mathcal{L}\}$ and solves this problem to compute the bid prices. For each test problem, we find the best value of β through an exhaustive numerical search over the interval $[0.6, 1.4]$ and report the results corresponding to the best value of β . By following the same approach that DLP uses, SCN recomputes the bid prices five times over the planning horizon.

Virtual Capacities with an Economic Model (VCE) The idea behind VCE is to construct a number of simple cost functions that capture the penalty costs incurred over the different flight legs and to incorporate these cost functions into a linear program similar to the one in (7)-(10). To construct the cost functions, VCE associates a virtual capacity with each flight leg and proceeds with three assumptions. First, if the virtual capacity on flight leg i is u_i , then we sell exactly u_i seats on flight leg i . Second, if a reservation uses the capacities on multiple flight legs, then we can allow boarding to this reservation over one flight leg, while denying boarding to the same reservation over another flight leg. Third, the probability that a reservation shows up over a flight leg and the penalty cost of denying boarding to a reservation over a flight leg are known. If we let Q_i be the probability that a reservation shows up over flight leg i and $B_i(u_i)$ be a binomial random variable with parameters u_i and Q_i , then the first and third assumptions imply that the number of reservations that show up over flight leg i is given by $B_i(u_i)$. If we let Γ_i be the penalty cost of denying boarding to a reservation over flight leg i , then the second and third assumptions imply that the total expected penalty cost that we incur at the departure time is given by $\sum_{i \in \mathcal{L}} \Gamma_i \mathbb{E}\{[B_i(u_i) - c_i]^+\}$. Therefore, letting w_j be the number of requests for itinerary j that we plan to accept over the planning horizon, and using w and u to respectively denote the vectors $\{w_j : j \in \mathcal{J}\}$ and $\{u_i : i \in \mathcal{L}\}$, VCE solves the problem

$$\begin{aligned} & \max_{(w,u) \in \mathfrak{R}_+^{|\mathcal{J}|+|\mathcal{L}|}} && \sum_{j \in \mathcal{J}} f_j q_j w_j - \sum_{i \in \mathcal{L}} \Gamma_i \mathbb{E}\{[B_i(u_i) - c_i]^+\} \\ & \text{subject to} && \sum_{j \in \mathcal{J}} a_{ij} w_j - u_i = 0 && \text{for all } i \in \mathcal{L} \\ & && w_j \leq \lambda_j && \text{for all } j \in \mathcal{J} \end{aligned}$$

and uses the optimal values of the dual variables associated with the first set of constraints as the bid prices. We note that we use the interpolations of $\mathbb{E}\{[B_i(u_i) - c_i]^+\}$ to compute this function at noninteger values for u_i . VCE is proposed by Karaesmen and van Ryzin (2004a).

Karaesmen and van Ryzin (2004a) suggest several choices for Q_i and Γ_i . For Q_i , we let $Q_i = \sum_{j \in \mathcal{J}} a_{ij} q_j / \sum_{j \in \mathcal{J}} a_{ij}$ so that Q_i is the average of the show up probabilities associated with the itineraries that use flight leg i . For Γ_i , we let $\tilde{\gamma}_j = \gamma_j / \sum_{l \in \mathcal{L}} a_{lj}$ to evenly distribute the penalty cost associated with itinerary j over the different flight legs that it uses, and similar to Q_i , let $\Gamma_i = \sum_{j \in \mathcal{J}} a_{ij} \tilde{\gamma}_j / \sum_{j \in \mathcal{J}} a_{ij}$. By following the same approach that DLP uses, VCE recomputes the bid prices five times over the planning horizon.

5.3 PERFORMANCE COMPARISON

Our main computational results are summarized in Tables 1 and 2. In particular, Tables 1 and 2 respectively show the results for the test problems where the penalty cost satisfies $\gamma \in \{600, 900\}$ and $\gamma \in \{1,200, 1,500\}$. The first column in these tables shows the characteristics of the test problem by using the tuple $(\gamma, \alpha, q^h, q^l)$. The second, third, fourth, fifth and sixth columns respectively show the total expected profits obtained by SAR, SAD, DLP, SCN and VCE. We estimate these total expected profits by simulating the performances of the different benchmark strategies under multiple demand trajectories. We use common random numbers when simulating the performances of the different

benchmark strategies. The seventh column shows the percent gap between the total expected profits obtained by SAD and SAR. This column also includes a “✓” if SAD performs better than SAR, a “×” if SAR performs better than SAD and a “⊙” if there does not exist a statistically significant difference between the total expected profits obtained by SAD and SAR at 5% significance level. The eighth, ninth and tenth columns do the same thing as the seventh column, but they respectively compare the performance of SAD with the performances of DLP, SCN and VCE. SAD turns out to be one of the better benchmark strategies and we use it as a reference point. Finally, the eleventh column shows the upper bound on the optimal total expected profit provided by the optimal objective value of problem (7)-(10). We can compare the total expected profit obtained by one of the benchmark strategies with this upper bound to get a feel for the optimality gap of the benchmark strategy.

The results indicate that both SAR and SAD perform better than DLP, SCN and VCE. The test problems with $\gamma \in \{600, 900\}$ represent the cases where the penalty cost of denying boarding to a reservation is relatively low, especially considering the fact that the penalty cost should include not only the tangible costs related to rescheduling and compensation, but also the intangible costs related to loss of goodwill and prestige. For these test problems, the average performance gap between SAD and DLP is 2.26%. SCN significantly improves on DLP, but the average gap between the total expected profits obtained by SAD and SCN is still 1.53%, which is considered a significant gap in the network revenue management setting. There are test problems where the performance gap between SAD and SCN is as high as 3.12%. The total expected profits obtained by VCE turn out to be comparable to those obtained by DLP. The performances of SAD and SAR are very close to each other. As mentioned above, SAD should be visualized only as a heuristic extension of SAR, but it is welcome news that SAD performs as well as SAR, as a randomized policy may not be appropriate in certain application settings. Similar to SAR and SAD, Topaloglu (2008) uses the bid prices that are obtained through a stochastic approximation algorithm in both randomized and deterministic policies and he reports that the deterministic policy performs at least as well as the randomized policy. The performance gaps become more noticeable when we move to the test problems with larger penalty costs. For the test problems with $\gamma \in \{1,200, 1,500\}$, the average performance gap between SAD and DLP comes out to be 4.89%. SCN provides improvements over DLP, but the average performance gap between SAD and SCN is still 2.00%. VCE can improve on DLP, but this is not the case for every test problem and its performance is not competitive to that of SAD.

From our computational experiments, the penalty cost of denying boarding to a reservation and the tightness of the leg capacities emerge as two factors that affect the performance gaps. Table 3 shows the performance gaps between SAD and the other four benchmark strategies averaged over different groups of test problems. The left portion of this table shows the performance gaps averaged over the test problems with a particular value for γ , whereas the right portion shows the performance gaps averaged over the test problems with a particular value for α . The results indicate that the performance gaps of SAD with DLP, SCN and VCE get larger as the penalty costs get larger and the leg capacities get tighter. For the test problems with large penalty costs and tight leg capacities, the “regret” associated with making an “incorrect” decision tends to be high. For example, if the penalty costs are large, then it is costly to deny boarding to a reservation and we cannot simply make up for an “incorrectly”

accepted itinerary request by denying boarding to the reservation at the departure time. Similarly, if the leg capacities are tight, then the only way for making up for an “incorrectly” accepted itinerary request may be denying boarding to a reservation. Furthermore, it is easy to see that as the penalty costs approach to zero or as the leg capacities approach to infinity, the optimal policy is to accept all itinerary requests. Therefore, the test problems with large penalty costs and tight leg capacities tend to require more careful planning and it is encouraging that SAD improves on DLP, SCN and VCE especially for such test problems.

Since our benchmark strategies are designed to maximize the total expected profit, we use this performance measure when comparing the different benchmark strategies in Tables 1 and 2. However, there are a number of other performance measures besides the total expected profit that can be of practical interest. In Tables 4 and 5, we provide additional performance measures for our benchmark strategies. In particular, Tables 4 and 5 respectively focus on the test problems with $\gamma \in \{600, 900\}$ and $\gamma \in \{1,200, 1,500\}$. The first column in these tables shows the characteristics of the test problem by using the tuple $(\gamma, \alpha, q^h, q^l)$. The second, third and fourth columns respectively show the total expected revenues obtained by SAD, DLP and SCN, whereas the fifth, sixth and seventh columns show the total expected penalty costs. We note that the difference between the total expected revenues and the total expected penalty costs in Tables 4 and 5 yields the total expected profits in Tables 1 and 2. The eighth, ninth and tenth columns respectively show what percents of the reservations that show up are allowed boarding by SAD, DLP and SCN. Finally, the eleventh, twelfth and thirteenth columns respectively show what percents of the seats are occupied when the flight legs depart under the policy used by SAD, DLP and SCN. The performances of SAR and SAD are very close to each other and VCE does not provide too much improvement over DLP. Therefore, we only focus on SAD, DLP and SCN for economy of space in Tables 4 and 5.

Comparing SAD and DLP in Tables 4 and 5, we observe that these two benchmark strategies obtain similar total expected revenues, but the total expected penalty costs incurred by SAD are noticeably smaller than those incurred by DLP. In addition, SAD ends up denying boarding to a smaller percent of reservations than DLP. Under the policy used by SAD, a larger percent of the seats are empty than under the policy used by DLP. These observations imply that SAD accepts fewer itinerary requests than DLP so that it denies boarding to a smaller percent of the reservations. Although SAD accepts fewer itinerary requests than DLP, it obtains comparable total expected revenues and this is due to the fact that SAD carefully reserves the capacity for high fare itinerary requests. Comparing SAD and SCN in Tables 4 and 5, we note that both the total expected revenues and the total expected penalty costs corresponding to SAD are larger than those corresponding to SCN. As shown by the results in Tables 1 and 2, the additional total expected revenues obtained by SAD justify the additional total expected penalty costs so that the total expected profits obtained by SAD end up being noticeably higher than those obtained by SCN.

Although we do not report detailed figures in the paper, our computational experiments show that the variation in the total profit from one sample path to the next is comparable for all of the benchmark strategies. Therefore, if a decision maker is risk averse and tries to minimize the variation in the total

profit, then the different benchmark strategies provide similar levels of comfort and it is still sensible to follow the benchmark strategy with the higher total expected profit. Another performance measure that can be of interest to a risk averse decision maker is the expected regret. In particular, if the numbers of requests for all of the itineraries and the show up decisions of all of the reservations are known in advance, then we can solve a simple integer program to decide which itinerary requests should be accepted. The regret of a benchmark strategy on a sample path is given by the percent gap between the total profit that is obtained by the benchmark strategy on the sample path and the total hindsight revenue that would be obtained when the numbers of requests for all of the itineraries and the show up decisions of all of the reservations on the sample path are known in advance. If we compare our benchmark strategies from the perspective of the expected regret, then we obtain results that are similar to those in Tables 1 and 2. For the test problems with $\gamma \in \{600, 900\}$, the average expected regret for SAD is 10.48%. In other words, the average percent gap between the total profit obtained by SAD and the total hindsight revenue is 10.48%. For DLP and SCN, the average expected regrets are respectively 12.47% and 11.83%. For the test problems with $\gamma \in \{1,200, 1,500\}$, the average expected regrets for SAD, DLP and SCN are respectively 12.82%, 17.00% and 14.54%. In all of our test problems, the expected regrets for SAD are smaller than those for DLP and SCN. Therefore, SAD provides improvements over DLP and SCN in terms of the expected regret as well.

Our results in this section indicate that the bid prices obtained by SAR and SAD perform noticeably better than those obtained by DLP, SCN and VCE. A natural question is whether we can obtain even better policies by searching beyond the class of traditional bid price policies. Erdelyi and Topaloglu (2009a) propose a method to decompose the dynamic programming formulation of the joint capacity allocation and overbooking problem by the flight legs. The policies that they ultimately obtain are more complicated to implement than a traditional bid price policy, as they obtain capacity dependent bid prices, where the bid price associated with a flight leg depends on how much capacity has been committed on the flight leg. It is also worthwhile to note that the capacity dependent bid prices of Erdelyi and Topaloglu (2009a) are obtained under the assumption that the total penalty cost is computed by solving problem (4)-(6). It is not clear whether their approach extends to another way of computing the total penalty cost. In any case, the capacity dependent bid prices contain more information than the traditional bid prices that we use in this paper. As a result, we expect them to perform better than the traditional bid prices. In Table 6, we compare the total expected profits obtained by SAD and the capacity dependent bid prices of Erdelyi and Topaloglu (2009a). We refer to the capacity dependent bid prices as CDB in this table. For the test problems with $\gamma \in \{600, 900\}$, the average performance gap between SAD and CDB is 0.70%. The average performance gap increases to 1.32% when we move to the test problems with $\gamma \in \{1,200, 1,500\}$. Considering the results in Tables 1, 2 and 6, SAD improves on DLP by somewhere between 0.94% and 8.00%. By searching for a good scaling factor, SCN improves on DLP, but the performance gap between SAD and SCN can still be as high as 3.12%. By switching to a more complicated policy structure, CDB improves on SAD by 0.17% to 2.70%. The improvements provided by SAD over DLP and SCN come without changing the structure of the bid price policy. We only need to change the way the bid prices are computed. By using a more complicated policy structure, CDB provides the opportunity for further improvement over SAD.

5.4 SENSITIVITY ANALYSIS

In this section, we focus on a limited set of test problems and provide detailed analyses for SAR and SAD. The results that we present are fairly representative and they capture the general behavior of SAR and SAD for a majority of our test problems.

The performances of SAR and SAD appear to be relatively insensitive to the choice of the initial bid prices in Step 1 of Algorithm 1. In our computational experiments, we choose the initial bid price of each flight leg as the fare associated with the low fare itinerary that connects the origin and destination locations of the flight leg. This is an arbitrary initialization procedure and our only goal in this initialization procedure is to come up with bid prices that are on the same order of magnitude as the fares. However, the performances of SAR and SAD do not change too much when we choose the initial bid prices differently. Figure 3 plots $\mathbb{E}\{V(s(x^k))\}$ as a function of the iteration counter k in Algorithm 1 for three different choices of initial bid prices and for two different test problems. The first choice of initial bid prices corresponds to the fares associated with the low fare itineraries that connect the origin and destination locations of the flight legs. The second choice of initial bid prices corresponds to using the trivial value of zero. The third choice of initial bid prices corresponds to using the bid prices provided by problem (7)-(10). The charts on the top and bottom sides of Figure 3 respectively focus on test problems (900, 1.05, 0.7, 0.7) and (1,500, 1.2, 0.7, 0.7). In Figure 3, the performances of the final bid prices obtained by starting from different initial bid prices are within 0.86% of each other. Since the initial bid prices provided by problem (7)-(10) are not as naive as the other two choices, we gain a slight advantage in the early iterations by starting from the bid prices provided by problem (7)-(10), but this advantage does not last long. Despite the encouraging results, we emphasize that the objective function of the problem $\max_x \mathbb{E}\{V(s(x))\}$ is not necessarily concave and the performance of Algorithm 1 may indeed depend on the choice of initial bid prices. Another observation from Figure 3 is that the performance of Algorithm 1 stabilizes after 500 to 1,000 iterations. However, considering the fact that there are really no good stopping conditions for stochastic approximation algorithms in general, we run Algorithm 1 for 2,500 iterations to make sure that we are on the safe side.

In all of our test problems, SAR and SAD use the function $\theta(h) = \frac{1}{2} - \frac{1}{2} e^{-\zeta[h]^+} + \frac{1}{2} e^{-\zeta[-h]^+}$ with $\zeta = 0.075$ when determining the probability of accepting an itinerary request. It turns out that there are a relatively wide range of possible values for ζ with which SAR and SAD perform well. Figure 4 plots $\mathbb{E}\{V(s(x^k))\}$ as a function of the iteration counter k in Algorithm 1 for five different values for ζ and for two different test problems. The charts on the top and bottom sides of this figure respectively focus on test problems (900, 1.05, 0.7, 0.7) and (1,500, 1.2, 0.7, 0.7). For test problem (900, 1.05, 0.7, 0.7), Figure 4 indicates that the values for ζ in the interval $[0.04, 0.075]$ perform well, but moving too much outside this interval in either direction can yield undesirable performance. For example, Algorithm 1 stalls with a total expected profit of about 133,000 when we use $\zeta = 0.3$ and this total expected profit is not as large as the total expected profit obtained by DLP, which is 137,746. We emphasize that Algorithm 1 converges w.p.1 for any finite value of ζ , but the issue is that the limiting bid prices may or may not yield desirable total expected profits if one does not pay attention to the choice of ζ . For test problem (1,500, 1.2, 0.7, 0.7), the values for ζ in the interval $[0.04, 0.15]$ perform well. The results indicate that

the performance of Algorithm 1 is relatively robust to the choice of ζ . Completely arbitrary choices of ζ can admittedly result in undesirable performance and we need to spend some effort to calibrate this parameter. The good news is that once we make a choice for ζ under a particular fare structure, the same value for ζ appears to perform well for a variety of penalty costs, leg capacities and show up probabilities. For example, setting $\zeta = 0.075$ works well for all of our test problems.

The results in Tables 1 and 2 indicate that the bid prices provided by Algorithm 1 tend to perform significantly better than those provided by problem (7)-(10). An interesting question is whether the bid prices provided by Algorithm 1 and problem (7)-(10) are indeed substantially different or they can be considered as minor adjustments to each other. Using $\{x_i^1 : i \in \mathcal{L}\}$ and $\{x_i^2 : i \in \mathcal{L}\}$ to respectively denote the bid prices provided by Algorithm 1 and problem (7)-(10), Figure 5 gives scatter plots of $\{(x_i^1, x_i^2) : i \in \mathcal{L}\}$ for two test problems. The charts on the left and right sides of Figure 5 respectively focus on test problems (900, 1.05, 0.7, 0.7) and (1,500, 1.2, 0.7, 0.7). The figure indicates that the bid prices provided by Algorithm 1 and problem (7)-(10) share similar trends, but there are some flight legs where the bid prices have significant deviations. The differences in the bid prices for these flight legs ultimately translate into significant performance gaps between Algorithm 1 and problem (7)-(10).

A useful advantage of SAR and SAD is that their performances do not depend too much on the number of times that we recompute the bid prices over the planning horizon. We recompute the bid prices five times over the planning horizon in all of our computational experiments. Figure 6 shows the total expected profits obtained by SAR, SAD, DLP, SCN and VCE for two test problems when we recompute the bid prices three, five and ten times over the planning horizon. The charts on the left and right sides respectively focus on test problems (900, 1.05, 0.7, 0.7) and (1,500, 1.2, 0.7, 0.7). The figure indicates that the performances of DLP and VCE can deteriorate when we recompute the bid prices fewer times, but the performances of SAR and SAD remain relatively stable. Furthermore, the performances of SAR and SAD with even three recomputations can be better than the performances of DLP, SCN and VCE with ten recomputations.

For different numbers of spokes in the airline network, Table 7 shows the CPU seconds required to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$ by using Algorithm 1. The CPU seconds in Table 7 correspond to the case where we terminate Algorithm 1 after 2,500 iterations. All of our computational experiments are carried out on an Intel Xeon 2 GHz CPU running Windows XP with 4 GB RAM. The first column in Table 7 shows the number of spokes in the airline network. The second column shows the total CPU seconds required to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$ by using Algorithm 1. The third column shows what portion of the CPU seconds is spent on computing $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$. By the discussion at the end of Section 3, we need to compute $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$ to be able to compute $\Psi_i(x^k, s^k(x^k))$ for all $i \in \mathcal{L}$. The fourth column shows what portion of the CPU seconds is spent on all other operations. We observe that a large portion of the CPU seconds is spent on computing $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$. This is due to the fact that computing $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$ requires solving $1 + |\mathcal{J}|$ optimization problems of the form (4)-(6) and the numbers of decision variables and constraints in problem (4)-(6) quickly increase with the number of spokes. We emphasize that our approach does not necessarily

require computing the total penalty cost by solving problem (4)-(6) and if we did not compute the total penalty cost by solving problem (4)-(6), then the portion of the CPU seconds spent on computing $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$ could be smaller. The portion of the CPU seconds spent on all other operations is quite small. Generally speaking, Algorithm 1 takes on the order of minutes to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$. Since the bid prices are computed in an offline manner, these CPU seconds are reasonable for practical implementation. The CPU seconds required by DLP, SCN and VCE to compute one set of bid prices are on the order of a fraction of a second. Nevertheless, the extra computational burden of Algorithm 1 is justified to a large extent by the significant performance improvements provided by SAR and SAD over the other three benchmark strategies. Furthermore, Figure 6 indicates that SAR and SAD can obtain better total expected profits than DLP, SCN and VCE by recomputing the bid prices fewer times. Due to the administrative burden associated with opening and closing the itineraries, a solution method that requires recomputing the bid prices fewer times but takes longer to run may be preferable to a solution method that requires recomputing the bid prices more frequently but takes shorter to run.

6 CONCLUSIONS

In this paper, we developed a stochastic approximation algorithm to make the capacity allocation and overbooking decisions over an airline network. Our algorithm searches for a good set of bid prices by using the stochastic gradients of the total expected profit function. We derived a simple expression that can be used to compute the stochastic gradient of the total expected profit function and showed that the iterates of our stochastic approximation algorithm converge to a stationary point of the total expected profit function w.p.1. Our computational experiments indicate that the bid prices computed by our approach perform significantly better than numerous benchmark strategies.

There are two practically useful features of the stochastic approximation algorithm that we develop in this paper. First, our stochastic approximation algorithm does not make any assumptions on how we compute the total penalty cost of denying boarding to the reservations. This is in contrast with the deterministic linear program, which strictly works under the assumption that we can oversee the whole airline network and jointly decide which of the reservations should be denied boarding to minimize the total penalty cost. Second, the bid prices computed by our stochastic approximation algorithm represent the traditional view of bid price policies and they do not depend on the committed capacities on the flight legs. This view of bid prices is in alignment with the current airline operations. Being a stochastic approximation algorithm, our approach potentially shares the shortcomings commonly associated with that class of algorithms. In particular, we need to make choices on the step size parameter and the stopping criterion in Algorithm 1 and there are no easy rules for making these choices. Similarly, we need to make a choice on the function $\theta(\cdot)$ that we use to determine the probability of accepting an itinerary request and we may need to do some experimentation to come up with a choice that works well. Although the choices for the step size parameter, stopping criterion and function $\theta(\cdot)$ require some setup runs, our experience has been that the same step size parameter, stopping criterion and function $\theta(\cdot)$ work well for a variety of penalty costs, leg capacities and show up probabilities. Therefore, it empirically appears that we do not have to make these choices completely from scratch for each new

data set. Finally, the randomization used by our bid price policies may be a concern when implementing these policies in practice, but SAD in our computational experiments indicates that we still obtain good performance by using the bid prices computed by our stochastic approximation algorithm in a deterministic decision rule.

The performance measure that drives our approach and the benchmark strategies that we use is the total expected profit. However, many airlines face the challenge of maximizing the total expected revenue while providing high service levels. One possible method to address this challenge is to vary the penalty cost of denying boarding to a reservation and observe the service level implication of different penalty costs. This method essentially computes the imputed penalty cost related to a certain service level and it can easily be implemented in conjunction with our stochastic approximation algorithm. On the other hand, a more comprehensive method to balance the total expected revenue with the service levels requires solving an optimization problem that maximizes the total expected revenue subject to constraints on the service levels. As a result, we need to deal with an optimization problem where both the objective function and the constraints involve random elements and such an optimization problem is clearly more difficult than the one that we consider in this paper.

There are several directions for future research and extensions. As mentioned above, one of the advantages of our approach is that it does not depend on the particular form of the total penalty cost. Nevertheless, since our stochastic approximation algorithm requires computing $p(s^k(x^k))$ and $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$ at each iteration, a significant portion of the computational requirements is due to computing the total penalty cost. One area for future research is to identify special cases for the total penalty cost where it becomes computationally easy to obtain $p(s^k(x^k) + \epsilon_j)$ for all $j \in \mathcal{J}$ after computing $p(s^k(x^k))$. This would significantly reduce the computational requirements. For an area of extension, we note that our construction in Section 3 is largely based on the observation that the sum of two independent Poisson random variables with means λ_1 and λ_2 is a Poisson random variable with mean $\lambda_1 + \lambda_2$. There are other random variables that satisfy similar properties. For example, the sum of two independent binomial random variables with parameters (n_1, p) and (n_2, p) is a binomial random variable with parameter $(n_1 + n_2, p)$. More generally, Karaesmen and van Ryzin (2004b) say that a random variable $Z(\rho)$ characterized by a vector of parameters ρ satisfies the semigroup property whenever the following condition holds. If Y_1 and Y_2 are independent random variables with Y_1 having the same distribution as $Z(\rho_1)$ and Y_2 having the same distribution as $Z(\rho_2)$, then $Y_1 + Y_2$ has the same distribution as $Z(\rho_1 + \rho_2)$. It is clear that Poisson and binomial random variables satisfy the semigroup property. Our analysis in this paper can indeed be generalized to cover the case where the numbers of reservations that show up at the departure time satisfy the semigroup property. We do not pursue this generalization in detail since we believe that the Poisson distribution already provides an appropriate and practically useful model of demand and the steps involved in the generalization essentially follow the outline of Section 3 line by line.

REFERENCES

Ball, M. O. and Queyranne, M. (2009), ‘Toward robust revenue management: Competitive analysis of online booking’, *Operations Research* **57**(4), 950–963.

- Bashyam, S. and Fu, M. C. (1998), ‘Optimization 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.
- Bertsimas, D. and Popescu, I. (2003), ‘Revenue management in a dynamic network environment’, *Transportation Science* **37**, 257–277.
- 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.
- Erdelyi, A. and Topaloglu, H. (2009a), ‘A dynamic programming decomposition method for making overbooking decisions over an airline network’, *INFORMS Journal on Computing* (to appear).
- Erdelyi, A. and Topaloglu, H. (2009b), ‘Separable approximations for joint capacity control and overbooking decisions in network revenue management’, *Journal of Revenue and Pricing Management* **8**(1), 3–20.
- Fu, M. (1994), ‘Sample path derivatives for (s, S) inventory systems’, *Operations Research* **42**(2), 351–363.
- 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. (2004a), Coordinating overbooking and capacity control decisions on a network, Technical report, Columbia Business School.
- Karaesmen, I. and van Ryzin, G. (2004b), ‘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. (2008a), ‘A tractable revenue management model for capacity allocation and overbooking over an airline network’, *Flexible Services and Manufacturing Journal* **20**(3), 125–147.
- Kunnumkal, S. and Topaloglu, H. (2008b), ‘Using stochastic approximation algorithms to compute optimal base-stock levels in inventory control problems’, *Operations Research* **56**(3), 646–664.
- Kushner, H. J. and Clark, D. S. (1978), *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, Berlin.
- Lan, Y., Ball, M. and Karaesmen, I. Z. (2007), Overbooking and fare-class allocation with limited information, Technical report, University of Maryland, Robert H. Smith School of Business.
- L’Ecuyer, P. and Glynn, P. (1994), ‘Stochastic optimization by simulation: Convergence proofs for the GI/G/1 queue in steady state’, *Management Science* **40**, 1245–1261.

- Mahajan, S. and van Ryzin, G. (2001), ‘Stocking retail assortments under dynamic customer substitution’, *Operations Research* **49**(3), 334–351.
- Subramanian, J., Stidham, S. and Lautenbacher, C. J. (1999), ‘Airline yield management with overbooking, cancellations and no-shows’, *Transportation Science* **33**(2), 147–167.
- Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer, New York, NY.
- Talluri, K. and van Ryzin, G. (1998), ‘An analysis of bid-price controls for network revenue management’, *Management Science* **44**(11), 1577–1593.
- 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 Vulcano, G. (2008a), ‘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. (2008b), ‘Simulation-based optimization of virtual nesting controls for network revenue management’, *Operations Research* **56**(4), 865–880.
- Williamson, E. L. (1992), *Airline Network Seat Control*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.




Figure 1: Structure of the airline network for the case where $N = 8$.




Figure 2: Our choice of $\theta(\cdot)$ for different values of ζ .

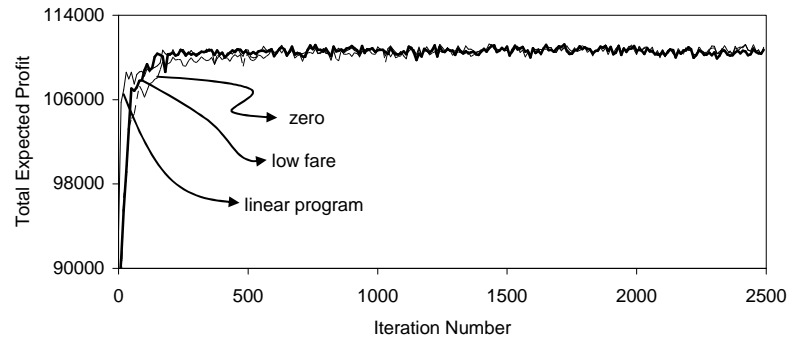
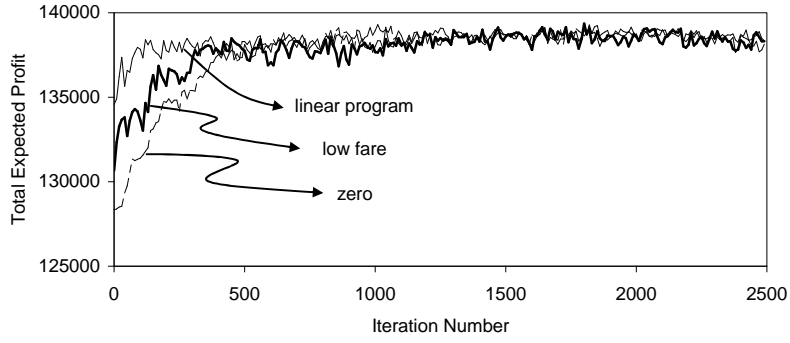


Figure 3: Performance of Algorithm 1 with different initial bid prices.

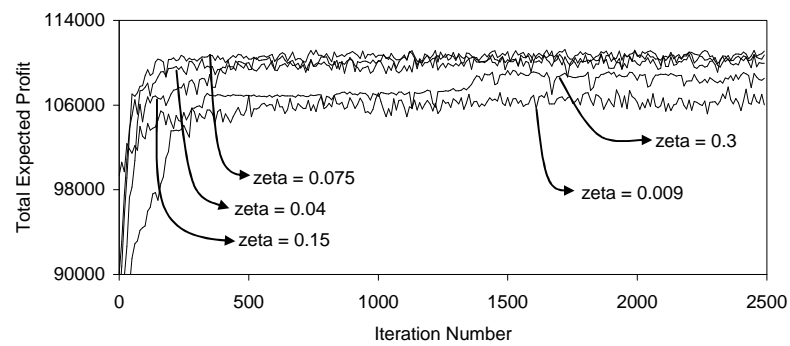



Figure 4: Performance of Algorithm 1 with different choices of ζ .




Figure 5: Bid prices provided by Algorithm 1 and problem (7)-(10).




Figure 6: Performances of SAR, SAD, DLP, SCN and VCE as a function of the number of times that we recompute the bid prices.

Test Problem				Tot. Exp. Profit					Percent Gap with SAD				DLP Bound
γ	α	q^h	q^l	SAR	SAD	DLP	SCN	VCE	SAR	DLP	SCN	VCE	
600	1.05	0.7	0.7	151,992	151,897	150,400	150,525	150,328	-0.06 \odot	0.99 \checkmark	0.90 \checkmark	1.03 \checkmark	164,132
600	1.05	0.7	0.9	149,563	149,531	147,840	147,840	148,546	-0.02 \odot	1.13 \checkmark	1.13 \checkmark	0.66 \checkmark	163,211
600	1.05	0.9	0.7	151,109	150,914	149,162	149,162	148,570	-0.13 \times	1.16 \checkmark	1.16 \checkmark	1.55 \checkmark	163,439
600	1.05	0.9	0.9	115,123	115,125	113,534	113,534	113,849	0.00 \odot	1.38 \checkmark	1.38 \checkmark	1.11 \checkmark	126,272
600	1.2	0.7	0.7	126,171	126,117	123,943	124,715	123,655	-0.04 \odot	1.72 \checkmark	1.11 \checkmark	1.95 \checkmark	139,668
600	1.2	0.7	0.9	132,616	133,085	131,832	131,832	132,499	0.35 \checkmark	0.94 \checkmark	0.94 \checkmark	0.44 \checkmark	148,052
600	1.2	0.9	0.7	121,329	121,242	119,132	119,790	118,858	-0.07 \odot	1.74 \checkmark	1.20 \checkmark	1.97 \checkmark	132,561
600	1.2	0.9	0.9	102,358	102,589	100,140	100,439	100,690	0.23 \odot	2.39 \checkmark	2.10 \checkmark	1.85 \checkmark	113,869
600	1.6	0.7	0.7	144,617	144,662	142,946	142,946	141,795	0.03 \odot	1.19 \checkmark	1.19 \checkmark	1.98 \checkmark	158,090
600	1.6	0.7	0.9	108,112	108,230	106,910	106,910	106,118	0.11 \odot	1.22 \checkmark	1.22 \checkmark	1.95 \checkmark	122,328
600	1.6	0.9	0.7	107,917	107,928	104,612	105,547	105,374	0.01 \odot	3.07 \checkmark	2.21 \checkmark	2.37 \checkmark	120,569
600	1.6	0.9	0.9	109,517	109,640	106,894	107,069	107,605	0.11 \odot	2.51 \checkmark	2.34 \checkmark	1.86 \checkmark	122,863
900	1.05	0.7	0.7	140,500	140,417	137,746	138,955	137,464	-0.06 \odot	1.90 \checkmark	1.04 \checkmark	2.10 \checkmark	155,966
900	1.05	0.7	0.9	145,382	145,344	141,910	143,122	144,027	-0.03 \odot	2.36 \checkmark	1.53 \checkmark	0.91 \checkmark	161,201
900	1.05	0.9	0.7	145,693	145,733	142,808	143,983	142,103	0.03 \odot	2.01 \checkmark	1.20 \checkmark	2.49 \checkmark	159,529
900	1.05	0.9	0.9	126,352	126,453	123,644	124,310	124,729	0.08 \odot	2.22 \checkmark	1.69 \checkmark	1.36 \checkmark	138,577
900	1.2	0.7	0.7	111,413	111,247	108,084	110,360	106,720	-0.15 \odot	2.84 \checkmark	0.80 \checkmark	4.07 \checkmark	127,143
900	1.2	0.7	0.9	120,492	120,705	117,652	119,485	118,310	0.18 \odot	2.53 \checkmark	1.01 \checkmark	1.98 \checkmark	135,794
900	1.2	0.9	0.7	126,271	126,408	121,962	123,977	120,789	0.11 \odot	3.52 \checkmark	1.92 \checkmark	4.44 \checkmark	140,828
900	1.2	0.9	0.9	129,961	130,394	126,189	128,715	127,125	0.33 \odot	3.22 \checkmark	1.29 \checkmark	2.51 \checkmark	142,330
900	1.6	0.7	0.7	95,344	95,350	90,613	92,372	89,156	0.01 \odot	4.97 \checkmark	3.12 \checkmark	6.50 \checkmark	109,658
900	1.6	0.7	0.9	96,577	96,342	94,644	94,644	92,240	-0.24 \odot	1.76 \checkmark	1.76 \checkmark	4.26 \checkmark	113,489
900	1.6	0.9	0.7	101,858	101,918	96,801	98,895	97,098	0.06 \odot	5.02 \checkmark	2.97 \checkmark	4.73 \checkmark	114,771
900	1.6	0.9	0.9	115,687	115,378	112,518	113,747	112,410	-0.27 \odot	2.48 \checkmark	1.41 \checkmark	2.57 \checkmark	131,842
Average									0.02	2.26	1.53	2.36	

Table 1: Total expected profits for the test problems with $\gamma = 600$ and $\gamma = 900$.

Test Problem				Tot. Exp. Profit					Percent Gap with SAD				DLP Bound
γ	α	q^h	q^l	SAR	SAD	DLP	SCN	VCE	SAR	DLP	SCN	VCE	
1,200	1.05	0.7	0.7	126,083	126,037	122,149	124,410	122,470	-0.04 \odot	3.08 \checkmark	1.29 \checkmark	2.83 \checkmark	140,869
1,200	1.05	0.7	0.9	143,037	143,208	139,234	140,627	141,307	0.12 \odot	2.78 \checkmark	1.80 \checkmark	1.33 \checkmark	161,248
1,200	1.05	0.9	0.7	125,320	125,585	120,400	122,822	118,230	0.21 \odot	4.13 \checkmark	2.20 \checkmark	5.86 \checkmark	139,107
1,200	1.05	0.9	0.9	122,729	122,509	118,138	120,724	119,036	-0.18 \odot	3.57 \checkmark	1.46 \checkmark	2.83 \checkmark	136,564
1,200	1.2	0.7	0.7	119,686	119,740	114,855	116,713	113,575	0.04 \odot	4.08 \checkmark	2.53 \checkmark	5.15 \checkmark	136,233
1,200	1.2	0.7	0.9	115,300	115,081	110,242	113,447	110,748	-0.19 \odot	4.20 \checkmark	1.42 \checkmark	3.76 \checkmark	132,127
1,200	1.2	0.9	0.7	145,979	146,249	141,210	143,423	138,957	0.18 \odot	3.45 \checkmark	1.93 \checkmark	4.99 \checkmark	162,005
1,200	1.2	0.9	0.9	131,648	132,006	126,282	129,287	127,723	0.27 \odot	4.34 \checkmark	2.06 \checkmark	3.24 \checkmark	146,062
1,200	1.6	0.7	0.7	106,682	106,995	103,165	105,033	99,957	0.29 \odot	3.58 \checkmark	1.83 \checkmark	6.58 \checkmark	126,501
1,200	1.6	0.7	0.9	85,703	86,012	81,254	83,616	81,101	0.36 \odot	5.53 \checkmark	2.79 \checkmark	5.71 \checkmark	102,076
1,200	1.6	0.9	0.7	115,137	115,214	109,766	112,311	111,178	0.07 \odot	4.73 \checkmark	2.52 \checkmark	3.50 \checkmark	132,560
1,200	1.6	0.9	0.9	104,536	104,410	99,406	101,635	100,232	-0.12 \odot	4.79 \checkmark	2.66 \checkmark	4.00 \checkmark	121,108
1,500	1.05	0.7	0.7	113,453	113,166	107,288	111,786	108,732	-0.25 \odot	5.19 \checkmark	1.22 \checkmark	3.92 \checkmark	129,070
1,500	1.05	0.7	0.9	142,813	143,151	137,510	141,143	140,625	0.24 \odot	3.94 \checkmark	1.40 \checkmark	1.76 \checkmark	161,469
1,500	1.05	0.9	0.7	119,186	119,244	114,074	117,385	110,934	0.05 \odot	4.34 \checkmark	1.56 \checkmark	6.97 \checkmark	133,143
1,500	1.05	0.9	0.9	112,931	112,691	106,903	110,170	108,526	-0.21 \odot	5.14 \checkmark	2.24 \checkmark	3.70 \checkmark	127,004
1,500	1.2	0.7	0.7	112,507	112,478	105,277	110,044	105,450	-0.03 \odot	6.40 \checkmark	2.16 \checkmark	6.25 \checkmark	128,141
1,500	1.2	0.7	0.9	107,517	107,670	101,841	106,105	104,425	0.14 \odot	5.41 \checkmark	1.45 \checkmark	3.01 \checkmark	125,088
1,500	1.2	0.9	0.7	111,709	111,272	103,855	109,106	103,054	-0.39 \odot	6.67 \checkmark	1.95 \checkmark	7.39 \checkmark	125,111
1,500	1.2	0.9	0.9	124,612	124,899	118,475	122,491	120,818	0.23 \odot	5.14 \checkmark	1.93 \checkmark	3.27 \checkmark	140,457
1,500	1.6	0.7	0.7	98,355	98,029	91,017	95,147	89,741	-0.33 \odot	7.15 \checkmark	2.94 \checkmark	8.45 \checkmark	115,376
1,500	1.6	0.7	0.9	97,342	97,493	91,563	95,285	92,962	0.16 \odot	6.08 \checkmark	2.27 \checkmark	4.65 \checkmark	115,574
1,500	1.6	0.9	0.7	114,800	114,594	105,431	111,570	108,237	-0.18 \odot	8.00 \checkmark	2.64 \checkmark	5.55 \checkmark	129,996
1,500	1.6	0.9	0.9	114,023	113,927	107,506	111,792	109,552	-0.08 \odot	5.64 \checkmark	1.87 \checkmark	3.84 \checkmark	132,531
Average									0.01	4.89	2.00	4.52	

Table 2: Total expected profits for the test problems with $\gamma = 1,200$ and $\gamma = 1,500$.

Pen. Cost	Avg. Perc. Gap with SAD			
	SAR	DLP	SCN	VCE
600	0.04	1.62	1.41	1.56
900	0.00	2.90	1.65	3.16
1,200	0.09	4.02	2.04	4.15
1,500	-0.06	5.76	1.97	4.90

Leg Tight.	Avg. Perc. Gap with SAD			
	SAR	DLP	SCN	VCE
1.05	-0.02	2.83	1.45	2.53
1.2	0.07	3.66	1.61	3.52
1.6	0.00	4.23	2.23	4.28

Table 3: Percent gaps between the total expected profits obtained by SAD and the other benchmark strategies averaged over different groups of test problems.

Test Problem				Tot. Exp. Revenue			Tot. Exp. Cost			Exp. Service Level			Exp. Occupancy		
γ	α	q^h	q^l	SAD	DLP	SCN	SAD	DLP	SCN	SAD	DLP	SCN	SAD	DLP	SCN
600	1.05	0.7	0.7	157,849	157,372	153,609	5,952	6,972	3,084	95.56	94.97	97.60	89.00	91.27	86.72
600	1.05	0.7	0.9	155,399	154,812	154,812	5,868	6,972	6,972	95.98	95.34	95.34	88.65	90.98	90.98
600	1.05	0.9	0.7	154,622	154,826	154,826	3,708	5,664	5,664	97.65	96.59	96.59	87.61	91.63	91.63
600	1.05	0.9	0.9	118,377	118,166	118,166	3,252	4,632	4,632	98.07	97.36	97.36	87.49	91.07	91.07
600	1.2	0.7	0.7	131,829	131,935	128,327	5,712	7,992	3,612	95.28	93.67	96.91	88.51	91.95	86.94
600	1.2	0.7	0.9	139,841	138,744	138,744	6,756	6,912	6,912	94.84	94.83	94.83	88.25	90.83	90.83
600	1.2	0.9	0.7	125,334	125,384	122,034	4,092	6,252	2,244	97.12	95.82	98.37	88.70	92.81	87.27
600	1.2	0.9	0.9	106,345	105,912	102,083	3,756	5,772	1,644	97.49	96.36	98.86	88.35	92.98	86.45
600	1.6	0.7	0.7	153,074	150,578	150,578	8,412	7,632	7,632	91.06	92.01	92.01	89.86	91.40	91.40
600	1.6	0.7	0.9	114,794	113,510	113,510	6,564	6,600	6,600	93.56	93.72	93.72	88.57	91.46	91.46
600	1.6	0.9	0.7	112,572	111,692	108,475	4,644	7,080	2,928	95.70	93.89	97.22	88.01	93.22	87.09
600	1.6	0.9	0.9	114,896	113,254	109,697	5,256	6,360	2,628	95.48	94.76	97.62	89.43	93.04	86.80
900	1.05	0.7	0.7	146,789	146,944	142,753	6,372	9,198	3,798	96.76	95.56	98.03	85.31	89.81	84.58
900	1.05	0.7	0.9	152,058	151,612	146,686	6,714	9,702	3,564	96.80	95.62	98.26	86.97	90.84	85.54
900	1.05	0.9	0.7	150,143	150,386	146,053	4,410	7,578	2,070	98.10	96.94	99.09	85.93	91.06	85.52
900	1.05	0.9	0.9	130,035	130,394	125,426	3,582	6,750	1,116	98.56	97.45	99.54	85.85	91.24	84.96
900	1.2	0.7	0.7	117,421	117,966	114,320	6,174	9,882	3,960	96.37	94.68	97.70	85.71	91.89	86.14
900	1.2	0.7	0.9	126,915	127,552	123,265	6,210	9,900	3,780	96.60	94.97	97.90	85.12	90.94	85.04
900	1.2	0.9	0.7	131,358	131,286	127,145	4,950	9,324	3,168	97.64	95.87	98.48	86.25	92.53	86.20
900	1.2	0.9	0.9	135,146	133,623	130,551	4,752	7,434	1,836	97.86	96.85	99.16	85.91	92.43	86.13
900	1.6	0.7	0.7	100,930	100,279	96,908	5,580	9,666	4,536	95.66	93.21	96.48	83.65	90.55	84.16
900	1.6	0.7	0.9	103,938	103,158	103,158	7,596	8,514	8,514	94.79	94.39	94.39	87.04	91.06	91.06
900	1.6	0.9	0.7	106,454	105,333	101,919	4,536	8,532	3,024	97.13	95.06	98.09	84.96	92.61	86.65
900	1.6	0.9	0.9	121,714	121,338	117,419	6,336	8,820	3,672	96.13	95.10	97.77	85.37	92.86	86.39
Average										96.26	95.21	97.14	87.10	91.69	87.54

Table 4: Detailed performance measures for the test problems with $\gamma = 600$ and $\gamma = 900$.

Test Problem				Tot. Exp. Revenue			Tot. Exp. Cost			Exp. Service Level			Exp. Occupancy		
γ	α	q^h	q^l	SAD	DLP	SCN	SAD	DLP	SCN	SAD	DLP	SCN	SAD	DLP	SCN
1,200	1.05	0.7	0.7	131,341	132,421	128,418	5,304	10,272	4,008	97.87	96.24	98.42	83.04	90.09	84.55
1,200	1.05	0.7	0.9	148,872	149,794	144,251	5,664	10,560	3,624	97.94	96.41	98.68	84.89	90.33	85.15
1,200	1.05	0.9	0.7	129,833	129,928	125,342	4,248	9,528	2,520	98.57	97.09	99.16	84.92	91.19	85.10
1,200	1.05	0.9	0.9	126,493	127,282	122,500	3,984	9,144	1,776	98.76	97.39	99.45	84.34	91.67	85.69
1,200	1.2	0.7	0.7	126,652	126,903	122,257	6,912	12,048	5,544	96.98	95.10	97.58	84.08	90.64	84.86
1,200	1.2	0.7	0.9	121,489	122,818	118,535	6,408	12,576	5,088	97.35	95.20	97.88	83.91	91.06	85.17
1,200	1.2	0.9	0.7	150,929	150,690	146,207	4,680	9,480	2,784	98.28	96.77	98.97	85.07	92.35	86.23
1,200	1.2	0.9	0.9	136,446	136,170	131,591	4,440	9,888	2,304	98.47	96.88	99.22	84.28	92.05	86.06
1,200	1.6	0.7	0.7	115,227	115,717	111,081	8,232	12,552	6,048	95.39	93.52	96.58	82.82	90.19	84.12
1,200	1.6	0.7	0.9	91,964	92,534	89,208	5,952	11,280	5,592	96.80	94.52	97.03	82.68	90.22	84.40
1,200	1.6	0.9	0.7	120,902	119,894	115,935	5,688	10,128	3,624	97.23	95.58	98.27	82.64	92.45	85.43
1,200	1.6	0.9	0.9	109,906	109,174	105,235	5,496	9,768	3,600	97.49	95.95	98.39	82.73	92.20	85.70
1,500	1.05	0.7	0.7	118,686	120,398	115,926	5,520	13,110	4,140	98.21	96.17	98.69	82.21	90.20	84.49
1,500	1.05	0.7	0.9	149,931	151,430	146,273	6,780	13,920	5,130	97.94	96.20	98.48	82.43	89.98	84.32
1,500	1.05	0.9	0.7	123,174	124,274	119,875	3,930	10,200	2,490	98.94	97.51	99.34	82.71	91.40	85.52
1,500	1.05	0.9	0.9	116,501	117,553	113,110	3,810	10,650	2,940	99.03	97.57	99.28	82.96	91.14	85.04
1,500	1.2	0.7	0.7	118,568	119,647	116,074	6,090	14,370	6,030	97.72	95.32	97.85	81.40	91.43	85.52
1,500	1.2	0.7	0.9	113,760	115,671	110,815	6,090	13,830	4,710	97.93	95.80	98.44	82.03	91.01	84.50
1,500	1.2	0.9	0.7	115,862	116,095	112,316	4,590	12,240	3,210	98.59	96.65	99.05	82.41	92.32	85.88
1,500	1.2	0.9	0.9	129,429	129,905	125,101	4,530	11,430	2,610	98.67	97.01	99.25	83.24	92.26	85.36
1,500	1.6	0.7	0.7	104,959	105,627	101,867	6,930	14,610	6,720	96.65	93.78	96.86	80.68	90.62	84.31
1,500	1.6	0.7	0.9	103,313	104,943	100,235	5,820	13,380	4,950	97.34	94.62	97.81	81.77	90.50	83.81
1,500	1.6	0.9	0.7	120,204	119,531	115,920	5,610	14,100	4,350	97.75	95.02	98.31	81.16	92.22	85.83
1,500	1.6	0.9	0.9	119,747	119,956	116,262	5,820	12,450	4,470	97.75	95.82	98.35	80.66	92.14	85.56
Average										97.82	95.92	98.39	82.88	91.24	85.11

Table 5: Detailed performance measures for the test problems with $\gamma = 1,200$ and $\gamma = 1,500$.

Test Problem				Tot. Exp. Profit		Perc.
γ	α	q^h	q^l	SAD	CDB	Gap
600	1.05	0.7	0.7	151,897	152,345	-0.30 \odot
600	1.05	0.7	0.9	149,531	149,951	-0.28 \odot
600	1.05	0.9	0.7	150,914	151,173	-0.17 \odot
600	1.05	0.9	0.9	115,125	115,810	-0.59 \times
600	1.2	0.7	0.7	126,117	126,627	-0.40 \odot
600	1.2	0.7	0.9	133,085	133,627	-0.41 \odot
600	1.2	0.9	0.7	121,242	121,609	-0.30 \odot
600	1.2	0.9	0.9	102,589	103,325	-0.72 \times
600	1.6	0.7	0.7	144,662	144,902	-0.17 \odot
600	1.6	0.7	0.9	108,230	109,042	-0.75 \times
600	1.6	0.9	0.7	107,928	108,422	-0.46 \odot
600	1.6	0.9	0.9	109,640	110,570	-0.85 \times
900	1.05	0.7	0.7	140,417	141,132	-0.51 \odot
900	1.05	0.7	0.9	145,344	146,538	-0.82 \times
900	1.05	0.9	0.7	145,733	146,791	-0.73 \times
900	1.05	0.9	0.9	126,453	127,386	-0.74 \times
900	1.2	0.7	0.7	111,247	112,398	-1.03 \times
900	1.2	0.7	0.9	120,705	122,023	-1.09 \times
900	1.2	0.9	0.7	126,408	127,578	-0.93 \times
900	1.2	0.9	0.9	130,394	131,874	-1.14 \times
900	1.6	0.7	0.7	95,350	95,310	0.04 \odot
900	1.6	0.7	0.9	96,342	98,392	-2.13 \times
900	1.6	0.9	0.7	101,918	102,224	-0.30 \odot
900	1.6	0.9	0.9	115,378	117,827	-2.12 \times
Average						-0.70

Test Problem				Tot. Exp. Profit		Perc.
γ	α	q^h	q^l	SAD	CDB	Gap
1,200	1.05	0.7	0.7	126,037	127,051	-0.80 \times
1,200	1.05	0.7	0.9	143,208	144,533	-0.92 \times
1,200	1.05	0.9	0.7	125,585	126,440	-0.68 \times
1,200	1.05	0.9	0.9	122,509	124,264	-1.43 \times
1,200	1.2	0.7	0.7	119,740	120,459	-0.60 \odot
1,200	1.2	0.7	0.9	115,081	116,774	-1.47 \times
1,200	1.2	0.9	0.7	146,249	147,130	-0.60 \times
1,200	1.2	0.9	0.9	132,006	133,902	-1.44 \times
1,200	1.6	0.7	0.7	106,995	108,294	-1.21 \times
1,200	1.6	0.7	0.9	86,012	87,312	-1.51 \times
1,200	1.6	0.9	0.7	115,214	116,356	-0.99 \times
1,200	1.6	0.9	0.9	104,410	106,915	-2.40 \times
1,500	1.05	0.7	0.7	113,166	114,662	-1.32 \times
1,500	1.05	0.7	0.9	143,151	144,736	-1.11 \times
1,500	1.05	0.9	0.7	119,244	120,577	-1.12 \times
1,500	1.05	0.9	0.9	112,691	114,882	-1.94 \times
1,500	1.2	0.7	0.7	112,478	114,063	-1.41 \times
1,500	1.2	0.7	0.9	107,670	109,367	-1.58 \times
1,500	1.2	0.9	0.7	111,272	112,477	-1.08 \times
1,500	1.2	0.9	0.9	124,899	126,707	-1.45 \times
1,500	1.6	0.7	0.7	98,029	99,539	-1.54 \times
1,500	1.6	0.7	0.9	97,493	99,171	-1.72 \times
1,500	1.6	0.9	0.7	114,594	115,390	-0.69 \odot
1,500	1.6	0.9	0.9	113,927	117,004	-2.70 \times
Average						-1.32

Table 6: Total expected profits obtained by SDA and CDB.

No. of Spokes	Total CPU Secs.	CPU Secs. on Problem (4)-(6)	CPU Secs. on Other Operations
2	1.29	0.68	0.61
4	5.59	4.41	1.17
8	46.76	44.09	2.67
12	188.37	181.29	7.08
16	557.03	540.45	16.58

Table 7: CPU seconds required to solve the problem $\max_x \mathbb{E}\{V(s(x))\}$ by using Algorithm 1.

A APPENDIX: PROOF OF PROPOSITION 4

In this section, we consider the limit in (2) as h approaches to zero from left and show that this limit is given by the expression on the right side of (2). The proof proceeds in three steps. In the first step, we use a construction that is similar to the one at the beginning of Section 3 to show that $s_j(x)$ can be visualized as the sum of two independent Poisson random variables. The mean of the Poisson random variable $s_j(x)$ is $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$. We fix $x \in \mathfrak{R}^{|\mathcal{L}|}$ and $i \in \mathcal{L}$, and for $h \geq 0$, the Taylor series expansion of $\theta_j(\cdot)$ at the point $f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h$ yields

$$\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) = \theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) + a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h).$$

Noting the fact that $f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k \geq f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h$ and $\theta(\cdot)$ is increasing, we obtain $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h) \geq 0$. Therefore, the mean of $s_j(x)$ can be written as the sum of the two nonnegative terms $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h)$ and $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h)$. The first term $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h)$ is the mean of the Poisson random variable $s_j(x + h e_i)$. We let $\Theta_j(h)$ be a Poisson random variable with mean $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h)$ that is independent of $s_j(x + h e_i)$. In this case, we can visualize $s_j(x)$ as the sum of $s_j(x + h e_i)$ and $\Theta_j(h)$. We assume that $\{\Theta_j(h) : j \in \mathcal{J}\}$ are independent of each other and use $\Theta(h)$ to denote the vector $\{\Theta_j(h) : j \in \mathcal{J}\}$. Replicating the proofs of Lemmas 2 and 3, it is straightforward to show that we have $\mathbb{P}\{\Theta(h) = 0\} = 1 - \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h)$ and $\mathbb{P}\{\Theta(h) = \epsilon_j\} = a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h)$ for all $j \in \mathcal{J}$. The total probability of the events that are not covered by these $1 + |\mathcal{J}|$ events is $o(h)$.

In the second step, we show that $\lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i))\} = \mathbb{E}\{V(s(x))\}$. Noting that we can visualize $s(x)$ as the sum of the two independent random variables $s(x + h e_i)$ and $\Theta(h)$, we have the conditional expectations $\mathbb{E}\{V(s(x)) | \Theta(h) = 0\} = \mathbb{E}\{V(s(x + h e_i))\}$ and $\mathbb{E}\{V(s(x)) | \Theta(h) = \epsilon_j\} = \mathbb{E}\{V(s(x + h e_i) + \epsilon_j)\}$. Since $V(s(x + h e_i)) \leq \bar{V}$ w.p.1 and the total probability of the events that are not covered by the events $\{\Theta(h) = 0\}$ and $\{\Theta(h) = \epsilon_j \text{ for some } j \in \mathcal{J}\}$ is $o(h)$, we obtain

$$\begin{aligned} \mathbb{E}\{V(s(x))\} &= \mathbb{E}\{\mathbb{E}\{V(s(x)) | \Theta(h)\}\} \leq \mathbb{P}\{\Theta(h) = 0\} \mathbb{E}\{V(s(x + h e_i))\} \\ &\quad + \sum_{j \in \mathcal{J}} \mathbb{P}\{\Theta(h) = \epsilon_j\} \mathbb{E}\{V(s(x + h e_i) + \epsilon_j)\} + \bar{V} o(h). \end{aligned}$$

Since $\lim_{h \downarrow 0} \mathbb{P}\{\Theta(h) = 0\} = 1$ and $\lim_{h \downarrow 0} \mathbb{P}\{\Theta(h) = \epsilon_j\} = 0$ for all $j \in \mathcal{J}$, taking the limit in the right side above as h approaches to zero from right, we obtain $\mathbb{E}\{V(s(x))\} \leq \lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i))\}$. It is possible to see that the reverse inequality also holds by following the same argument, but using the fact that $V(s(x + h e_i)) \geq -\bar{V}$ w.p.1. Therefore, we have $\lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i))\} = \mathbb{E}\{V(s(x))\}$. One can use a similar argument to show that $\lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i) + \epsilon_j)\} = \mathbb{E}\{V(s(x) + \epsilon_j)\}$ for all $j \in \mathcal{J}$.

In the third step, we consider the limit in (2) as h approaches to zero from left. If we replace h by $-h$, then we obtain

$$\lim_{h \uparrow 0} [\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(x - h e_i))\}]/h = \lim_{h \downarrow 0} [\mathbb{E}\{V(s(x + h e_i))\} - \mathbb{E}\{V(s(x))\}]/h.$$

Assuming that $h \geq 0$, since we can visualize $s(x)$ as the sum of the two independent random variables $s(x + h e_i)$ and $\Theta(h)$, we have the conditional expectations, $\mathbb{E}\{V(s(x + h e_i)) - V(s(x)) | \Theta(h) = 0\} = 0$

and $\mathbb{E}\{V(s(x + h e_i)) - V(s(x)) \mid \Theta(h) = \epsilon_j\} = \mathbb{E}\{V(s(x + h e_i)) - V(s(x + h e_i) + \epsilon_j)\}$. Noting that $V(s(x)) - V(s(x + h e_i)) \leq 2\bar{V}$ w.p.1 and the total probability of the events that are not covered by the events $\{\Theta(h) = 0\}$ and $\{\Theta(h) = \epsilon_j \text{ for some } j \in \mathcal{J}\}$ is $o(h)$, we obtain

$$\begin{aligned} \mathbb{E}\{V(s(x + h e_i)) - V(s(x))\} &= \mathbb{E}\{\mathbb{E}\{V(s(x + h e_i)) - V(s(x)) \mid \Theta(h)\}\} \\ &\leq \sum_{j \in \mathcal{J}} \mathbb{P}\{\Theta(h) = \epsilon_j\} \mathbb{E}\{V(s(x + h e_i)) - V(s(x + h e_i) + \epsilon_j)\} + 2\bar{V}o(h). \end{aligned}$$

Dividing both sides of the inequality above by h , using the probability $\mathbb{P}\{\Theta(h) = \epsilon_j\} = a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) h + o(h)$ and taking the limit as h approaches to zero from right, we obtain

$$\begin{aligned} \lim_{h \downarrow 0} \frac{\mathbb{E}\{V(s(x + h e_i)) - V(s(x))\}}{h} &\leq \sum_{j \in \mathcal{J}} a_{ij} \lim_{h \downarrow 0} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k - a_{ij} h) \lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i)) - V(s(x + h e_i) + \epsilon_j)\} \\ &\leq \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j (f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) \mathbb{E}\{V(s(x)) - V(s(x) + \epsilon_j)\}, \end{aligned}$$

where we use the fact that $\dot{\theta}_j(\cdot)$ is Lipschitz continuous, $\lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i))\} = \mathbb{E}\{V(s(x))\}$ and $\lim_{h \downarrow 0} \mathbb{E}\{V(s(x + h e_i) + \epsilon_j)\} = \mathbb{E}\{V(s(x) + \epsilon_j)\}$. It is possible to see that the reverse inequality also holds by following the same argument, but using the fact that $V(s(x)) - V(s(x + h e_i)) \geq -2\bar{V}$ w.p.1. This establishes that the limit in (2) as h approaches to zero from left is given by the expression on the right side of (2).

B APPENDIX: PROPOSITION 4.1 FROM BERTSEKAS AND TSITSIKLIS (1996)

For a function $F(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$, we consider solving the problem $\max_x F(x)$ by using the algorithm

$$x^{k+1} = x^k + \sigma^k d^k,$$

where $\{\sigma^k\}_k$ is a sequence of step size parameters and $\{d^k\}_k$ is a sequence of possibly random step directions. We let \mathcal{F}^k be the filtration generated by the random variables $\{x^1, d^1, \dots, d^{k-1}\}$ in this algorithm and assume that the following statements hold.

- (B.1) There exists a finite scalar \bar{F} that satisfies $F(x) \leq \bar{F}$ for all $x \in \mathbb{R}^n$.
- (B.2) We have $\mathbb{E}\{d^k \mid \mathcal{F}^k\} = \nabla F(x^k)$ w.p.1 for all $k = 1, 2, \dots$
- (B.3) There exists a finite scalar \bar{d} such that we have $\|d^k\| \leq \bar{d}$ w.p.1 for all $k = 1, 2, \dots$
- (B.4) There exists a finite scalar L_F such that we have $\|\nabla F(x) - \nabla F(y)\| \leq L_F \|x - y\|$ for all $x, y \in \mathbb{R}^n$.

In this case, the next convergence result is from Proposition 4.1 in Bertsekas and Tsitsiklis (1996).

Proposition 6 *Assume that the sequence of step size parameters $\{\sigma^k\}_k$ are \mathcal{F}^k -measurable and satisfy $\sigma^k \geq 0$ for all $k = 1, 2, \dots$, $\sum_{k=1}^{\infty} \sigma^k = \infty$ and $\sum_{k=1}^{\infty} [\sigma^k]^2 < \infty$ w.p.1. If the sequence $\{x^k\}_k$ is generated by the algorithm above and (B.1)-(B.4) hold, then we have $\lim_{k \rightarrow \infty} \nabla F(x^k) = 0$ w.p.1 and every limit point x^* of the sequence $\{x^k\}_k$ satisfies $\nabla F(x^*) = 0$ w.p.1.*

C APPENDIX: LIPSCHITZ CONTINUITY OF $\mathbb{E}\{V(s(x))\}$ AND $\mathbb{E}\{V(s(x) + \epsilon_j)\}$

The proof of Proposition 5 uses that fact that $\mathbb{E}\{V(s(x))\}$ and $\mathbb{E}\{V(s(x) + \epsilon_j)\}$ are Lipschitz functions of x . The goal of this section is to show that $\mathbb{E}\{V(s(x))\}$ and $\mathbb{E}\{V(s(x) + \epsilon_j)\}$ are indeed Lipschitz functions of x . We begin with the next lemma.

Lemma 7 *We let e_i be the n -dimensional unit vector with a one in the component corresponding to i and consider the function $F(\cdot) : \mathfrak{R}^n \rightarrow \mathfrak{R}$. If there exists a finite scalar L_F that satisfies $|F(x) - F(x - h e_i)| \leq L_F h$ for all $h \geq 0$, $x \in \mathfrak{R}^n$, $i = 1, \dots, n$, then we have $|F(x) - F(y)| \leq 2nL_F \|x - y\|$ for all $x, y \in \mathfrak{R}^n$.*

Proof We consider $x, y \in \mathfrak{R}^n$ such that $x_i \geq y_i$ for all $i = 1, \dots, n$. Writing $F(x) - F(y)$ as a telescoping sum, we have

$$\begin{aligned} |F(x) - F(y)| &= \left| \sum_{i=1}^n \{F(x - \sum_{j=1}^{i-1} [x_j - y_j] e_j) - F(x - \sum_{j=1}^i [x_j - y_j] e_j)\} \right| \\ &\leq \sum_{i=1}^n |F(x - \sum_{j=1}^{i-1} [x_j - y_j] e_j) - F(x - \sum_{j=1}^i [x_j - y_j] e_j)| \leq \sum_{i=1}^n L_F [x_i - y_i], \end{aligned} \quad (11)$$

where the second inequality follows from the assumption that $|F(x) - F(x - h e_i)| \leq L_F h$ for all $h \geq 0$, $x \in \mathfrak{R}^n$ and $i = 1, \dots, n$. In this case, for arbitrary $x, y \in \mathfrak{R}^n$, we let $z_i = \max\{x_i, y_i\}$ and use z to denote the vector $\{z_i : i = 1, \dots, n\}$. Since $z_i \geq x_i$ and $z_i \geq y_i$ for all $i = 1, \dots, n$, (11) implies that

$$\begin{aligned} |F(x) - F(y)| &\leq |F(z) - F(x)| + |F(z) - F(y)| \leq \sum_{i=1}^n L_F [z_i - x_i] + \sum_{i=1}^n L_F [z_i - y_i] \\ &\leq \sum_{i=1}^n L_F |x_i - y_i| + \sum_{i=1}^n L_F |x_i - y_i|, \end{aligned}$$

where the third inequality follows from the fact that $\max\{x_i, y_i\} - x_i \leq |x_i - y_i|$ and $\max\{x_i, y_i\} - y_i \leq |x_i - y_i|$. The result follows by noting that the last expression in the chain of inequalities above is bounded by $2nL_F \|x - y\|$. \square

Proposition 8 *There exists a finite scalar L_V that satisfies $|\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(y))\}| \leq L_V \|x - y\|$ for all $x, y \in \mathfrak{R}^{|\mathcal{L}|}$.*

Proof Noting Lemma 7, it is enough to show that there exists a finite scalar \bar{L}_V that satisfies $|\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(x - h e_i))\}| \leq \bar{L}_V h$ for all $h \geq 0$, $x \in \mathfrak{R}^{|\mathcal{L}|}$, $i \in \mathcal{L}$, in which case, the result follows by letting $L_V = 2|\mathcal{L}|\bar{L}_V$. We fix $x \in \mathfrak{R}^{|\mathcal{L}|}$ and $i \in \mathcal{L}$, and begin with a construction that is similar to the one at the beginning of Section 3. The mean of the Poisson random variable $s_j(x - h e_i)$ is $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + a_{ij} h)$. On the other hand, the mean value theorem implies that

$$\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + a_{ij} h) = \theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k) + a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + \nu a_{ij} h) h,$$

where $\nu \in [0, 1]$ potentially depends on x, i, j and h . Therefore, the mean of $s_j(x - h e_i)$ can be written as the sum of the two nonnegative terms $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$ and $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + \nu a_{ij} h) h$.

The first term $\theta_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k)$ is the mean of the Poisson random variable $s_j(x)$. We let $\Xi_j(h)$ be a Poisson random variable with mean $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + \nu a_{ij} h) h$ that is independent of $s_j(x)$ so that we can visualize $s_j(x - h e_i)$ as the sum of $s_j(x)$ and $\Xi_j(h)$. We assume that $\{\Xi_j(h) : j \in \mathcal{J}\}$ are independent of each other and use $\Xi(h)$ to denote the vector $\{\Xi_j(h) : j \in \mathcal{J}\}$.

Similar to the argument in the proof of Proposition 4, since we can visualize $s(x - h e_i)$ as the sum of the two independent random variables $s(x)$ and $\Xi(h)$, we have the conditional expectation $\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) | \Xi(h) = 0\} = 0$, which implies that

$$\begin{aligned} \mathbb{E}\{V(s(x)) - V(s(x - h e_i))\} &= \mathbb{E}\{\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) | \Xi(h)\}\} \\ &= \mathbb{P}\{\Xi(h) \neq 0\} \mathbb{E}\{V(s(x)) - V(s(x - h e_i)) | \Xi(h) \neq 0\}. \end{aligned} \quad (12)$$

If X is a Poisson random variable with mean β , then we have $\mathbb{P}\{X \neq 0\} = 1 - e^{-\beta} \leq \beta$. The inequality can be seen by noting that the functions $f(\beta) = 1 - e^{-\beta}$ and $g(\beta) = \beta$ take the same value at $\beta = 0$ and the derivative of $f(\cdot)$ is smaller than the derivative of $g(\cdot)$ for all $\beta \geq 0$. Therefore, for the probability on the right side of the expression above, we obtain the bound

$$\begin{aligned} \mathbb{P}\{\Xi(h) \neq 0\} &= \mathbb{P}\{\Xi_j(h) \neq 0 \text{ for some } j \in \mathcal{J}\} \leq \sum_{j \in \mathcal{J}} \mathbb{P}\{\Xi_j(h) \neq 0\} \\ &\leq \sum_{j \in \mathcal{J}} a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + \nu a_{ij} h) h \leq \sum_{j \in \mathcal{J}} a_{ij} q_j \lambda_j L_\theta h, \end{aligned}$$

where second inequality follows from the fact that $\Xi_j(h)$ is a Poisson random variable with mean $a_{ij} \dot{\theta}_j(f_j - \sum_{k \in \mathcal{L}} a_{kj} x_k + \nu a_{ij} h) h$ and the third inequality follows from the fact that $\theta_j(\cdot)$ is Lipschitz with modulus $q_j \lambda_j L_\theta$ so that its derivative is bounded by $q_j \lambda_j L_\theta$. Using the fact that $|V(s(x)) - V(s(x - h e_i))| \leq 2\bar{V}$ w.p.1, (12) implies that

$$\begin{aligned} |\mathbb{E}\{V(s(x)) - V(s(x - h e_i))\}| &= \mathbb{P}\{\Xi(h) \neq 0\} |\mathbb{E}\{V(s(x)) - V(s(x - h e_i)) | \Xi(h) \neq 0\}| \\ &\leq \sum_{j \in \mathcal{J}} a_{ij} q_j \lambda_j L_\theta h 2\bar{V}. \end{aligned}$$

Letting $\bar{a}_j = \max_{i \in \mathcal{L}} \{a_{ij}\}$ and $\bar{L}_V = \sum_{j \in \mathcal{J}} \bar{a}_j q_j \lambda_j L_\theta 2\bar{V}$, we obtain $|\mathbb{E}\{V(s(x)) - V(s(x - h e_i))\}| = |\mathbb{E}\{V(s(x))\} - \mathbb{E}\{V(s(x - h e_i))\}| \leq \bar{L}_V h$ and the result follows. \square

Proposition 9 *There exists a finite scalar L'_V that satisfies $|\mathbb{E}\{V(s(x) + \epsilon_j)\} - \mathbb{E}\{V(s(y) + \epsilon_j)\}| \leq L'_V \|x - y\|$ for all $x, y \in \mathbb{R}^{|\mathcal{L}|}$, $j \in \mathcal{J}$.*

Proof The proof follows from the same argument in the proof of Proposition 8. \square