

A Stochastic Approximation Method for the Single-Leg Revenue Management Problem with Discrete Demand Distributions

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

Tel: 1-607-255-0698 Fax: 1-607-255-9129

(Corresponding Author)

June 13, 2008

Abstract

We consider the problem of optimally allocating the seats on a single flight leg to the demands from multiple fare classes that arrive sequentially. It is well-known that the optimal policy for this problem is characterized by a set of protection levels. In this paper, we develop a new stochastic approximation method to compute the optimal protection levels under the assumption that the demand distributions are not known and we only have access to the samples from the demand distributions. The novel aspect of our method is that it works with the nonsmooth version of the problem where the capacity can only be allocated in integer quantities. We show that the sequence of protection levels generated by our method converges to a set of optimal protection levels with probability one. We discuss applications to the case where the demand information is censored by the seat availability. Computational experiments indicate that our method is especially advantageous when the total expected demand exceeds the capacity by a significant margin and we do not have good a priori estimates of the optimal protection levels.

Keywords: Revenue management, stochastic approximation, dynamic programming, subgradient optimization.

A canonical problem in the revenue management literature involves optimally allocating the seats on a single flight leg to the demands from multiple fare classes that arrive sequentially. Given the demand from the current fare class and the number of unsold seats, the decision is how many seats to sell to the current fare class. It is well-known that the optimal policy to this problem is characterized by one protection level per fare class. In particular, letting n be the number of fare classes, there exists a set of protection levels $\{y_j^* : j = 1, \dots, n\}$ such that the optimal policy keeps the number of unsold seats just after making the decisions for fare class j as close as possible to y_j^* . In other words, letting x_j be the number of unsold seats just before making the decisions for fare class j and $[\cdot]^+ = \max\{0, \cdot\}$, it is optimal to make $[x_j - y_j^*]^+$ seats available for sale to fare class j . If the demand from fare class j does not exceed $[x_j - y_j^*]^+$, then we satisfy all of the demand. Otherwise, we sell $[x_j - y_j^*]^+$ seats. This structure of the optimal policy arises from the fact that the value functions in the dynamic programming formulation of the problem are concave in the number of unsold seats. In this case, the computation of the optimal protection levels through the Bellman equations requires solving a number of convex optimization problems, which is a simple task as long as the demand distributions are known.

In this paper, we develop a stochastic approximation method to compute the optimal protection levels when the demand distributions are not known and we only have access to the samples from the demand distributions. We work with a particular version of the problem where the demand distributions are discrete and the fare classes that generate lower revenues arrive earlier than the ones that generate higher revenues. We develop a new method that uses the dynamic programming formulation of the problem in conjunction with the samples from the demand distributions to approximate the stochastic subgradients of the value functions. By showing that our approximate stochastic subgradients are indeed accurate in the limit, we establish that the iterates of our stochastic approximation method converge to a set of optimal protection levels with probability one (w.p.1). Furthermore, we consider the case where the demand information is censored by the seat availability. We show that our method remains applicable when we only observe the number of seats sold to a fare class and whether any demand from this fare class is turned down. In particular, if a portion of the demand from a fare class cannot be satisfied, then we do not need to know the exact quantity of the demand that is turned down.

There has been work on using stochastic approximation methods to compute the optimal protection levels, but our paper makes several unique contributions. Brumelle and McGill (1993) characterize the conditions that should be satisfied by the optimal protection levels and van Ryzin and McGill (2000) exploit these conditions to develop a stochastic approximation method. However, this method is tightly related to the conditions in Brumelle and McGill (1993) and it is not clear whether it can be extended to another problem class. In contrast, we work with the dynamic programming formulation of the problem and it is possible to extend our method to numerous inventory control problems where the value functions are convex and the base stock policies are optimal; see Kunnumkal and Topaloglu (2007). Furthermore, the step directions used by our method are related to the stochastic subgradients of the value functions, whereas this is not the case for the method proposed by van Ryzin and McGill (2000). Stochastic subgradients of the value functions are useful when making tactical decisions such as setting the capacity of the flight leg. It is also possible to establish that the theoretical convergence rate of a variant of our method is more favorable than that of the method proposed by van Ryzin and

McGill (2000). Huh and Rusmevichientong (2006) also propose a stochastic approximation method to compute the optimal protection levels. There are similarities between their method and ours as both exploit the dynamic programming formulation of the problem, but Huh and Rusmevichientong (2006) use the results from the online convex optimization literature pioneered by Zinkevich (2003), whereas we use the stochastic approximation theory. Finally, both van Ryzin and McGill (2000) and Huh and Rusmevichientong (2006) assume that the demand distributions are continuous. To our knowledge, our method is the only one that works with discrete demand distributions and has a convergence guarantee for the performance of the policy. To deal with discrete demand distributions, van Ryzin and McGill (2000) propose a randomized version of their method, but the performance of the policy obtained by this method is not necessarily optimal.

It is important to note that the total expected revenue for the seat allocation problem is not concave when viewed as a function of the protection levels. To illustrate, we assume that there are three fare classes and let $\{r_1, r_2, r_3\}$, $\{D_1, D_2, D_3\}$ and $\{y_1, y_2, y_3\}$ respectively be the revenues, demand random variables and protection levels for the three fare classes. If there are x_j unsold seats just before making the decisions for fare class j , then we make $[x_j - y_j]^+$ seats available for sale to fare class j . In this case, the number of seats sold to fare class j is given by the random variable $P_j(x_j, y_j) = \min\{[x_j - y_j]^+, D_j\}$ and the number of unsold seats just before making the decisions for fare class $j + 1$ is given by the random variable $X_{j+1}(x_j, y_j) = x_j - P_j(x_j, y_j)$. Therefore, letting c be the initial capacity, the total expected revenue obtained by using the protection levels $\{y_1, y_2, y_3\}$ can be written as

$$R(c, y_1, y_2, y_3) = r_1 \mathbb{E}\{P_1(c, y_1)\} + r_2 \mathbb{E}\{P_2(X_2(c, y_1), y_2)\} + r_3 \mathbb{E}\{P_3(X_3(X_2(c, y_1), y_2), y_3)\}$$

and the optimal protection levels can be found by solving the problem

$$\max_{(y_1, y_2, y_3)} R(c, y_1, y_2, y_3). \tag{1}$$

Figure 1 plots $R(c, y_1, y_2, 0)$ as a function of (y_1, y_2) for a particular problem instance and indicates that the total expected revenue may not be concave when viewed as a function of the protection levels. Therefore, if we naively attempt to compute the optimal protection levels by solving problem (1) through a stochastic approximation method, then we may get stuck at a locally optimal solution. Our results in this paper, however, show that it is indeed possible to develop a stochastic approximation method to compute the optimal protection levels as long as we use step directions that are based on the dynamic programming formulation of the problem.

Besides computing the optimal protection levels, our method allows us to exploit the well-known advantages of stochastic approximation methods. In particular, our method only requires the ability to obtain samples from the demand distributions. Parametric representations of the demand distributions are not necessary. Furthermore, our method remains applicable when we only observe the number of seats sold to a fare class and whether any demand from this fare class is turned down. In other words, we do not need to know the exact quantity of the demand that is turned down when a portion of the demand from a fare class cannot be satisfied. Therefore, we can still compute the optimal protection levels by using our method when the demand information is censored by the seat availability.

The rest of the paper is organized as follows. Section 1 briefly reviews the other related literature. Section 2 gives a dynamic programming formulation of the seat allocation problem and summarizes the important results from the existing literature. Section 3 describes our stochastic approximation method and Section 4 proves its convergence. Section 5 considers the case where the demand information is censored by the seat availability. Section 6 provides computational experiments. Section 7 concludes.

1 REVIEW OF OTHER RELATED LITERATURE

Most of the literature on the single flight leg problem assumes that there are multiple fare classes and the demands from different fare classes occur over nonoverlapping time intervals. This ensures that we can formulate the problem as a dynamic program with the number of decision epochs being equal to the number of fare classes. Motivated by the fact that the leisure travelers tend to book earlier than the business travelers, it is also a common assumption that the fare classes that generate lower revenues arrive earlier than the ones that generate higher revenues. An important consequence of the second assumption is that the optimal protection levels are nested. In other words, the optimal protection level for a given fare class is larger than the optimal protection levels for the fare classes that arrive later.

Littlewood (1972), Curry (1990), Wollmer (1992) and Brumelle and McGill (1993) employ the two assumptions in the previous paragraph and show the optimality of protection level policies and the nestedness of the optimal protection levels. Robinson (1995) shows that protection level policies are still optimal under the assumption that the demands from different fare classes occur over nonoverlapping time intervals, but the fare classes that generate lower revenues do not necessarily arrive earlier. Lee and Hersh (1993) and Lautenbacher and Stidham (1999) focus on the single leg problem when the demands from different fare classes do not necessarily occur over nonoverlapping time intervals and show that variations of protection level policies are still optimal. We refer the reader to Talluri and van Ryzin (2004) for a coverage of the related revenue management literature.

Although airlines generally operate hub-and-spoke networks, many local airlines are very interested in the seat allocation problem on a single flight leg. Furthermore, the seat allocation problem on a single flight leg has important implications for the hub-and-spoke networks. To begin with, it justifies, at least to a certain extent, the use of protection level policies for complex networks. In addition, there are a variety of techniques to decompose the revenue management problem over a network into a sequence of single flight leg problems; see Section 3.4.4 in Talluri and van Ryzin (2004). In this case, it becomes important to solve the seat allocation problem on a single flight leg under a variety of assumptions.

The use of stochastic approximation methods for solving optimization problems under uncertainty is well-known. Kushner and Clark (1978), Ermoliev (1988) and Bertsekas and Tsitsiklis (1996) give a coverage of the theory of stochastic approximation methods. There are numerous papers that use these methods for solving the revenue management problem over a network. In particular, van Ryzin and Vulcano (2004), Bertsimas and de Boer (2005) and van Ryzin and Vulcano (2006) describe methods to compute protection levels, Topaloglu (2007) describes a method to compute bid prices and Karaesmen and van Ryzin (2004) describe a method to compute overbooking limits. As far as other application

areas are concerned, L'Ecuyer and Glynn (1994), Fu (1994), Glasserman and Tayur (1995), Bashyam and Fu (1998) and Mahajan and van Ryzin (2001) focus on queueing and inventory control. Kunnumkal and Topaloglu (2007) show that a method similar to ours can be used to compute the optimal base stock levels in inventory control problems. However, since they work with continuous demand distributions, their proof technique is considerably different from ours.

2 PROBLEM FORMULATION

We want to use c seats available on a single flight leg to satisfy the demands from n fare classes that arrive sequentially. We index the fare classes such that the demand from fare class 1 arrives first and the demand from fare class n arrives last. If we sell a seat to fare class j , then we generate a revenue of r_j . We assume that the revenues satisfy $0 < r_1 \leq r_2 \leq \dots \leq r_n$ so that the demands from the cheaper fare classes arrive earlier. The demands from different fare classes are random and we use D_j to denote the demand from fare class j . We assume that D_j is a positive and integer random variable, and $\{D_j : j = 1, \dots, n\}$ are independent of each other. We are interested in maximizing the total expected revenue from n fare classes.

Using x_j to denote the remaining capacity just before making the decisions for fare class j , u_j to denote the number of seats sold to fare class j and d_j to denote a particular realization of D_j , the optimal policy can be found by solving the optimality equation

$$v_j(x_j, d_j) = \max_{0 \leq u_j \leq \min\{x_j, d_j\}} \left\{ r_j u_j + \mathbb{E}\{v_{j+1}(x_j - u_j, D_{j+1})\} \right\}, \quad (2)$$

with $v_{n+1}(\cdot, \cdot) = 0$. The constraints in the problem above ensure that the number of seats sold do not exceed the remaining capacity and the demand from fare class j . Alternatively, if we let $y_j = x_j - u_j$ be the remaining capacity just after making the decisions for fare class j , then (2) can be written as

$$v_j(x_j, d_j) = \max_{[x_j - d_j]^+ \leq y_j \leq x_j} \left\{ -r_j y_j + \mathbb{E}\{v_{j+1}(y_j, D_{j+1})\} \right\} + r_j x_j, \quad (3)$$

where the constraints follow from the fact that $y_j = x_j - u_j \geq x_j - \min\{x_j, d_j\} = \max\{0, x_j - d_j\}$ and $y_j = x_j - u_j \leq x_j$.

It is possible to show that $\{v_j(\cdot, D_j) : j = 1, \dots, n\}$ are piecewise-linear concave functions with points of nondifferentiability being a subset of integers; see Brumelle and McGill (1993). Using this result, Brumelle and McGill (1993) show that the optimal policy is characterized by a set of protection levels $\{y_j^* : j = 1, \dots, n\}$, where y_j^* can be computed as a maximizer of the function

$$f_j(y_j) = -r_j y_j + \mathbb{E}\{v_{j+1}(y_j, D_{j+1})\} \quad (4)$$

over the interval $[0, c]$. This is to say that if the remaining capacity just before making the decisions for fare class j is x_j , then it is optimal to make $[x_j - y_j^*]^+$ seats available for sale to fare class j . In this case, if the demand from fare class j is d_j , then we sell $\min\{[x_j - y_j^*]^+, d_j\}$ seats to fare class j . The protection level terminology is due to the fact that it is optimal to protect y_j^* seats for the demand from fare classes $\{j + 1, \dots, n\}$ when making the decisions for fare class j .

Using the fact that the demands from the cheaper fare classes arrive earlier, Brumelle and McGill (1993) also show that the optimal protection levels are nested. That is, the optimal number of seats to protect for the demand from fare classes $\{j, \dots, n\}$ is at least as large as the optimal number of seats to protect for the demand from fare classes $\{j+1, \dots, n\}$. To state this mathematically, we let

$$\mathcal{Y}_j^* = \operatorname{argmax}_{0 \leq y_j \leq c} f_j(y_j). \quad (5)$$

Therefore, we can use any element of \mathcal{Y}_j^* as the optimal protection level when making the decisions for fare class j . The fact that the optimal protection levels are nested implies that $\min_{y_j \in \mathcal{Y}_j^*} y_j \geq \min_{y_{j+1} \in \mathcal{Y}_{j+1}^*} y_{j+1}$ and $\max_{y_j \in \mathcal{Y}_j^*} y_j \geq \max_{y_{j+1} \in \mathcal{Y}_{j+1}^*} y_{j+1}$ for all $j = 1, \dots, n-1$. In this case, we can choose $y_1^* \in \mathcal{Y}_1^*$, $y_2^* \in \mathcal{Y}_2^*$, \dots , $y_n^* \in \mathcal{Y}_n^*$ such that $y_1^* \geq y_2^* \geq \dots \geq y_n^*$.

We note that our dynamic programming formulation differs from the existing literature in two relatively minor aspects. First, we index the fare classes such that fare classes 1 and n respectively correspond to the cheapest and most expensive fare classes, whereas the existing literature usually indexes the fare classes in the reverse order. The motivation for our choice is that it is common to refer to a cheaper fare class as a lower fare class and it is more consistent to index a cheaper fare class with a smaller integer. Second, we use a two-dimensional state variable in (2), whereas the existing literature generally uses a one-dimensional state variable. It is possible to use a one-dimensional state variable in (2) by simply letting $\hat{v}_j(x_j) = \mathbb{E}\{v_j(x_j, D_j)\}$ and taking the expectations of both sides to obtain

$$\hat{v}_j(x_j) = \mathbb{E} \left\{ \max_{0 \leq u_j \leq \min\{x_j, D_j\}} \left\{ r_j u_j + \hat{v}_{j+1}(x_j - u_j) \right\} \right\}.$$

However, the optimality equation in (2) will be more useful for the development in the paper. We also note that it is possible to relax the assumption that the cheaper fare classes arrive earlier, but we obtain weaker results without this assumption. We briefly describe this extension in Section 7.

3 STOCHASTIC APPROXIMATION METHOD

In this section, we use the dynamic programming formulation of the seat allocation problem to develop a stochastic approximation method that computes the optimal protection levels.

The optimal protection level for fare class j is a maximizer of $f_j(\cdot)$ over the interval $[0, c]$. On the other hand, by (4), we can compute a stochastic subgradient of $f_j(\cdot)$ at y_j through

$$\Delta_j(y_j, d_{j+1}) = -r_j + \dot{v}_{j+1}(y_j, d_{j+1}), \quad (6)$$

where we use $\dot{v}_{j+1}(y_j, d_{j+1})$ to denote a subgradient of $v_{j+1}(\cdot, d_{j+1})$ at y_j . In other words, if we use $\partial v_{j+1}(y_j, d_{j+1})$ to denote the subdifferential of $v_{j+1}(\cdot, d_{j+1})$ at y_j , then we have $\dot{v}_{j+1}(y_j, d_{j+1}) \in \partial v_{j+1}(y_j, d_{j+1})$. Interchanging the orders of all expectations and subgradients throughout the paper trivially follows from the fact that the demand distributions are discrete and the capacity is finite. In this case, letting $\{y_j^k : j = 1, \dots, n\}$ be the estimates of the optimal protection levels at iteration k , $\{D_j^k : j = 1, \dots, n\}$ be the demand random variables at iteration k and $\{\alpha_j^k : j = 1, \dots, n\}_k$ be a

sequence of step size parameters, we can update our estimates of the optimal protection levels by the iterative algorithm

$$y_j^{k+1} = \min \{ [y_j^k + \alpha_j^k \Delta_j(y_j^k, D_{j+1}^k)]^+, c \}, \quad (7)$$

where the operator $\min\{\cdot\}^+, c$ ensures that the estimates of the optimal protection levels always lie in the interval $[0, c]$. The algorithm in (7) is motivated by the fact that the optimal protection level for fare class j is a maximizer of $f_j(\cdot)$ over the interval $[0, c]$ and $\Delta_j(y_j^k, D_{j+1}^k)$ is a stochastic subgradient of $f_j(\cdot)$ at y_j^k . If the sequence of protection levels $\{y_j^k : j = 1, \dots, n\}_k$ is generated by the algorithm in (7), then we can use standard results on stochastic approximation methods to show that $\{y_j^k : j = 1, \dots, n\}_k$ converges to a set of optimal protection levels w.p.1.

The algorithm in (7) is clearly not realistic since the computation in (6) requires the knowledge of $\{\dot{v}_j(\cdot, \cdot) : j = 1, \dots, n\}$. The stochastic approximation method that we propose in this section is based on constructing tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \dots, n\}$. We begin by constructing tractable approximations to $\{\dot{v}_j(\cdot, \cdot) : j = 1, \dots, n\}$. Since $f_j(\cdot)$ is concave and the optimal protection level y_j^* is a maximizer of $f_j(\cdot)$ over the interval $[0, c]$, we have

$$\max_{[x_j - d_j]^+ \leq y_j \leq x_j} \{f_j(y_j)\} = \begin{cases} f_j([x_j - d_j]^+) & \text{if } y_j^* < [x_j - d_j]^+ \\ f_j(y_j^*) & \text{if } [x_j - d_j]^+ \leq y_j^* \leq x_j \\ f_j(x_j) & \text{if } x_j < y_j^* \end{cases} \quad (8)$$

for all $x_j \in [0, c]$. On the other hand, we can write (3) as

$$v_j(x_j, d_j) = \max_{[x_j - d_j]^+ \leq y_j \leq x_j} \{f_j(y_j)\} + r_j x_j,$$

in which case (8) and the definition of $f_j(\cdot)$ in (4) imply that

$$v_j(x_j, d_j) = \begin{cases} -r_j [x_j - d_j]^+ + r_j x_j + \mathbb{E}\{v_{j+1}([x_j - d_j]^+, D_{j+1})\} & \text{if } y_j^* < [x_j - d_j]^+ \\ -r_j y_j^* + r_j x_j + \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} & \text{if } [x_j - d_j]^+ \leq y_j^* \leq x_j \\ \mathbb{E}\{v_{j+1}(x_j, D_{j+1})\} & \text{if } x_j < y_j^* \end{cases} \quad (9)$$

for all $x_j \in [0, c]$. Noting that $y_j^* \geq 0$, we have $0 < [x_j - d_j]^+$ whenever the condition in the first case above holds. Therefore, we can replace $[x_j - d_j]^+$ in the first case by $x_j - d_j$. On the other hand, the condition in the second case above is equivalent to $x_j - d_j \leq y_j^* \leq x_j$ and $0 \leq y_j^* \leq x_j$. Since we always have $y_j^* \geq 0$, we can replace the condition in the second case by $x_j - d_j \leq y_j^* \leq x_j$. These imply that we can write (9) as

$$v_j(x_j, d_j) = \begin{cases} r_j d_j + \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} & \text{if } y_j^* < x_j - d_j \\ r_j [x_j - y_j^*] + \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} & \text{if } x_j - d_j \leq y_j^* \leq x_j \\ \mathbb{E}\{v_{j+1}(x_j, D_{j+1})\} & \text{if } x_j < y_j^*. \end{cases} \quad (10)$$

In this case, it is easy to see that we can compute a subgradient of $v_j(\cdot, d_j)$ at x_j through the recursion over the fare classes

$$\dot{v}_j(x_j, d_j) = \begin{cases} \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} & \text{if } y_j^* < x_j - d_j \\ r_j & \text{if } x_j - d_j \leq y_j^* \leq x_j \\ \mathbb{E}\{\dot{v}_{j+1}(x_j, D_{j+1})\} & \text{if } x_j < y_j^*. \end{cases} \quad (11)$$

This recursion can heuristically be justified by differentiating both sides of (10) with respect to x_j . However, the value function is not necessarily differentiable and we formally show in the appendix that the recursion in (11) indeed gives a subgradient of $v_j(\cdot, d_j)$ at x_j .

To construct tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \dots, n\}$, we “mimic” the computation in (11) by using the estimates of the optimal protection levels. In particular, letting $\{y_j^k : j = 1, \dots, n\}$ be the estimates of the optimal protection levels at iteration k and using $\mathcal{O}(\cdot)$ to denote the operator that rounds a scalar to a nearest integer by breaking ties arbitrarily, we recursively define

$$\rho_j^k(x_j, d_j, d_{j+1}, \dots, d_n) = \begin{cases} \rho_{j+1}^k(x_j - d_j, d_{j+1}, \dots, d_n) & \text{if } \mathcal{O}(y_j^k) < x_j - d_j \\ r_j & \text{if } x_j - d_j \leq \mathcal{O}(y_j^k) \leq x_j \\ \rho_{j+1}^k(x_j, d_{j+1}, \dots, d_n) & \text{if } x_j < \mathcal{O}(y_j^k), \end{cases} \quad (12)$$

with $\rho_{n+1}^k(\cdot, \cdot, \dots, \cdot) = 0$. We propose using $\rho_j^k(x_j, d_j, d_{j+1}, \dots, d_n)$ to approximate $\dot{v}_j(x_j, d_j)$. More specifically, at iteration k , we replace $\dot{v}_{j+1}(y_j, d_{j+1})$ in (6) with $\rho_{j+1}^k(y_j, d_{j+1}, \dots, d_n)$ and use

$$s_j^k(y_j, d_{j+1}, \dots, d_n) = -r_j + \rho_{j+1}^k(y_j, d_{j+1}, \dots, d_n) \quad (13)$$

to approximate a stochastic subgradient of $f_j(\cdot)$ at y_j . Therefore, we propose the following algorithm to compute the optimal protection levels.

Algorithm 1

Step 1. Initialize the estimates of the optimal protection levels $\{y_j^1 : j = 1, \dots, n\}$ such that $c \geq y_1^1 \geq y_2^1 \geq \dots \geq y_n^1 = 0$. Initialize the iteration counter by setting $k = 1$.

Step 2. Letting $\{D_j^k : j = 1, \dots, n\}$ be the demand random variables at iteration k , set

$$y_j^{k+1} = \max \left\{ \min \left\{ [y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)]^+, c \right\}, \mathcal{O}(y_{j+1}^{k+1}) \right\} \quad (14)$$

for all $j = 1, \dots, n$, with the convention that $y_{n+1}^{k+1} = 0$.

Step 3. Increase k by 1 and go to Step 2.

We let \mathcal{F}^k be the filtration generated by $\{\{y_1^1, \dots, y_n^1\}, \{D_1^1, \dots, D_n^1\}, \dots, \{D_1^{k-1}, \dots, D_n^{k-1}\}\}$. Given \mathcal{F}^k , we assume that the conditional distribution of $\{D_j^k : j = 1, \dots, n\}$ is the same as the distribution of $\{D_j : j = 1, \dots, n\}$. We assume that the step size parameters $\{\alpha_j^k : j = 1, \dots, n\}$ are positive and \mathcal{F}^k -measurable, in which case the estimates of the optimal protection levels $\{y_j^k : j = 1, \dots, n\}$ are also \mathcal{F}^k -measurable. In the next section, we show that if the sequence $\{y_j^k : j = 1, \dots, n\}_k$ is generated by Algorithm 1, then it converges to a set of optimal protection levels w.p.1.

Several remarks are in order for our approximation to $\dot{v}_j(x_j, d_j)$ and Algorithm 1. First, we do not need to compute expectations when computing $\rho_j^k(x_j, d_j, d_{j+1}, \dots, d_n)$. We also round the estimates of the optimal protection levels to nearest integers when computing $\rho_j^k(x_j, d_j, d_{j+1}, \dots, d_n)$. In Section 5, we show that rounding becomes useful when the demand information is censored by the seat availability. Second, comparing (6) with (13) indicates that if $\dot{v}_j(\cdot, \cdot)$ and $\mathbb{E}\{\rho_j^k(\cdot, \cdot, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\}$ are “close” to each other for all $j = 1, \dots, n$, then the expected step directions $\mathbb{E}\{\Delta_j(\cdot, D_{j+1}^k) | \mathcal{F}^k\}$ and

$\mathbb{E}\{s_j^k(\cdot, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\}$ are also “close” to each other. This implies that using the step direction $s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$ instead of $\Delta_j(y_j^k, D_{j+1}^k)$ does not bring too much error in expectation. Indeed, our convergence proof is heavily based on this observation, which we make mathematically precise in Lemma 2 below. Third, the way that we update our estimates of the optimal protection levels in (14) ensures that $y_j^{k+1} \geq \mathcal{O}(y_{j+1}^{k+1})$ for all $k = 1, 2, \dots$, which implies that $\mathcal{O}(y_j^{k+1}) \geq \mathcal{O}(y_{j+1}^{k+1})$. Therefore, we have $c \geq \mathcal{O}(y_1^k) \geq \mathcal{O}(y_2^k) \geq \dots \geq \mathcal{O}(y_n^k) \geq 0$ for all $k = 1, 2, \dots$ and the rounded estimates of the optimal protection levels at each iteration are nested. It is also important to note that the update in (14) is of Gauss-Seidel variant. More specifically, we need the value of y_{j+1}^{k+1} to compute the value of y_j^{k+1} . Therefore, Step 2 in Algorithm 1 has to be carried out starting from fare class n and moving backwards through the fare classes.

4 CONVERGENCE PROOF

In this section, we show that the iterates of Algorithm 1 converge to a set of optimal protection levels w.p.1. We begin with some preliminary results in Section 4.1 and complete the proof in Section 4.2

4.1 PRELIMINARIES

The next lemma establishes a uniform bound on our step directions.

Lemma 1 *There exists a finite scalar M such that we have*

$$|\rho_j^k(x_j, D_j^k, D_{j+1}^k, \dots, D_n^k)| \leq M \text{ and } |s_j^k(x_j, D_{j+1}^k, \dots, D_n^k)| \leq M$$

w.p.1 for all $x_j \in [0, c]$, $j = 1, \dots, n$, $k = 1, 2, \dots$

Proof If we let $R = \max_{j \in \{1, \dots, n\}} r_j$, then by using (12) and moving backwards through the fare classes, it is easy to see that $|\rho_j^k(x_j, D_j^k, D_{j+1}^k, \dots, D_n^k)| \leq R$. By (13), the result follows by letting $M = 2R$. \square

The next lemma shows that if $\{y_j^k : j = 1, \dots, n\}$ get close to the optimal protection levels, then the step directions in Algorithm 1 are related to the stochastic subgradients of the value functions. In Lemma 2 and throughout the rest of the paper, since $\{y_j^k : j = 1, \dots, n\}$ are \mathcal{F}^k -measurable, we treat $\{y_j^k : j = 1, \dots, n\}$ as known constants when dealing with a conditional expectation of the form $\mathbb{E}\{\cdot | \mathcal{F}^k\}$. Furthermore, since $\{v_j(\cdot, D_j) : j = 1, \dots, n\}$ are piecewise-linear concave functions with points of nondifferentiability being a subset of integers, it is easy to see that \mathcal{Y}_j^* in (5) is a closed interval with integer end points. We let $\mathcal{Y}_j^* = [L_j^*, U_j^*]$ throughout the rest of the paper, where L_j^* and U_j^* are integers.

Lemma 2 *Assume that the sequence $\{y_j^k : j = 1, \dots, n\}_k$ is generated by Algorithm 1. If it holds that*

$$y_j^k \in (L_j^* - \frac{1}{2}, U_j^* + \frac{1}{2}), y_{j+1}^k \in (L_{j+1}^* - \frac{1}{2}, U_{j+1}^* + \frac{1}{2}), \dots, y_n^k \in (L_n^* - \frac{1}{2}, U_n^* + \frac{1}{2}),$$

then we have $\mathbb{E}\{\rho_j^k(x_j, D_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_j^k\} \in \partial v_j(x_j, D_j^k)$ w.p.1 for all $x_j \in [0, c]$.

Proof We show the result by induction over the fare classes. It is easy to show the result for fare class n . Assuming that the result holds for fare class $j + 1$, we now show that the result holds for fare class j . The assumption in the lemma implies that we can find $y_j^* \in \mathcal{Y}_j^*$ such that $y_j^* = \mathcal{O}(y_j^k)$. Furthermore, since we have $y_{j+1}^k \in (L_{j+1}^* - \frac{1}{2}, U_{j+1}^* + \frac{1}{2}), \dots, y_n^k \in (L_n^* - \frac{1}{2}, U_n^* + \frac{1}{2})$ by the assumption in the lemma, the induction assumption implies that $\mathbb{E}\{\rho_{j+1}^k(\cdot, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} \in \partial v_{j+1}(\cdot, D_{j+1}^k)$.

Taking the conditional expectations in (12) and recalling that we use $\dot{v}_{j+1}(x_j, d_{j+1})$ to denote an element of $\partial v_{j+1}(x_j, d_{j+1})$, we obtain

$$\begin{aligned} & \mathbb{E}\{\rho_j^k(x_j, d_j, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} \\ &= \mathbb{E}\{\mathbb{E}\{\rho_j^k(x_j, d_j, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} | \mathcal{F}^k\} \\ &= \begin{cases} \mathbb{E}\{\mathbb{E}\{\rho_{j+1}^k(x_j - d_j, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} | \mathcal{F}^k\} & \text{if } y_j^* < x_j - d_j \\ r_j & \text{if } x_j - d_j \leq y_j^* \leq x_j \\ \mathbb{E}\{\mathbb{E}\{\rho_{j+1}^k(x_j, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} | \mathcal{F}^k\} & \text{if } x_j < y_j^* \end{cases} \\ &= \begin{cases} \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1}^k) | \mathcal{F}^k\} & \text{if } y_j^* < x_j - d_j \\ r_j & \text{if } x_j - d_j \leq y_j^* \leq x_j \\ \mathbb{E}\{\dot{v}_{j+1}(x_j, D_{j+1}^k) | \mathcal{F}^k\} & \text{if } x_j < y_j^*, \end{cases} \end{aligned} \quad (15)$$

where the second equality follows from the fact that we can find $y_j^* \in \mathcal{Y}_j^*$ such that $y_j^* = \mathcal{O}(y_j^k)$ and the third equality follows from the induction assumption that $\mathbb{E}\{\rho_{j+1}^k(\cdot, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} \in \partial v_{j+1}(\cdot, D_{j+1}^k)$. Comparing (15) with (11) and noting that the distribution of D_{j+1}^k conditional on \mathcal{F}^k is the same as the distribution of D_{j+1} , we obtain $\mathbb{E}\{\rho_j^k(x_j, d_j, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} \in \partial v_j(x_j, d_j)$. \square

Roughly speaking, the next lemma shows that if the estimates of the optimal protection levels at iteration k are close to the optimal protection levels and the step sizes are small, then the estimates of the optimal protection levels at iteration $k + 1$ are also close to the optimal protection levels.

Lemma 3 *Assume that the sequence $\{y_j^k : j = 1, \dots, n\}_k$ is generated by Algorithm 1 and M is as in Lemma 1. If it holds that*

$$\begin{aligned} & y_j^k \in (L_j^* - \frac{1}{4}, U_j^* + \frac{1}{4}), y_{j+1}^k \in (L_{j+1}^* - \frac{1}{4}, U_{j+1}^* + \frac{1}{4}), \dots, y_n^k \in (L_n^* - \frac{1}{4}, U_n^* + \frac{1}{4}) \text{ and} \\ & \alpha_j^k \in [0, \frac{1}{4M}], \alpha_{j+1}^k \in [0, \frac{1}{4M}], \dots, \alpha_n^k \in [0, \frac{1}{4M}], \end{aligned}$$

then we have $\mathcal{O}(y_j^{k+1}) \in \mathcal{Y}_j^$ w.p.1.*

Proof All statements in the proof are in w.p.1 sense. We show the result by induction over the fare classes. Since $r_n > 0$, we have $\mathcal{Y}_n^* = \{0\}$ by (5) and $s_n^k(\cdot) < 0$ by (13). Therefore, (14), the initial condition that $y_n^1 = 0$ in Step 1 of Algorithm 1 and the fact that $s_n^k(\cdot) < 0$ imply that $y_n^k = 0$ for all $k = 1, 2, \dots$. In this case, we have $y_n^k = 0 \in \mathcal{Y}_n^*$ for all $k = 1, 2, \dots$ and the result trivially holds for fare class n . Assuming that the result holds for fare class $j + 1$, we now show that the result holds for fare class j . By Lemma 1 and the assumption in the current lemma, we have

$$L_j^* - \frac{1}{2} < y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) < U_j^* + \frac{1}{2}. \quad (16)$$

Since we have $y_{j+1}^k \in (L_{j+1}^* - \frac{1}{4}, U_{j+1}^* + \frac{1}{4}), \dots, y_n^k \in (L_n^* - \frac{1}{4}, U_n^* + \frac{1}{4})$ and $\alpha_{j+1}^k \in [0, \frac{1}{4M}], \dots, \alpha_n^k \in [0, \frac{1}{4M}]$ by the assumption in the lemma, the induction assumption implies that $\mathcal{O}(y_{j+1}^{k+1}) \in \mathcal{Y}_{j+1}^*$. We consider three cases.

Case 1. Assume that $y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq c \geq \mathcal{O}(y_{j+1}^{k+1})$. Since we have $\mathcal{O}(y_{j+1}^{k+1}) \in \mathcal{Y}_{j+1}^* \subset [0, c]$ by the induction assumption and $y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq c$, we obtain $y_j^{k+1} = c$ by (14). On the other hand, we have $U_j^* + \frac{1}{2} > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq c$ by (16). Since U_j^* is an integer smaller than c , we obtain $U_j^* = c$. Therefore, we have $\mathcal{O}(y_j^{k+1}) = c \in [c, c] \subset [L_j^*, U_j^*]$.

Case 2. Assume that $c > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq \mathcal{O}(y_{j+1}^{k+1})$. Since Algorithm 1 ensures that $y_j^k \geq 0$ for all $j = 1, \dots, n, k = 1, 2, \dots$, we have $c > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq \mathcal{O}(y_{j+1}^{k+1}) \geq 0$ and we obtain $y_j^{k+1} = y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$ by (14). Therefore, we have $\mathcal{O}(y_j^{k+1}) \in [L_j^*, U_j^*]$ by (16).

Case 3. Assume that $c \geq \mathcal{O}(y_{j+1}^{k+1}) > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. We have $y_j^{k+1} = \mathcal{O}(y_{j+1}^{k+1})$ by (14) and $\mathcal{O}(y_{j+1}^{k+1}) > L_j^* - \frac{1}{2}$ by (16). Since we have $\mathcal{O}(y_{j+1}^{k+1}) \in \mathcal{Y}_{j+1}^*$ by the induction assumption, the fact that the optimal protection levels are nested implies that $L_{j+1}^* \leq \mathcal{O}(y_{j+1}^{k+1}) \leq U_{j+1}^* \leq U_j^*$. Therefore, we obtain $L_j^* - \frac{1}{2} < \mathcal{O}(y_{j+1}^{k+1}) = y_j^{k+1} = \mathcal{O}(y_{j+1}^{k+1}) \leq U_j^*$, which implies that $\mathcal{O}(y_j^{k+1}) \in [L_j^*, U_j^*]$. \square

In Section 4.2, we give a convergence result for Algorithm 1 that shows that the distance between y_j^k and the optimal protection level that is closest to y_j^k converges to zero w.p.1. For fare class j , we define the optimal protection level that is closest to y_j^k as

$$\mathcal{C}_j(y_j^k) = \operatorname{argmin}_{y_j^* \in \mathcal{Y}_j^*} |y_j^* - y_j^k| = \operatorname{argmin}_{y_j^* \in [L_j^*, U_j^*]} |y_j^* - y_j^k|. \quad (17)$$

The next lemma shows a contraction type of result for Algorithm 1.

Lemma 4 *Assume that the sequence $\{y_j^k : j = 1, \dots, n\}_k$ is generated by Algorithm 1 and M is as in Lemma 1. If it holds that*

$$y_{j+1}^k \in (L_{j+1}^* - \frac{1}{4}, U_{j+1}^* + \frac{1}{4}), y_{j+2}^k \in (L_{j+2}^* - \frac{1}{4}, U_{j+2}^* + \frac{1}{4}), \dots, y_n^k \in (L_n^* - \frac{1}{4}, U_n^* + \frac{1}{4}) \text{ and} \\ \alpha_{j+1}^k \in [0, \frac{1}{4M}], \alpha_{j+2}^k \in [0, \frac{1}{4M}], \dots, \alpha_n^k \in [0, \frac{1}{4M}],$$

then we have $|y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})| \leq |y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) - \mathcal{C}_j(y_j^k)|$ w.p.1.

Proof All statements in the proof are in w.p.1 sense. We consider the same three cases in the proof of Lemma 3.

Case 1. Assume that $y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq c \geq \mathcal{O}(y_{j+1}^{k+1})$. By the same argument in Lemma 3, we have $y_j^{k+1} = c$. Since $\mathcal{C}_j(y_j^{k+1})$ is the closest optimal protection level to y_j^{k+1} and $U_j^* \leq c$, we obtain $\mathcal{C}_j(y_j^{k+1}) = U_j^*$. Since we trivially have $\mathcal{C}_j(y_j^k) \leq U_j^* \leq c$, we obtain $y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq c = y_j^{k+1} = c \geq U_j^* = \mathcal{C}_j(y_j^{k+1}) = U_j^* \geq \mathcal{C}_j(y_j^k)$ and the result follows.

Case 2. Assume that $c > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \geq \mathcal{O}(y_{j+1}^{k+1})$. By the same argument in Lemma 3, we have $y_j^{k+1} = y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. Therefore, we have $|y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) - \mathcal{C}_j(y_j^k)| = |y_j^{k+1} - \mathcal{C}_j(y_j^k)| \geq |y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})|$, where the last inequality follows from (17).

Case 3. Assume that $c \geq \mathcal{O}(y_{j+1}^{k+1}) > y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. By the same argument in Lemma 3, we have $y_j^{k+1} = \mathcal{O}(y_{j+1}^{k+1})$. If $\mathcal{O}(y_{j+1}^{k+1}) \in \mathcal{Y}_j^*$, then we have $|y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})| = |\mathcal{O}(y_{j+1}^{k+1}) - \mathcal{C}_j(\mathcal{O}(y_{j+1}^{k+1}))| = 0$ by (17) and the result follows. We now assume that either $\mathcal{O}(y_{j+1}^{k+1}) > U_j^*$ or $\mathcal{O}(y_{j+1}^{k+1}) < L_j^*$. We immediately eliminate the former possibility by noting that we have $\mathcal{O}(y_{j+1}^{k+1}) \in \mathcal{Y}_{j+1}^*$ by Lemma 3 and the assumption in the current lemma, which, together with the fact that the optimal protection levels are nested, implies that $\mathcal{O}(y_{j+1}^{k+1}) \leq U_{j+1}^* \leq U_j^*$. Therefore, we have $y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) < \mathcal{O}(y_{j+1}^{k+1}) = y_j^{k+1} = \mathcal{O}(y_{j+1}^{k+1}) < L_j^*$. This implies that $|y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})| = L_j^* - y_j^{k+1} < L_j^* - y_j^k - \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) \leq \mathcal{C}_j(y_j^k) - y_j^k - \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. \square

4.2 CONVERGENCE OF ALGORITHM 1

We have the following convergence result for Algorithm 1.

Proposition 5 *Assume that the sequence $\{y_j^k : j = 1, \dots, n\}_k$ is generated by Algorithm 1. If the sequence of step size parameters $\{\alpha_j^k : j = 1, \dots, n\}_k$ is positive and satisfies $\sum_{k=1}^{\infty} \alpha_j^k = \infty$ and $\sum_{k=1}^{\infty} [\alpha_j^k]^2 < \infty$ w.p.1 for all $j = 1, \dots, n$, then we have $\lim_{k \rightarrow \infty} |y_j^k - \mathcal{C}_j(y_j^k)| = 0$ w.p.1 for all $j = 1, \dots, n$.*

Proof All statements in the proof are in w.p.1 sense. We show the result by induction over the fare classes. Since we have $\mathcal{Y}_n^* = \{0\}$ and $y_n^k = 0$ for all $k = 1, 2, \dots$ by the argument in the proof of Lemma 3, the result holds for fare class n . Assuming that the result holds for fare classes $\{j+1, j+2, \dots, n\}$, we now show that the result holds for fare class j . The proof is in three parts. The first part shows that an inequality of the form $\mathbb{E}\{Y^{k+1} | \mathcal{F}^k\} \leq Y^k - X^k + Z^k$ holds for appropriately defined sequences $\{X^k\}_k$, $\{Y^k\}_k$ and $\{Z^k\}_k$. The second part shows that $\{X^k\}_k$, $\{Y^k\}_k$ and $\{Z^k\}_k$ are positive and \mathcal{F}^k -measurable, and $\{Z^k\}_k$ satisfies $\sum_{k=1}^{\infty} Z^k < \infty$. In this case, we conclude by using the supermartingale convergence theorem that the sequence $\{Y^k\}_k$ converges and $\sum_{k=1}^{\infty} X^k < \infty$; see Proposition 4.2 in Bertsekas and Tsitsiklis (1996). The third part uses these results to complete the proof.

Part 1. To capture the cases where the assumption of Lemma 4 holds, we define the event A_j^k as

$$A_j^k = \left\{ y_{j+1}^k \in \left(L_{j+1}^* - \frac{1}{4}, U_{j+1}^* + \frac{1}{4} \right), y_{j+2}^k \in \left(L_{j+2}^* - \frac{1}{4}, U_{j+2}^* + \frac{1}{4} \right), \dots, y_n^k \in \left(L_n^* - \frac{1}{4}, U_n^* + \frac{1}{4} \right) \right. \\ \left. \text{and } \alpha_{j+1}^k \in \left[0, \frac{1}{4M} \right], \alpha_{j+2}^k \in \left[0, \frac{1}{4M} \right], \dots, \alpha_n^k \in \left[0, \frac{1}{4M} \right] \right\}.$$

Using $\mathbf{1}(\cdot)$ to denote the indicator function, since $|y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})| \leq c$, Lemma 4 implies that

$$\begin{aligned} & |y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})|^2 \\ & \leq \mathbf{1}(A_j^k) |y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) - \mathcal{C}_j(y_j^k)|^2 + [1 - \mathbf{1}(A_j^k)] c^2 \\ & \leq |y_j^k - \mathcal{C}_j(y_j^k)|^2 - \mathbf{1}(A_j^k) 2 \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) [\mathcal{C}_j(y_j^k) - y_j^k] + [\alpha_j^k]^2 M^2 + [1 - \mathbf{1}(A_j^k)] c^2, \end{aligned}$$

where the last inequality follows from Lemma 1. Taking the conditional expectations and noting that $\mathbf{1}(A_j^k)$ is \mathcal{F}^k -measurable, we obtain

$$\begin{aligned} \mathbb{E}\{|y_j^{k+1} - \mathcal{C}_j(y_j^{k+1})|^2 | \mathcal{F}^k\} & \leq |y_j^k - \mathcal{C}_j(y_j^k)|^2 \\ & - \mathbf{1}(A_j^k) 2 \alpha_j^k [\mathcal{C}_j(y_j^k) - y_j^k] \mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} + [\alpha_j^k]^2 M^2 + [1 - \mathbf{1}(A_j^k)] c^2. \end{aligned}$$

If we let $Y^k = |y_j^k - \mathcal{C}_j(y_j^k)|^2$, $X^k = \mathbf{1}(A_j^k) 2\alpha_j^k [\mathcal{C}_j(y_j^k) - y_j^k] \mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\}$ and $Z^k = [\alpha_j^k]^2 M^2 + [1 - \mathbf{1}(A_j^k)] c^2$, then the inequality above is of the form $\mathbb{E}\{Y^{k+1} | \mathcal{F}^k\} \leq Y^k - X^k + Z^k$.

Part 2. Clearly, $\{Y^k\}_k$ and $\{Z^k\}_k$ are positive and $\{X^k\}_k$, $\{Y^k\}_k$ and $\{Z^k\}_k$ are \mathcal{F}^k -measurable. We now show that $\{X^k\}_k$ is positive. If $\mathbf{1}(A_j^k) = 0$, then we have $X^k = 0$. If, on the other hand, we have $\mathbf{1}(A_j^k) = 1$, then we obtain $\mathbb{E}\{\rho_{j+1}^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} \in \partial v_{j+1}(y_j^k, D_{j+1}^k)$ by Lemma 2 and the definition of the event A_j^k . That is, if $\mathbf{1}(A_j^k) = 1$, then we have $\mathbb{E}\{\rho_{j+1}^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} = \dot{v}_{j+1}(y_j^k, D_{j+1}^k)$. Therefore, by (13), if $\mathbf{1}(A_j^k) = 1$, then we have

$$\begin{aligned} \mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} &= -r_j + \mathbb{E}\{\mathbb{E}\{\rho_{j+1}^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k, D_{j+1}^k\} | \mathcal{F}^k\} \\ &= -r_j + \mathbb{E}\{\dot{v}_{j+1}(y_j^k, D_{j+1}^k) | \mathcal{F}^k\}, \end{aligned}$$

in which case (4) implies that $\mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\}$ is a subgradient of $f_j(\cdot)$ at y_j^k . Consequently, we have $\mathbf{1}(A_j^k) [\mathcal{C}_j(y_j^k) - y_j^k] \mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} \geq \mathbf{1}(A_j^k) [f_j(\mathcal{C}_j(y_j^k)) - f_j(y_j^k)] \geq 0$, where the last inequality follows from the fact that $\mathcal{C}_j(y_j^k) \in \mathcal{Y}_j^*$ and (5). Therefore, $\{X^k\}_k$ is positive.

We now show that $\sum_{k=1}^{\infty} Z^k < \infty$. Noting the induction assumption that $\lim_{k \rightarrow \infty} |y_{j+1}^k - \mathcal{C}_{j+1}(y_{j+1}^k)| = 0$, $\lim_{k \rightarrow \infty} |y_{j+2}^k - \mathcal{C}_{j+2}(y_{j+2}^k)| = 0, \dots, \lim_{k \rightarrow \infty} |y_n^k - \mathcal{C}_n(y_n^k)| = 0$ and the fact that $\lim_{k \rightarrow \infty} \alpha_j^k = 0$ for all $j = 1, \dots, n$, there exists a finite iteration counter K such that $\mathbf{1}(A_j^k) = 1$ for all $k = K, K+1, \dots$. Therefore, we have $\sum_{k=1}^{\infty} Z^k \leq \sum_{k=1}^{\infty} [\alpha_j^k]^2 M^2 + Kc^2 < \infty$.

Part 3. Since we have $\mathbb{E}\{Y^{k+1} | \mathcal{F}^k\} \leq Y^k - X^k + Z^k$ and $\sum_{k=1}^{\infty} Z^k < \infty$, we can conclude by the supermartingale convergence theorem that the sequence $\{Y^k\}_k = \{|y_j^k - \mathcal{C}_j(y_j^k)|^2\}_k$ converges and we have $\sum_{k=1}^{\infty} X^k < \infty$. By the discussion in Part 2, for all $k = K, K+1, \dots$, we have $\mathbf{1}(A_j^k) = 1$ and

$$X^k = 2\alpha_j^k [\mathcal{C}_j(y_j^k) - y_j^k] \mathbb{E}\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) | \mathcal{F}^k\} \geq 2\alpha_j^k |f_j(\mathcal{C}_j(y_j^k)) - f_j(y_j^k)|.$$

Therefore, we have $\sum_{k=K}^{\infty} \alpha_j^k |f_j(\mathcal{C}_j(y_j^k)) - f_j(y_j^k)| \leq \sum_{k=1}^{\infty} X^k < \infty$, which, together with the fact that $\sum_{k=1}^{\infty} \alpha_j^k = \infty$, implies that $\liminf_{k \rightarrow \infty} |f_j(\mathcal{C}_j(y_j^k)) - f_j(y_j^k)| = 0$. Consequently, there exists a subsequence $\{\hat{y}_j^k\}_k$ of $\{y_j^k\}_k$ such that $\lim_{k \rightarrow \infty} |f_j(\mathcal{C}_j(\hat{y}_j^k)) - f_j(\hat{y}_j^k)| = 0$. Since the sequence $\{\hat{y}_j^k\}_k$ takes values in the bounded interval $[0, c]$, we can take a further subsequence $\{\tilde{y}_j^k\}_k$ of $\{\hat{y}_j^k\}_k$ such that $\lim_{k \rightarrow \infty} \tilde{y}_j^k = \tilde{y}_j$ for some $\tilde{y}_j \in [0, c]$.

Noting the definition of $\mathcal{C}_j(\cdot)$ and letting $F_j^* = \max_{0 \leq y_j \leq c} f_j(y_j)$, we clearly have $f_j(\mathcal{C}_j(\tilde{y}_j^k)) = F_j^*$ for all $k = 1, 2, \dots$. Therefore, by the fact that $\lim_{k \rightarrow \infty} |f_j(\mathcal{C}_j(\tilde{y}_j^k)) - f_j(\tilde{y}_j^k)| = 0$, we have $\lim_{k \rightarrow \infty} f_j(\tilde{y}_j^k) = F_j^*$. On the other hand, by the continuity of $f_j(\cdot)$ and the fact that $\lim_{k \rightarrow \infty} \tilde{y}_j^k = \tilde{y}_j$, we have $\lim_{k \rightarrow \infty} f_j(\tilde{y}_j^k) = f_j(\tilde{y}_j)$. From the last two statements, we obtain $f_j(\tilde{y}_j) = F_j^*$ so that $\tilde{y}_j \in \mathcal{Y}_j^*$, which implies that $|\tilde{y}_j^k - \mathcal{C}_j(\tilde{y}_j^k)| \leq |\tilde{y}_j^k - \tilde{y}_j|$ for all $k = 1, 2, \dots$. Therefore, since $\{|\tilde{y}_j^k - \tilde{y}_j|^2\}_k$ converges to zero, $\{|\tilde{y}_j^k - \mathcal{C}_j(\tilde{y}_j^k)|^2\}_k$ also converges to zero. Recalling that the whole sequence $\{Y^k\}_k = \{|y_j^k - \mathcal{C}_j(y_j^k)|^2\}_k$ converges, we obtain $\lim_{k \rightarrow \infty} |y_j^k - \mathcal{C}_j(y_j^k)|^2 = 0$. \square

A simple corollary to Proposition 5 is that there exists a finite iteration number \bar{K} w.p.1 such that we have $L_j^* - \frac{1}{2} < y_j^k < U_j^* + \frac{1}{2}$ for all $j = 1, \dots, n$, $k = \bar{K}, \bar{K} + 1, \dots$. Therefore, we have $\mathcal{O}(y_j^k) \in \mathcal{Y}_j^*$ for all $j = 1, \dots, n$, $k = \bar{K}, \bar{K} + 1, \dots$ and the policy that uses $\{\mathcal{O}(y_j^k) : j = 1, \dots, n\}$ as the protection levels is optimal w.p.1 after a finite number of iterations.

5 CENSORED DEMANDS

Demand censorship refers to the situation where we can observe the number of seats sold to a fare class, but not the actual amount of demand from a fare class. In this case, our demand observations are “truncated” when the amount of demand from a fare class exceeds the number of seats that we make available for sale. In this section, we show that we can compute the step direction in (13) as long as we can observe the number of seats sold to a fare class and whether any demand from this fare class is turned down. In particular, if a portion of the demand from a fare class cannot be satisfied, then we do not need to know the exact quantity of the demand that is turned down. This result implies that Algorithm 1 remains applicable when the demand information is censored.

If the demand information is censored, then we do not observe the demand random variables $\{D_j^k : j = 1, \dots, n\}$ in Step 2 of Algorithm 1. Instead, we simulate the behavior of the policy characterized by the protection levels $\{\mathcal{O}(y_j^k) : j = 1, \dots, n\}$, and observe the number of seats sold to different fare classes and whether any demand from different fare classes is turned down. In this case, Step 2 of Algorithm 1 has to be replaced with the following steps.

Step 2.a. Set the initial capacity x_1^k to c and set $j = 1$.

Step 2.b. Make $[x_j^k - \mathcal{O}(y_j^k)]^+$ seats available for sale to fare class j .

Step 2.c. Observe the number of seats sold to fare class j as $P_j^k = \min\{[x_j^k - \mathcal{O}(y_j^k)]^+, D_j^k\}$. Observe whether any demand from fare class j is turned down as

$$B_j^k = \mathbf{1}(D_j^k > [x_j^k - \mathcal{O}(y_j^k)]^+).$$

Compute the capacity just before making the decisions for fare class $j + 1$ as $x_{j+1}^k = x_j^k - P_j^k$.

Step 2.d. If $j < n$, then increase j by 1 and go to Step 2.b.

Step 2.e. For all $j = 1, \dots, n$, set

$$y_j^{k+1} = \max \left\{ \min \left\{ [y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)]^+, c \right\}, \mathcal{O}(y_{j+1}^{k+1}) \right\}.$$

When the demand information is censored, we only have access to $\{P_j^k : j = 1, \dots, n\}$ and $\{B_j^k : j = 1, \dots, n\}$, but not to the demand random variables themselves. The next proposition shows that this information is adequate to compute the step direction in (13).

Proposition 6 *Having access to $\{y_j^k : j = 1, \dots, n\}$, $\{x_j^k : j = 1, \dots, n\}$, $\{P_j^k : j = 1, \dots, n\}$ and $\{B_j^k : j = 1, \dots, n\}$ is adequate to compute $s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$ for all $j = 1, \dots, n$.*

Proof It is possible to show the final result by induction over the fare classes, but we use a constructive proof, which is easier to follow and shows the computations involved more clearly. We begin by using induction over the fare classes to show that

$$x_j^k \geq \mathcal{O}(y_{j-1}^k) \geq \mathcal{O}(y_j^k) \tag{18}$$

for all $j = 2, \dots, n$. We only focus on the first inequality above since Algorithm 1 ensures that the rounded estimates of the optimal protection levels are nested and we always have $\mathcal{O}(y_1^k) \geq \mathcal{O}(y_2^k) \geq \dots \geq$

$\mathcal{O}(y_n^k)$ for all $k = 1, 2, \dots$. The result in (18) is easy to show for the second fare class. Assuming that the result holds for fare class j , we have $x_{j+1}^k = x_j^k - \min\{x_j^k - \mathcal{O}(y_j^k)\}^+, D_j^k\} = x_j^k - \min\{x_j^k - \mathcal{O}(y_j^k), D_j^k\} = \max\{\mathcal{O}(y_j^k), x_j^k - D_j^k\} \geq \mathcal{O}(y_j^k) \geq \mathcal{O}(y_{j+1}^k)$, where the last inequality uses the fact that Algorithm 1 ensures that the rounded estimates of the optimal protection levels are nested. Therefore, (18) holds.

We now focus on computing $s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. We note that by (13), this requires computing $\rho_{j+1}^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$. We consider two cases.

Case 1. Assume that $B_{j+1}^k = 1$. In this case, we deduce that $D_{j+1}^k > [x_{j+1}^k - \mathcal{O}(y_{j+1}^k)]^+$. Therefore, we have $D_{j+1}^k \geq [x_{j+1}^k - \mathcal{O}(y_{j+1}^k)]^+ + 1 \geq \mathcal{O}(y_j^k) - \mathcal{O}(y_{j+1}^k) + 1 \geq y_j^k - \mathcal{O}(y_{j+1}^k)$, where the second inequality follows from (18). This chain of inequalities and (14) imply that $y_j^k - D_{j+1}^k \leq \mathcal{O}(y_{j+1}^k) \leq y_j^k$ and we obtain $\rho_{j+1}^k(y_j^k, D_{j+1}^k, D_{j+2}^k, \dots, D_n^k) = r_{j+1}$ by (12).

Case 2. Assume that $B_{j+1}^k = 0$. In this case, we deduce that $D_{j+1}^k \leq [x_{j+1}^k - \mathcal{O}(y_{j+1}^k)]^+$ so that $D_{j+1}^k = \min\{[x_{j+1}^k - \mathcal{O}(y_{j+1}^k)]^+, D_{j+1}^k\} = P_{j+1}^k$. Therefore, since we have access to the value of P_{j+1}^k , we also have access to the value of D_{j+1}^k . We consider two subcases.

Case 2.a. Assume that $D_{j+1}^k \geq y_j^k - \mathcal{O}(y_{j+1}^k)$. We have $y_j^k - D_{j+1}^k \leq \mathcal{O}(y_{j+1}^k) \leq y_j^k$, where the second inequality follows from (14) and we obtain $\rho_{j+1}^k(y_j^k, D_{j+1}^k, D_{j+2}^k, \dots, D_n^k) = r_{j+1}$ by (12).

Case 2.b. Assume that $D_{j+1}^k < y_j^k - \mathcal{O}(y_{j+1}^k)$. We have $\rho_{j+1}^k(y_j^k, D_{j+1}^k, D_{j+2}^k, \dots, D_n^k) = \rho_{j+2}^k(y_j^k - D_{j+1}^k, D_{j+2}^k, \dots, D_n^k)$ by (12).

Therefore, if Cases 1 or 2.a holds, then we are done. Otherwise, it remains to compute $\rho_{j+2}^k(y_j^k - D_{j+1}^k, D_{j+2}^k, \dots, D_n^k)$ for a known value of $y_j^k - D_{j+1}^k$. The result follows by continuing in the same fashion for the subsequent fare classes. \square

Our view of demand censorship in this section is more restrictive than the view of demand censorship in the existing literature. In the existing literature, if the demand information is censored, then we observe the number of seats sold to a fare class, but not necessarily whether any demand from this fare class is turned down. In this section, however, we assume that we observe both the number of seats sold to a fare class and whether any demand from this fare class is turned down. This discrepancy is purely due to the fact that we work with discrete demand distributions. When the demand distributions are continuous, the probability that the demand from a fare class is equal to the number of seats that we make available for sale is zero. In this case, if the number of seats that we make available for sale is equal to the number of seats that we sell, then we can safely assume that a portion of the demand from this fare class is not satisfied. However, it is not possible to assume the same thing when the demand distributions are discrete. Indeed, many discussions of demand censorship with discrete demand distributions assume that we can observe both the number of seats sold to a fare class and whether any demand from this fare class is turned down; see Huh, Levi, Rusmevichientong and Orlin (2008). In this sense, our view of demand censorship is in alignment with these discussions. In the next section, we give a heuristic modification of our stochastic approximation method that only requires the ability to observe the number of seats sold to a fare class. This heuristic modification does not have a convergence guarantee, but it still performs remarkably well.

From a purely arithmetical standpoint, it is straightforward to see that having access to $\{B_j^k : j = 1, \dots, n\}$ is necessary to compute $s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)$ and having access to only $\{y_j^k : j = 1, \dots, n\}$, $\{x_j^k : j = 1, \dots, n\}$ and $\{P_j^k : j = 1, \dots, n\}$ is not adequate. To illustrate, we use a numerical example with $n = 3$, $c = 4$ and

$$\begin{aligned} y_1^k &= 3.2, y_2^k = 2.1, y_3^k = 0 \\ D_1^k &= 1, D_2^k = 1, D_3^k = 2. \end{aligned}$$

By Steps 2.a-2.d above, we have $x_1^k = 4$, $P_1^k = \min\{[x_1^k - \mathcal{O}(y_1^k)]^+, D_1^k\} = 1$, $x_2^k = 4 - 1 = 3$, $P_2^k = \min\{[x_2^k - \mathcal{O}(y_2^k)]^+, D_2^k\} = 1$, $x_3^k = 3 - 1 = 2$ and $P_3^k = \min\{[x_3^k - \mathcal{O}(y_3^k)]^+, D_3^k\} = 2$. By (13), computing $s_1^k(y_1^k, D_2^k, D_3^k)$ requires computing $\rho_2^k(y_1^k, D_2^k, D_3^k)$ and we have

$$\rho_2^k(y_1^k, D_2^k, D_3^k) = \begin{cases} \rho_3^k(3.2 - D_2^k, D_3^k) & \text{if } \mathcal{O}(2.1) < 3.2 - D_2^k \\ r_2 & \text{if } 3.2 - D_2^k \leq \mathcal{O}(2.1) \leq 3.2 \end{cases}$$

by (12). Therefore, to compute $s_1^k(y_1^k, D_2^k, D_3^k)$, we need to know whether $D_2^k < 1.2$ or $D_2^k \geq 1.2$. However, if we only have access to $\{y_1^k, y_2^k, y_3^k\}$, $\{x_1^k, x_2^k, x_3^k\}$ and $\{P_1^k, P_2^k, P_3^k\}$, then we know that $1 = P_2^k = \min\{[x_2^k - \mathcal{O}(y_2^k)]^+, D_2^k\} = \min\{1, D_2^k\} \leq D_2^k$, but not whether $D_2^k < 1.2$ or $D_2^k \geq 1.2$.

6 COMPUTATIONAL EXPERIMENTS

In this section, we compare the performance of our stochastic approximation method with that of the stochastic approximation method proposed by van Ryzin and McGill (2000).

6.1 EXPERIMENTAL SETUP

Our test problems are based on the test problems in van Ryzin and McGill (2000), but we introduce some variety by working with 4, 8 or 12 fare classes. We assume that the demand from each fare class is normally distributed. To satisfy our assumptions, we discretize the demand random variables by rounding them to the nearest integer. Table 1 gives the revenues associated with the fare classes, along with the means and standard deviations of the demand random variables. For the test problems with 4 fare classes, we use $c \in \{124, 164\}$, in which case the total expected demand is 25% more or 5% less than the initial capacity. Similarly, for the test problems with 8 and 12 fare classes, we respectively use $c \in \{260, 344\}$ and $c \in \{409, 541\}$.

We choose the initial protection levels in three different ways. In particular, we use

$$y_j^1 = \frac{r_{j+1} + \dots + r_n}{r_1 + \dots + r_n} c \quad (19)$$

$$y_j^1 = \frac{\mu_{j+1} + \dots + \mu_n}{\mu_1 + \dots + \mu_n} c \quad (20)$$

$$y_j^1 = \frac{\mu_{j+1} r_{j+1} + \dots + \mu_n r_n}{\mu_1 r_1 + \dots + \mu_n r_n} c, \quad (21)$$

where μ_j is the mean of the demand random variable for fare class j . These initial protection levels are motivated by the observation that if the revenue associated with a fare class is large and the demand

from a fare class is likely to be large, then we should protect a large number of seats for this fare class. We note that (19) uses the revenues, (20) uses the means of the demand random variables and (21) uses both the revenues and means of the demand random variables. Accordingly, we refer to the sets of initial protection levels computed by (19), (20) and (21) respectively as Y_R^1 , Y_M^1 and Y_{RM}^1 . After some experimentation, we decided to use the step size parameter $\alpha_j^k = (n - j + 1) \frac{200}{r_n(10+k)}$ in our stochastic approximation method. We note that this choice of step size parameters results in a bit more aggressive updates for the protection levels for the cheaper fare classes.

We label our test cases by $(n, \kappa, Y^1) \in \{4, 8, 12\} \times \{0.95, 1.25\} \times \{Y_R^1, Y_M^1, Y_{RM}^1\}$, where n is the number of fare classes, κ is the ratio of the total expected demand to the initial capacity and Y^1 is the initial protection levels. This provides 18 test cases to consider in our experimental setup. We run all of our benchmark strategies for 100 iterations on 25 sample paths and present the average results over the 25 sample paths. We use common random numbers when comparing the performances of different benchmark strategies.

6.2 BENCHMARK STRATEGIES

We compare the performances of three benchmark strategies. The first benchmark strategy is Algorithm 1 described in Section 3. By the discussion in Section 5, this benchmark strategy requires that we can observe both the number of seats sold to different fare classes and whether any demand from different fare classes is turned down. We refer to this benchmark strategy as A1.

The second benchmark strategy is a heuristic modification of A1 and it only requires that we can observe the number of seats sold to different fare classes. This heuristic modification is obtained by replacing Step 2 of Algorithm 1 with the following steps.

Step 2.a. Set the initial capacity x_1^k to c and set $j = 1$.

Step 2.b. Make $[x_j^k - \mathcal{O}(y_j^k)]^+$ seats available for sale to fare class j .

Step 2.c. Observe the number of seats sold to fare class j as $P_j^k = \min\{[x_j^k - \mathcal{O}(y_j^k)]^+, D_j^k\}$. Heuristically compute whether any demand from fare class j is turned down as

$$\tilde{B}_j^k = \mathbf{1}(P_j^k = [x_j^k - \mathcal{O}(y_j^k)]^+).$$

Compute the capacity just before making the decisions for fare class $j + 1$ as $x_{j+1}^k = x_j^k - P_j^k$.

Step 2.d. If $j < n$, then increase j by 1 and go to Step 2.b.

Step 2.e. Compute $\{s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k) : j = 1, \dots, n\}$ under the assumption that the number of seats sold to different fare classes is given by $\{P_j^k : j = 1, \dots, n\}$ and whether any demand from different fare classes is turned down is given by $\{\tilde{B}_j^k : j = 1, \dots, n\}$. For all $j = 1, \dots, n$, set

$$y_j^{k+1} = \max \left\{ \min \left\{ [y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)]^+, c \right\}, \mathcal{O}(y_{j+1}^{k+1}) \right\}.$$

In Step 2.c, if we have $P_j^k < [x_j^k - \mathcal{O}(y_j^k)]^+$, then no demand from fare class j is turned down and we set $\tilde{B}_j^k = 0$. On the other hand, if we have $P_j^k = [x_j^k - \mathcal{O}(y_j^k)]^+$, then either a portion of the demand from fare class j is turned down or the demand from fare class j is equal to the number of seats that

we make available for sale. In this case, we heuristically assume that a portion of the demand from fare class j is turned down and set $\tilde{B}_j^k = 1$, irrespective of whether a portion of the demand from fare class j is actually turned down or not. We refer to this benchmark strategy as A2.

The third benchmark strategy is the stochastic approximation method proposed by van Ryzin and McGill (2000). We refer to this benchmark strategy as VM.

6.3 COMPUTATIONAL RESULTS

In Figure 2, we begin by comparing the performances of A1 and A2 on a number of test cases. We give the label for each test case in the charts. The dashed and solid data series in the charts respectively plot the total expected revenues corresponding to the protection levels obtained by A1 and A2 as a function of the iteration counter. All of the total expected revenues are given in terms of the percentage of the optimal total expected revenues. Figure 2 indicates that the performance of A2 is indistinguishable from that of A1. The data series corresponding to A1 and A2 in many charts essentially coincide and it is difficult to see that there are actually two data series. This observation uniformly holds for all of the test cases that we worked with. Since A2 only requires that we can observe the number of seats sold to different fare classes, we compare A2 with VM in the subsequent discussion.

Figures 3, 4 and 5 respectively show the results for the test cases with 4, 8 and 12 fare classes. In these figures, the charts on the left, in the middle and on the right respectively correspond to the test cases where the initial protection levels are Y_R^1 , Y_M^1 and Y_{RM}^1 . The top rows correspond to the test cases where $\kappa = 0.95$, whereas the bottom rows correspond to the test cases where $\kappa = 1.25$. The thin and thick data series in the charts respectively plot the total expected revenues corresponding to the protection levels obtained by A2 and VM as a function of the iteration counter.

We begin by considering the charts in the bottom rows of Figures 3, 4 and 5. These charts correspond to the test cases where $\kappa = 1.25$ and they indicate that A2 tends to perform better than VM when the total expected demand exceeds the capacity by a significant margin. We observe that the performance gaps between A2 and VM are substantial when the initial protection levels are Y_R^1 . If the initial protection levels are Y_{RM}^1 , then A2 performs better than VM by a small but consistent margin. The performances of A2 and VM are quite similar when the initial protection levels are Y_M^1 . We note that the protection levels Y_R^1 are not very good as the total expected revenues obtained by these protection levels are less than 90% of the optimal total expected revenues. This can be seen by checking the first data points in the charts. On the other hand, the total expected revenues obtained by Y_{RM}^1 and Y_M^1 are respectively about 95% and 98% of the optimal total expected revenues. Therefore, A2 appears to perform better than VM especially when the total expected demand exceeds the capacity by a significant margin and the initial protection levels are not very close to the optimal protection levels.

The charts in the top rows of Figures 3, 4 and 5 indicate that the performance of A2 is comparable to that of VM when the total expected demand is below the capacity. If the initial protection levels are Y_M^1 or Y_{RM}^1 , then VM tends to perform slightly better than A2, whereas if the initial protection levels are Y_R^1 , then A2 tends to perform slightly better than VM. These charts are consistent with the earlier

observation that A2 appears to perform better than VM when the initial protection levels are not close to the optimal protection levels. We also note that the performance gap between A2 and VM can still be significant as shown in the top row of Figure 5 with the initial protection levels Y_R^1 .

In the bottom rows of Figures 4 and 5, it appears that VM cannot obtain a good set of protection levels within a reasonable number of iterations when the initial protection levels are Y_R^1 . To make sure that VM does not prematurely stop making progress, we run VM for 1,000 iterations for these two test cases and Figure 6 plots the total expected revenues corresponding to the protection levels obtained by VM as a function of the iteration counter. It turns out that VM eventually obtains good protection levels but this may take a large number of iterations.

Our computational experiments indicate that the performances of A1 and A2 are at least comparable to that of VM and there are situations with tight capacities where A1 and A2 perform significantly better than VM. Furthermore, A1 and A2 tend to perform better than VM when the initial protection levels are not very close to the optimal protection levels. Therefore, A1 and A2 appear to be reasonable substitutes for VM. We also note that A1 and A2 provide stochastic subgradients of the value functions with respect to the seat availability, which may be useful when making tactical decisions such as setting the capacity of the flight leg.

7 CONCLUSIONS AND EXTENSIONS

In this paper, we developed a stochastic approximation method to compute the optimal protection levels for the seat allocation problem under the assumption that the demand distributions are discrete. Although the problem that we consider is nonsmooth and the total expected revenue is not concave when viewed as a function of the protection levels, we were able to show that the iterates of our method converge to a set of optimal protection levels. We also showed that our method remains applicable when we can observe the number of seats sold to a fare class and whether any demand from this fare class is turned down. Computational experiments demonstrated that our method can be advantageous when the total expected demand exceeds the capacity by a significant margin and we do not have good a priori estimates of the optimal protection levels.

It is possible to extend our method to cover the case where the revenues do not necessarily satisfy $r_1 \leq r_2 \leq \dots \leq r_n$. If we do not have $r_1 \leq r_2 \leq \dots \leq r_n$, then the optimal protection levels are not necessarily nested and it is not guaranteed that $y_1^* \geq y_2^* \geq \dots \geq y_n^*$. However, if we modify (14) in Step 2 of Algorithm 1 as

$$y_j^{k+1} = \min \{ [y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \dots, D_n^k)]^+, c \}, \quad (22)$$

then we can use the same argument in Section 4 to show that the iterates of Algorithm 1 still converge to a set of optimal protection levels w.p.1. This variant of Algorithm 1 works when we do not necessarily have $r_1 \leq r_2 \leq \dots \leq r_n$.

For the variant of Algorithm 1 in (22), if we assume that the step size parameter α_j^k is the same for all fare classes and it is equal to α^k , then we can follow the same line of reasoning in van Ryzin and

McGill (2000) to show the convergence rate

$$\mathbb{E}\{|y_j^k - C_j(y_j^k)|^2\} \leq \Lambda_j [\alpha^k]^\beta \quad (23)$$

for all $j = 1, \dots, n$, where β and Λ_j are finite scalars, $\Lambda_n = 0$ and Λ_j increases exponentially fast as j decreases from n to 1. In contrast, van Ryzin and McGill (2000) show the convergence rate

$$\mathbb{E}\{|y_j^k - C_j(y_j^k)|^2\} \leq \Lambda_j [\alpha^k]^{\beta/2^{n-j}},$$

which is slower than the convergence rate in (23).

Finally, we note that the proof of Proposition 6 depends on the fact that the rounded estimates of the optimal protection levels are nested. Therefore, the variant of Algorithm 1 in (22) does not work under the assumption that the demand information is censored by the seat availability. Developing stochastic approximation methods that work when the demand information is censored and we do not necessarily have $r_1 \leq r_2 \leq \dots \leq r_n$ is an important area for further research.

REFERENCES

- 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.
- 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.
- Brumelle, S. L. and McGill, J. I. (1993), ‘Airline seat allocation with multiple nested fare classes’, *Operations Research* **41**, 127–137.
- Curry, R. E. (1990), ‘Optimal airline seat allocation with fare classes nested by origins and destinations’, *Transportation Science* **24**, 193–204.
- Ermoliev, Y. (1988), Stochastic quasigradient methods, in Y. Ermoliev and R. Wets, eds, ‘*Numerical Techniques for Stochastic Optimization*’, Springer-Verlag, Berlin.
- Fu, M. (1994), ‘Sample path derivatives for (s, S) inventory systems’, *Operations Research* **42**(2), 351–363.
- Glasserman, P. and Tayur, S. (1995), ‘Sensitivity analysis for base-stock levels in multiechelon production-inventory systems’, *Management Science* **41**(2), 263–281.
- Huh, W. T., Levi, R., Rusmevichientong, P. and Orlin, J. B. (2008), Adaptive data-driven inventory control policies based on Kaplan-Meier estimator, Technical report, School of Operations Research and Information Engineering, Cornell University.
- Huh, W. T. and Rusmevichientong, P. (2006), Adaptive capacity allocation with censored demand data: Application of concave umbrella functions, Technical report, Cornell University, School of Operations Research and Information Engineering.
- Karaesmen, I. and van Ryzin, G. (2004), ‘Overbooking with substitutable inventory classes’, *Operations Research* **52**(1), 83–104.
- Kunnumkal, S. and Topaloglu, H. (2007), ‘Using stochastic approximation algorithms to compute optimal base-stock levels in inventory control problems’, *Operations Research* (to appear).

- Kushner, H. J. and Clark, D. S. (1978), *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, Berlin.
- Lautenbacher, C. J. and Stidham, S. (1999), ‘The underlying markov decision process in the single-leg airline yield management problem’, *Transportation Science* **33**, 136–146.
- 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.
- Lee, T. and Hersh, M. (1993), ‘A model for dynamic airline seat inventory control with multiple seat bookings’, *Transportation Science* **27**, 252–265.
- Littlewood, K. (1972), Forecasting and control of passengers, in ‘Proceedings of the 12th AGIFORS’, New York, pp. 95–117.
- Mahajan, S. and van Ryzin, G. (2001), ‘Stocking retail assortments under dynamic customer substitution’, *Operations Research* **49**(3), 334–351.
- Robinson, L. (1995), ‘Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes’, *Operations Research* **43**, 252–263.
- Talluri, K. T. and van Ryzin, G. J. (2004), *The Theory and Practice of Revenue Management*, Kluwer Academic Publishers.
- Topaloglu, H. (2007), ‘A stochastic approximation method to compute bid prices in network revenue management problems’, *INFORMS Journal on Computing* (to appear).
- van Ryzin, G. and McGill, J. (2000), ‘Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels’, *Management Science* **46**(6), 760–775.
- van Ryzin, G. and Vulcano, G. (2004), Computing virtual nesting controls for network revenue management under customer choice behavior, Technical Report DRO-2004-09, Columbia Business School.
- van Ryzin, G. and Vulcano, G. (2006), ‘Simulation-based optimization of virtual nesting controls for network revenue management’, *Operations Research* (to appear).
- Wollmer, R. D. (1992), ‘An airline seat management model for a single leg route when lower fare classes book first’, *Operations Research* **40**, 26–37.
- Zinkevich, M. (2003), Online convex programming and generalized infinitesimal gradient ascent, in ‘Proceedings of the Twentieth International Conference on Machine Learning (ICML-2003)’, Washington, D.C.




Figure 1: Total expected revenue $R(c, y_1, y_2, 0)$ as a function of the protection levels (y_1, y_2) . Problem parameters are $c = 20$, $r_1 = 8$, $r_2 = 10$, $r_3 = 14$, $D_1 \sim \text{Uniform}\{0, \dots, 20\}$, $D_2 \sim \text{Uniform}\{0, \dots, 15\}$, $D_3 \sim \text{Uniform}\{0, \dots, 20\}$.

fare class	1	2	3	4
revenue	1050	567	527	350
mean	17.3	45.1	73.6	19.8
std. dev.	5.8	15	17.4	6.6

fare class	1	2	3	4	5	6	7	8
revenue	1155	1050	623.7	579.7	567.7	527	385	350
mean	19	17.3	49.6	81	45.1	73.6	21.8	19.8
std. dev.	6.4	5.8	16.5	19.1	15	17.4	7.3	6.6

fare class	1	2	3	4	5	6	7	8	9	10	11	12
revenue	1260	1155	1050	680.4	632.4	623.7	579.7	567	527	420	385	350
mean	20.8	19	17.3	54.1	88.3	49.6	81	45.1	73.6	23.8	21.8	19.8
std. dev.	7	6.4	5.8	18	20.9	16.5	19.1	15	17.4	7.9	7.3	6.6

Table 1: Parameters of our test problems.




Figure 2: Comparison of A1 and A2.




Figure 3: Computational results for the test cases with 4 fare classes.

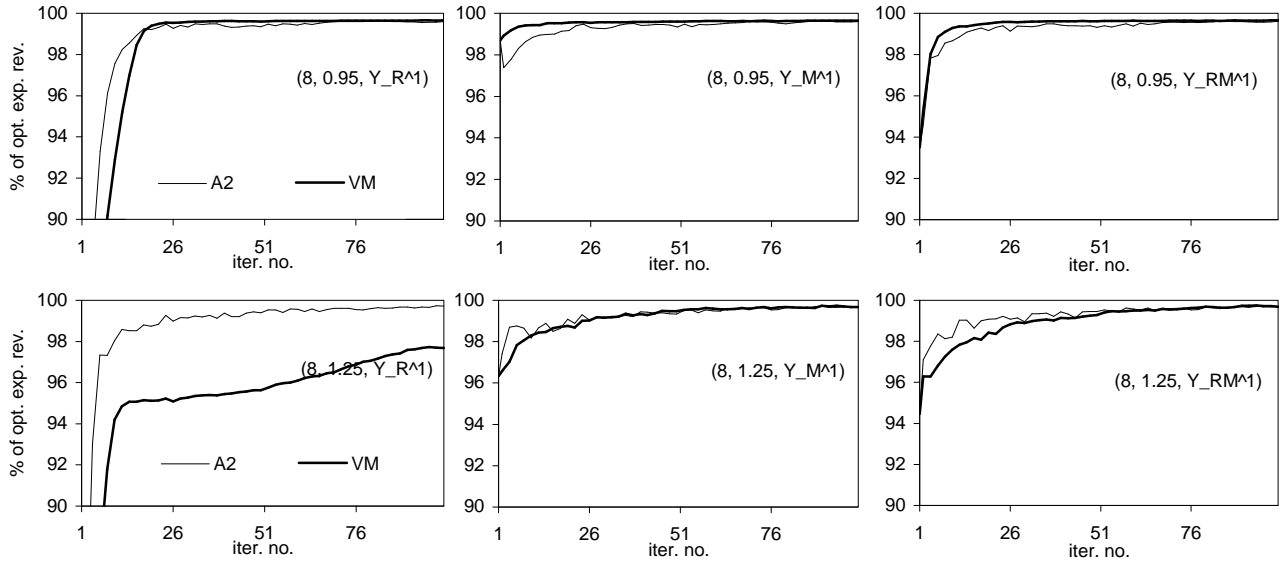


Figure 4: Computational results for the test cases with 8 fare classes.




Figure 5: Computational results for the test cases with 12 fare classes.




Figure 6: Performance of VM on test cases $(8, 1.25, Y_R^1)$ and $(12, 1.25, Y_R^1)$.

APPENDIX: OBTAINING A SUBGRADIENT OF $v_j(\cdot, d_j)$

In this section, we use induction over the fare classes to show that the recursion in (11) gives a subgradient of $v_j(\cdot, d_j)$ at x_j . It is easy to show the result for fare class n . Assuming that the result holds for fare class $j + 1$ and $\dot{v}_j(x_j, d_j)$ is defined as in (11), we show that

$$v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) \leq \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j]$$

for all $x_j, \tilde{x}_j \in [0, c]$. Since the roles of x_j and \tilde{x}_j are interchangeable, we consider six cases.

Case 1. Assume that $y_j^* < x_j - d_j$ and $y_j^* < \tilde{x}_j - d_j$. We have $v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) = \mathbb{E}\{v_{j+1}(\tilde{x}_j - d_j, D_{j+1})\} - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [\tilde{x}_j - x_j] = \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j]$, where the first equality is by (10), the inequality is by the induction assumption that $\dot{v}_{j+1}(x_j - d_j, D_{j+1})$ is a subgradient of $v_{j+1}(\cdot, D_{j+1})$ at $x_j - d_j$ and the second equality is by (11).

Case 2. Assume that $y_j^* < x_j - d_j$ and $\tilde{x}_j - d_j \leq y_j^* \leq \tilde{x}_j$. By the induction assumption that $\dot{v}_{j+1}(x_j - d_j, D_{j+1})$ is a subgradient of $v_{j+1}(\cdot, D_{j+1})$ at $x_j - d_j$, we have $\mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [y_j^* - x_j + d_j]$. Since y_j^* is a maximizer of $f_j(\cdot)$ over $[0, c]$ and $y_j^* < x_j - d_j$, we also have $\mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} \leq r_j$ by (4). Noting that $\tilde{x}_j - d_j - y_j^* \leq 0$, we obtain $r_j [\tilde{x}_j - d_j - y_j^*] \leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [\tilde{x}_j - d_j - y_j^*]$. By (10), we have

$$\begin{aligned} v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) &= r_j [\tilde{x}_j - y_j^*] + \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} - r_j d_j - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \\ &\leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [\tilde{x}_j - d_j - y_j^*] + \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [y_j^* - x_j + d_j] \\ &= \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [\tilde{x}_j - x_j] \\ &= \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j], \end{aligned}$$

where the inequality follows from the two inequalities that we derive at the beginning of this case and the last equality is by (11).

Case 3. Assume that $y_j^* < x_j - d_j$ and $\tilde{x}_j < y_j^*$. By the induction assumption, we have $\mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [y_j^* - x_j + d_j]$. Since y_j^* is a maximizer of $f_j(\cdot)$ over $[0, c]$, we have $\mathbb{E}\{v_{j+1}(\tilde{x}_j, D_{j+1})\} - \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} \leq r_j [\tilde{x}_j - y_j^*]$. Adding these two inequalities, we obtain $\mathbb{E}\{v_{j+1}(\tilde{x}_j, D_{j+1})\} - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [y_j^* - x_j + d_j] + r_j [\tilde{x}_j - y_j^*]$. Similar to Case 2, since y_j^* is a maximizer of $f_j(\cdot)$ over $[0, c]$ and $y_j^* < x_j - d_j$, we also have $\mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} \leq r_j$. In this case, by (10), we have

$$\begin{aligned} v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) &= \mathbb{E}\{v_{j+1}(\tilde{x}_j, D_{j+1})\} - r_j d_j - \mathbb{E}\{v_{j+1}(x_j - d_j, D_{j+1})\} \\ &\leq -r_j d_j + \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [y_j^* - x_j + d_j] + r_j [\tilde{x}_j - y_j^*] \\ &\leq \mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} [\tilde{x}_j - x_j] \\ &= \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j], \end{aligned}$$

where the first inequality follows from the inequality that we derive at the beginning of this case, the second inequality follows from the fact that $\mathbb{E}\{\dot{v}_{j+1}(x_j - d_j, D_{j+1})\} \leq r_j$, $\tilde{x}_j < y_j^*$ and $d_j \geq 0$, and the last equality is by (11).

Case 4. Assume that $x_j - d_j \leq y_j^* \leq x_j$ and $\tilde{x}_j - d_j \leq y_j^* \leq \tilde{x}_j$. By (10) and (11), we have $v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) = r_j [\tilde{x}_j - x_j] = \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j]$.

Case 5. Assume that $x_j - d_j \leq y_j^* \leq x_j$ and $\tilde{x}_j < y_j^*$. Since y_j^* is a maximizer of $f_j(\cdot)$ over $[0, c]$, we have $\mathbb{E}\{v_{j+1}(\tilde{x}_j, D_{j+1})\} - \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} \leq r_j [\tilde{x}_j - y_j^*]$. Therefore, by (10), we have

$$\begin{aligned} v_j(\tilde{x}_j, d_j) - v_j(x_j, d_j) &= \mathbb{E}\{v_{j+1}(\tilde{x}_j, D_{j+1})\} - r_j [x_j - y_j^*] - \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} \\ &\leq r_j [\tilde{x}_j - x_j] = \dot{v}_j(x_j, d_j) [\tilde{x}_j - x_j], \end{aligned}$$

where the last equality is by (11).

Case 6. Assume that $x_j < y_j^*$ and $\tilde{x}_j < y_j^*$. This case follows from the same argument in Case 1.