

Revenue Management Without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels

Garrett van Ryzin • Jeff McGill

Graduate School of Business, Columbia University, New York, New York 10027

School of Business, Queen's University, Kingston, Ontario, Canada

gvanryzi@research.gsb.columbia.edu • jmcgill@business.queensu.ca

We investigate a simple adaptive approach to optimizing seat protection levels in airline revenue management systems. The approach uses only historical observations of the relative frequencies of certain seat-filling events to guide direct adjustments of the seat protection levels in accordance with the optimality conditions of Brumelle and McGill (1993). Stochastic approximation theory is used to prove the convergence of this adaptive algorithm to the optimal protection levels. In a simulation study, we compare the revenue performance of this adaptive approach to a more traditional method that combines a censored forecasting method with a common seat allocation heuristic (EMSR-b).

(*Yield Management; Revenue Management; Airlines; Forecasting; Optimization; Fare Class Allocation; Distribution Free; Adaptive Algorithms; Stochastic Approximation*)

Introduction

Modern airlines must decide thousands of times per day whether or not to accept discount seat booking requests or refuse them in the hope of later, higher-fare bookings. Their objective is to manage the opening and closing of discount fare classes in such a way that overall expected revenues are maximized. This *revenue management* (also *yield management*) problem is greatly complicated by such factors as volatile, stochastic demand for air travel, fluctuations in fare prices, multiple-leg passenger itineraries, and diversion of passengers to other fare classes or flights. While no operations research model has succeeded in dealing with all of these complexities, simplified models and heuristics have been applied with remarkable success at many airlines. (See, for example, Smith et al. 1992). See Etschmaier and Rothstein (1974), Belobaba

(1987b), Weatherford and Bodily (1992), and McGill and van Ryzin (1999) for overviews and surveys.

One fundamental yield management model considers a single flight leg with booking requests arriving in order of booking class. This single-leg model can be analyzed to determine the structure of optimal booking policies and the optimal policy parameters. Belobaba (1987a, b, 1989), Brumelle and McGill (1993), Curry (1989), and Wollmer (1992) provide analyses of the single-leg model in which the booking fares increase monotonically from low to high as the time of flight departure approaches, and Robinson (1991) generalizes Brumelle and McGill's optimality conditions to the case that fares are nonmonotonic. Many yield management systems (called leg-based) use solutions from this elementary model to guide heuristic solutions for more realistic situations.

To illustrate, consider the simplest version of the single-leg model in which there are only two fare classes, demand in the discount fare class arrives before high-fare demand, and the high-fare demand is statistically independent of the discount demand. It is known that an optimal nonanticipating policy has the following structure: Set a fixed *protection* level θ for the high-fare seats. Seats can then be sold to the discount class as long as there are more than θ seats remaining. If high-fare demand, X , is modeled as a continuous random variable, the optimal protection level θ^* can be determined from the “newsvendor” type optimality condition first proposed (in discrete form) in Littlewood (1972):

$$P(X > \theta^*) = r, \quad (1)$$

where r is the ratio of the discount to the high fare.

Condition (1) stipulates that if the discount fare is, for example, 60% of the full fare, then the optimal protection level will be such that full-fare demand exceeds its protection level on 60% of all flights over the long run. Note that the frequency of such *fill events* does not correspond to the rate of *lost* full-fare bookings since some overflow demand can be accommodated when discount seats remain unbooked. The conditions for more than two fare classes are more complex than this but are conceptually similar.

Typically, application of conditions like (1) requires three steps. First, historical demand data are studied to determine suitable models for the demand distributions. Second, forecasting techniques are applied to estimate the parameters of these distributions. Because the booking limits themselves—or aircraft capacity constraints—cause censoring of the demand data, special techniques must be employed to “uncensor” the demand data. Third, the demand statistics from the forecast are passed to an optimization routine that solves for protection levels like θ^* . The resulting protection levels are then used to make individual accept-deny decisions as reservations come in. In practice, bookings from similar flights are fed back into the forecasting system, and the process is repeated cyclically over time.

Note that in each of these cycles, bookings data from the current departure are being converted—via the

forecasting and optimization procedures—into updated policy parameters for the next departure. This raises an interesting question: Is it possible to *directly* update booking policy parameters for the next departure based on simple observation of the performance of the parameters in previous departures, without recourse to the complex cycles of forecasting and optimization? Such ad hoc adjustment of protection levels was commonly used in the early days of yield management (and still is today in some airlines). However, most human analysts find it difficult to guess at revenue-maximizing protection levels.

In this paper, we show how to construct a simple and effective adjustment scheme by using properties of the optimal policy. Moreover, we show that under stationary demand conditions, the repeated application of our updating scheme eventually produces optimal booking policy parameters.

In §1 we propose a simple adaptive updating scheme that relies on observations of certain *fill events*, which correspond to subsets of fare classes reaching their respective protection levels. These fill events are easily determined from booking records data. Protection levels are updated based on a multivariate version of the stochastic approximation method of Robbins and Monro (1951), applied to the single-leg optimality conditions of Brumelle and McGill (1993). We prove in §3 that our proposed algorithm converges (almost surely) to an optimal set of protection levels, and we obtain bounds on the rate of convergence. Some modifications of the algorithm to handle practical issues like nonstationarity, integrality, and booking lead times are discussed in §4. In §5, we report results of a numerical comparison of our adaptive algorithm against a traditional procedure that combines censored forecasting and expected marginal seat revenue (EMSR) protection levels. (See Belobaba (1989)). Conclusions are provided in §6.

1. Notation, Model Assumptions, and Optimality Conditions

We let $\mathbf{1}(E)$ denote the indicator function of the event E ; that is, $\mathbf{1}(E) = 1$ if event E occurs and $\mathbf{1}(E) = 0$ otherwise. The expression $\sum_{n=1}^{\infty} x_n$ is abbreviated as $\sum_n x_n$, and (a.s.) is short for *almost surely*. Superscripts

on vectors or on elements of vectors index the members of a sequence of vectors; for example, $\{X^1, X^2, \dots\}$ is a sequence of demand vectors; while X_i^n is the demand for fare class i in the n -th demand vector. Subscripts will index sequences of scalar quantities; for example, $\{\gamma_1, \gamma_2, \dots\}$ is a sequence of scalar step sizes. Superscripts on such scalar quantities will have the usual interpretation as exponentiation.

Much of our analysis establishes upper bounds involving "sufficiently large" arbitrary constants on the right-hand sides of inequalities. To avoid a proliferation of such constants, we let C denote a generic, sufficiently large constant. The value of C changes throughout the paper depending on context. For example, a statement of the form $D \leq CE + C^2F$ can be replaced with $D \leq C(E + F)$, where $C \leftarrow \max\{C, C^2\}$.

We consider a model in which $k + 1$ fare classes book on a single-leg, fare-class allocation is nested (described below), low-fare classes book strictly before higher fare classes, fare-class demands are mutually independent, and there are no cancellations or no-shows. Let f_i denote the fare (or expected contribution) from fare class i , where we assume $f_1 > f_2 > \dots > f_{k+1}$. The demand for fare class i is denoted X_i . We assume $\{X_1, \dots, X_{k+1}\}$ are mutually independent, the probability distributions of demand are continuous, and that seat capacity is a continuous quantity. (See §4.2 for a discussion of integral demand and capacity.)

The stochastic process $\{X^1, X^2, \dots\}$ of demand vectors from successive flights is assumed to be stationary in the sense that the joint probability distributions of demand remain constant over time. This implies that the same protection levels are optimal for all flights in the sequence and that successive flights are comparable; for example, midweek morning commuter flights between two specific centers in "high season." Extensions to nonstationary demand processes are discussed in §4.3.

A *fixed protection level* policy for fare classes 1 through k is defined by a static set of protection levels given by the vector $\theta = (\theta_1, \dots, \theta_k)$, where $\theta_1 \leq \theta_2 \leq \dots \leq \theta_k$. (There is no protection level for the lowest fare class, $k + 1$.) Protection levels are *nested* in the sense that θ_i represents the number of seats to reserve (protect) for all of fare classes 1, 2, \dots , i . Reservations

for fare class $i + 1$ are accepted if and only if the number of seats remaining is strictly greater than the protection limit θ_i . For the single-leg problem under the assumptions stated above, it is shown in Brumelle and McGill (1993) that such fixed protection level policies are optimal among all nonanticipating policies.

Define the nested sequence of *fill events*:

$$\begin{aligned} A_1(\theta, X) &= \{X_1 > \theta_1\}; \\ A_2(\theta, X) &= \{X_1 > \theta_1, X_1 + X_2 > \theta_2\}; \dots \\ A_i(\theta, X) &= \{X_1 > \theta_1, X_1 + X_2 \\ &> \theta_2, \dots, X_1 + \dots + X_i > \theta_i\}. \end{aligned} \quad (2)$$

We refer to these events as fill events because, when the current inventory of unbooked seats is θ_i , all remaining seats are sold if and only if the event $A_i(\theta, X)$ occurs.

Let the vector $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_k^*)$ denote an optimal set of protection levels. Brumelle and McGill (1993) show that, when all passenger demand distributions are continuous, an optimal θ^* satisfies

$$r_{i+1} = P(A_i(\theta^*, X)), \quad \text{for } i = 1, \dots, k, \quad (3)$$

where $r_{i+1} = f_{i+1}/f_1$ is the *discount ratio* of fare $i + 1$ relative to the full fare. The parallel with Littlewood's Rule (1) is clear: For example, if discount fare i is 60% of the *full* fare, optimal protection levels will be such that event $A_i(\theta, X)$ occurs on 60% of future flights in the long run.

The conditions in (3) are independent of aircraft capacity and assume continuity of demand. In practice, protection levels that exceed aircraft capacity can be simply truncated to the capacity without loss of optimality. This will be discussed further below. Also, the continuous protection levels obtained by (3) are generally good approximations for the true optimal protection levels for integer-valued demand. Indeed, Brumelle and McGill (1993) show that there will be at least one set of fixed, integer-valued, optimal protection levels when demand is integer-valued. Moreover, the revenue management problem is known to be robust to small departures from the optimal protection levels. (See Brumelle and McGill (1993) and Robinson

(1991).) Wollmer (1992) analyzes integer-optimal protection levels when demand forecasts are available as discrete probability distributions. (See §4.2 for further discussion of integrality.)

In principle, it is easy to determine the frequency of the events $A_i(\theta, X)$ from a record of past flights. No uncensoring of the demand is required—it is only necessary to observe if demand reached the protection levels, not the degree to which it exceeded them. There are two important exceptions to this. First, if θ_i happens to exceed the maximum number of seats available for sale (usually the physical capacity plus an overbooking “pad”), then the event $X_1 + \dots + X_i > \theta_i$ is not observable.¹ Second, if protection levels are revised during the lead time prior to flight departure it can easily happen that a new protection level exceeds the remaining capacity on the aircraft (a problem similar to the first one), or that earlier, high protection levels constrained demand during part of the booking period in one or more discount fare classes. In this case, total demand is not observed relative to θ_i (a variant of censorship of the demand data).

Initially, we will make the simplifying assumption that $A_i(\theta, X)$ is always observable. We show in §4 how the algorithm and results can be modified to allow for the physical capacity constraint and booking lead times in practical implementations.

The conditions in (3) are appealing on practical grounds because one can check the optimality of protection levels in a series of departed flights retroactively by simply comparing the fraction of flights on which $A_i(\theta, X)$ occurred to the discount ratio r_{i+1} . This approach has the distinct advantage that it requires no assumptions on the specific nature of the probability distributions of the demands. Observing fill event frequencies therefore provides a distribution-free test of the optimality of protection levels.

However, what happens if the observed frequencies do not equal the discount ratio? Next, we show that these same fill event conditions can be used to adap-

tively adjust protection levels using an extension of the classical stochastic approximation algorithm of Robbins and Monro (1951).

2. An Adaptive Algorithm

For $i = 1, \dots, k$ define $H_i(\theta, X) = r_{i+1} - \mathbf{1}(A_i(\theta, X))$. The quantity $H_i(\theta, X)$ will be negative if the event A_i occurs, and positive otherwise. If protection levels are being adjusted, occurrence of the event A_i (all of classes 1 through i reached their protection levels) suggests that the protection level θ_i should be adjusted upwards. Thus $-H_i(\theta, X)$ can be viewed as an *adjustment direction* for protection level θ_i . The corresponding *adjustment vector* is $H(\theta, X) = (H_1(\theta, X), \dots, H_k(\theta, X))$.

Now define $h_i(\theta) = r_{i+1} - P(A_i(\theta, X))$, $i = 1, \dots, k$; and let $h(\theta) = (h_1(\theta), \dots, h_k(\theta))$. Note that $h(\theta) = EH(\theta, X)$. Thus, $-h(\theta)$ can be properly viewed as the *expected adjustment vector* for protection levels given current levels θ . The optimality Condition (3) stipulates that we should seek a θ^* such that the expected adjustment for all protection levels is zero; or, $h(\theta^*) = 0$.

The Robbins-Monro procedure (generalized here for vector quantities) constructs a sequence of parameter estimates, $\{\theta^1, \theta^2, \dots, \theta^n, \dots\}$, from a sequence of independent trials, $\{X^1, X^2, \dots, X^n, \dots\}$, using

$$\theta^{n+1} = \theta^n - \gamma_n H(\theta^n, X^n), \tag{4}$$

where γ_n is a sequence of nonnegative step sizes satisfying

$$\sum_n \gamma_n = +\infty \quad \text{and} \quad \sum_n \gamma_n^2 < +\infty. \tag{5}$$

(The simplest example of a suitable step-size sequence is defined by $\gamma_n = 1/n$, however, this simple averaging sequence takes small steps early in the procedure, which can delay convergence. In the development to follow, we use a sequence of the form $A/(n + B)$, where A and B are constants chosen to effect larger early steps.)

The directions $H(\theta^n, X^n)$ can be determined after the departure of each flight. If the fill event A_i occurs, $H_i^n = r_{i+1} - 1 < 0$ and the protection level θ_i^n is increased by $\gamma_n(1 - r_{i+1})$; if not, then $H_i^n = r_{i+1} > 0$,

¹ One could still observe this event if rejected sales were recorded, but this information is not available to most airlines. Rejections occur at the point of sale (e.g., the travel agent) and are not recorded in the reservation system.

and θ_i^n is reduced by $\gamma_n r_{i+1}$. Thus protection levels are stepped up when high demand is observed and stepped down when low demand is observed, with the step size becoming smaller as the algorithm progresses.

A key theoretical and practical problem is determining conditions under which such an adaptive algorithm will converge to optimal protection levels. It is known that almost sure convergence to a point θ^* satisfying $h(\theta^*) = 0$ is guaranteed if $h(\theta)$ is the gradient of a concave function with a unique maximum at θ^* or, more generally, if

$$\inf_{\epsilon < \|\theta - \theta^*\| < (1/\epsilon)} \{(\theta - \theta^*)^T h(\theta)\} > 0 \quad \text{for all } \epsilon > 0. \quad (6)$$

(See Benveniste et al. (1990) and Blum (1954). Loosely speaking, this condition stipulates that the expected adjustment vector $-h(\theta)$ always points into the half-space containing θ^* .)

However, with a nested allocation policy, the expected revenue is not jointly concave in the protection levels. More to the point, the function $h(\theta)$ is not even the gradient of the expected revenue function; rather, it is a collection of gradients from a sequence of scalar subproblems, one for each stage in the booking process during which reservations for a particular fare class arrive. (See Brumelle and McGill (1993).) Therefore, we cannot assume that the iterates satisfy a joint stability condition like (6).

A second difficulty is that the Algorithm (4) can produce vectors θ that violate the monotonicity condition $\theta_1 \leq \theta_2, \dots, \leq \theta_k$ required for the events A_i to be appropriately nested. To address this problem, we define interim protection levels p_i , where

$$p_i(\theta) = \max\{\theta_j : 1 \leq j \leq i\}, \quad (7)$$

and let $p(\theta) = (p_1(\theta), \dots, p_k(\theta))$. (Note that if $\theta_1 \leq \theta_2, \dots, \leq \theta_k$, then $p(\theta) = \theta$.) The algorithm uses the protection levels $p_i(\theta)$ to control the availability of seats, but monitors the events $A_i(\theta, X)$ to update the parameters θ . Because $\theta_i \leq p_i(\theta)$, this modification preserves our ability to observe the event $A_i(\theta, X)$ when the components of θ are not monotonically increasing, and allows us to treat θ as unconstrained.

We will show in the next section that, in spite of these difficulties, reasonable regularity conditions en-

sure that the Procedure (4) does converge (a.s.) to a value θ^* satisfying $h(\theta^*) = 0$, and that $p(\theta^*) = \theta^*$ (a.s.).

3. Convergence Proof

To prove the convergence of the Iteration (4) we need several preliminary lemmas. The first provides a sufficient condition for the almost sure convergence of a series of random variables:

LEMMA 1. (Lukacs (1975), Theorem 4.2.1) Let $\{Y_n\}$ be a sequence of random variables with $EY_n < +\infty$ for all n and $\sum_n E|Y_n| < +\infty$. Then $\sum_n |Y_n| < +\infty$ (a.s.).

The following supermartingale lemma is essential to many convergence proofs in stochastic approximation. It is due originally to Robbins and Siegmund (1971) (see also Benveniste et al. (1990), p. 344):

LEMMA 2. Let $\{\Omega, F, F_n, P\}$ be a probability space with an increasing family of σ -fields F_n . Suppose Z_n, B_n, C_n , and D_n are finite, nonnegative random variables, adapted to the σ -field F_n , which satisfy $E(Z_{n+1}|F_n) \leq (1 + B_n)Z_n + C_n - D_n$. Then on the set $\{\sum_n B_n < \infty, \sum_n C_n < \infty, \sum_n D_n < \infty$ (a.s.), and $Z_n \rightarrow Z < \infty$ (a.s.).

We shall require that θ remain bounded (a.s.). The following lemma follows easily from (4) and the fact that the step sizes γ_n are nonnegative and decreasing:

LEMMA 3. If demand X^n has bounded support for $n = 1, 2, \dots$ (i.e., $X_i^n \leq C$ (a.s.) for $i = 1, 2, \dots$), then $\|\theta^n\|$ is bounded (a.s.).

Finally, we will need the following lemma, which is adapted from Benveniste et al. (1990 Lemma 23, p. 245):

LEMMA 4. Let $\gamma_n = A/(n + B)$, where A and B are constants such that $A > 0$ and $B \geq 0$, and let $u_n = \lambda \gamma_n^\beta$, where λ and β are nonnegative. Let $\delta > 0$ and $C > 0$ be arbitrary constants. Then there exists a $\beta > 0$, $\lambda < +\infty$, and a n_0 such that the inequality

$$u_{n+1} \geq (1 - 2\delta\gamma_n)u_n + C\gamma_n^{\beta+1}$$

is satisfied for all $n > n_0$.

PROOF. By substituting $u_n = \lambda \gamma_n^\beta$ in the above inequality and rearranging, we obtain

$$\lambda(\gamma_{n+1}^\beta - \gamma_n^\beta) \geq (C - 2\delta\lambda)\gamma_n^{\beta+1}.$$

Now, if $\gamma_n = A/(n + B)$, then by taking a Taylor series expansion, we find that

$$\gamma_{n+1}^\beta = \gamma_n^\beta - \beta A^\beta (n + B)^{-(\beta+1)} + O(n^{-(\beta+2)}).$$

Substituting this expression into the first term on the left-hand side above and simplifying we obtain

$$\beta - O(n^{-2}) \leq A \left(2\delta - \frac{C}{\lambda} \right).$$

Because $\delta > 0$, we can choose λ sufficiently large and $\beta > 0$ sufficiently small so that $\beta < A(2\delta - (C/\lambda))$. The existence of a n_0 satisfying the conditions of the lemma then follows. \square

We are now ready to prove our main convergence result:

THEOREM 1. *Let θ^n be defined as in (4), and let γ_n be defined as in Lemma 4. Suppose the following assumptions hold:*

A1) *Each X_i has bounded support (i.e., $P(X_i < C) = 1$ for some constant C).*

A2) *There exists a $\delta > 0$ such that, for all $\theta_i, (\theta_i - \theta_i^*)h_i(\theta_i, \theta_{i-1}^*, \dots, \theta_1^*) \geq \delta|\theta_i - \theta_i^*|^2$.*

A3) *The distributions of the partial sums, $X_1 + \dots + X_i$ are Lipschitz continuous. That is, for all $i = 1, \dots, k|P(X_1 + \dots + X_i < x) - P(X_1 + \dots + X_i < y)| \leq C|x - y|$.*

Then, for $i = 1, \dots, k$,

$$\theta_i^n \rightarrow \theta_i^* \text{ (a.s.)} \tag{8}$$

$$\text{and } p_i(\theta^*) = \theta_i^* \tag{9}$$

Furthermore, there exists a $\beta > 0$ such that for $i = 1, 2, \dots, k$,

$$E|\theta_i^n - \theta_i^*|^2 \leq C\gamma_n^{\beta/2^{i-1}}. \tag{10}$$

PROOF. The proof is by induction on the fare classes.

First, consider Fare Class 1 and note that the event $A_1(\theta, X)$ is not a function of $\theta_i, i > 1$. It is easy to see that the sequence $\{\theta_1^n\}$ is a classical (scalar) Robbins–Monro process and hence converges (a.s.). Indeed, from A2 we have that for all $\epsilon > 0$,

$$\inf_{\epsilon < \|\theta_1 - \theta_1^*\| < (1/\epsilon)} \{(\theta_1 - \theta_1^*)h_1(\theta_1)\} > 0,$$

and $H(\cdot)$ is uniformly bounded, so $E|H_1(\theta, X)|^2 \leq C$ for all θ . These conditions together with the fact that the gain sequence $\{\gamma_n\}$ satisfies (5) guarantee that $\theta_1^n \rightarrow \theta_1^*$ (a.s.). (See Benveniste et al. 1990 and Robbins and Sgund (1971) for proofs.) The fact that $p_1(\theta^*) = \theta_1^*$ follows trivially from (7). Finally, from Benveniste et al. (Theorem 22, p. 244) we have that with A2 and the gain sequence $\{\gamma_n\}$ there exists a constant λ such that for all $0 < \beta < 1$, we have $E|\theta_1^n - \theta_1^*|^2 \leq C\lambda\gamma_n^\beta$. Thus for $i = 1$, (8)–(10) are all satisfied.

Now, suppose (8)–(10) hold for fare class i and consider fare class $i + 1$. Define $T_n = \theta_{i+1}^n - \theta_{i+1}^*$ and $Z_n = |T_n|^2$. Then by (4)

$$\begin{aligned} Z_{n+1} &= Z_n - 2\gamma_n T_n H_{i+1}(\theta^n, X_{i+1}^n, X_i^n, \dots, X_1^n) \\ &\quad + \gamma_n^2 |H_{i+1}(\theta^n, X_{i+1}^n, X_i^n, \dots, X_1^n)|^2. \end{aligned}$$

Taking expectations conditioned on F_n yields

$$\begin{aligned} E(Z_{n+1}|F_n) &= Z_n - 2\gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n) \\ &\quad + \gamma_n^2 E[|H_{i+1}(\theta^n, X_{i+1}^n, X_i^n, \dots, X_1^n)|^2]. \end{aligned}$$

Since $H(\cdot)$ is uniformly bounded, we have that $E[|H_{i+1}(\theta, X_{i+1}^n, X_i^n, \dots, X_1^n)|^2] \leq C$ for all θ , hence

$$\begin{aligned} E(Z_{n+1}|F_n) &\leq Z_n - 2\gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n) \\ &\quad + C\gamma_n^2 = Z_n - 2\gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_i^*, \dots, \theta_1^*) + C\gamma_n^2 \\ &\quad + 2\gamma_n T_n [h_{i+1}(\theta_{i+1}^n, \theta_i^*, \dots, \theta_1^*) \\ &\quad - h_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n)]. \end{aligned} \tag{11}$$

We next bound the last term in (11). From the definition of $h_i(\cdot)$, we have

$$\begin{aligned} &|h_{i+1}(\theta_{i+1}^n, \theta_i^*, \dots, \theta_1^*) - h_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n)| \\ &\leq E|H_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n) \\ &\quad - H_{i+1}(\theta_{i+1}^n, \theta_i^*, \dots, \theta_1^*)| \\ &\leq E[|\mathbf{1}(X_1^n + \dots + X_{i+1}^n > \theta_{i+1}^n)| \\ &\quad \cdot |\mathbf{1}(A_i(\theta^n)) - \mathbf{1}(A_i(\theta^*))|] \\ &\leq E|\mathbf{1}(A_i(\theta^n)) - \mathbf{1}(A_i(\theta^*))| \\ &= P(\mathbf{1}(A_i(\theta^n)) \neq \mathbf{1}(A_i(\theta^*))). \end{aligned}$$

Now define the events

$$E_j = \{\min\{\theta_j^*, \theta_j^n\} < X_1 + \dots + X_j \leq \max\{\theta_j^*, \theta_j^n\}\},$$

and observe that if $\mathbf{1}(A_i(\theta^n)) \neq \mathbf{1}(A_i(\theta^*))$, then at least one of the events $E_j, j = 1, \dots, i$ must occur. Combining this observation with A3, we obtain

$$P(\mathbf{1}(A_i(\theta^n)) \neq \mathbf{1}(A_i(\theta^*))) \leq \sum_{j=1}^i P(E_j) \leq C \sum_{j=1}^i |\theta_j^n - \theta_j^*|.$$

Therefore, we have

$$|h_{i+1}(\theta_{i+1}^n, \theta_{i'}^* \dots, \theta_1^*) - h_{i+1}(\theta_{i+1}^n, \theta_i^n, \dots, \theta_1^n)| \leq C \sum_{j=1}^i |\theta_j^n - \theta_j^*|.$$

Substituting this bound in (11) and using the fact that by A1 and Lemma 3, $|T_n| \leq C$ (a.s.) we obtain

$$E(Z_{n+1}|F_n) \leq Z_n - 2\gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_{i'}^* \dots, \theta_1^*) + C\gamma_n^2 + C \sum_{j=1}^i \gamma_n |\theta_j^n - \theta_j^*|. \tag{12}$$

We next show that the last term above is finite (a.s.). Indeed, we have by A1 and Lemma 3 that $E(\gamma_n |\theta_i^n - \theta_i^*|)$ is bounded. Also, by the induction hypothesis, (10) holds for i , and therefore

$$\sum_n E(\gamma_n |\theta_i^n - \theta_i^*|) \leq \sum_n \gamma_n \sqrt{E(|\theta_i^n - \theta_i^*|^2)} \leq C \sum_n \gamma_n^{1+\beta/2^i} < +\infty,$$

where the last inequality follows from the definition of $\{\gamma_n\}$ and the fact that $\sum_n 1/n^p$ converges for all $p > 1$. Applying Lemma 1, we conclude that $\sum_n \gamma_n |\theta_i^n - \theta_i^*| < +\infty$ (a.s.). Since $\sum_n \gamma_n^2$ is bounded, we can apply A2 and Lemma 2 to (12) and conclude that $Z_n \rightarrow Z < +\infty$ (a.s.) and $\sum_n \gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_{i'}^* \dots, \theta_1^*) < +\infty$ (a.s.).

We next show that $Z = 0$. Indeed, if $Z > 0$ then Assumption A2 ensures that there exists an N such that $T_n h_{i+1}(\theta_{i+1}^n, \theta_{i'}^* \dots, \theta_1^*) > 0$ for all $n > N$. But this in turn would imply $\sum_n \gamma_n T_n h_{i+1}(\theta_{i+1}^n, \theta_{i'}^* \dots, \theta_1^*)$ is unbounded, which is a contradiction. Therefore, $Z = 0$ and (8) is proven for $i + 1$.

We next show (9) by contradiction. Indeed, if (9) is not true, then since $p_i(\theta^*) = \theta_i^*$ by the induction hypothesis, we must have $\theta_{i+1}^* < \theta_i^*$. This in turn implies $A_{i+1}(\theta^*, X) = A_i(\theta^*, X)$, which violates (3) if $f_{i+2} < f_{i+1}$. Thus, we must have $p_{i+1}(\theta^*) = \theta_{i+1}^*$.

Finally, to show (10) holds for $i + 1$ we apply A2 to (12), which yields

$$E(Z_{n+1}|F_n) \leq (1 - 2\gamma_n \delta) Z_n + C\gamma_n^2 + C \sum_{j=1}^i \gamma_n |\theta_j^n - \theta_j^*|.$$

Unconditioning and using the induction hypothesis that (10) holds for i we obtain,

$$E(Z_{n+1}) \leq (1 - 2\gamma_n \delta) E(Z_n) + C\gamma_n^2 + C \sum_{j=1}^i \gamma_n E(|\theta_j^n - \theta_j^*|) \leq (1 - 2\gamma_n \delta) E(Z_n) + C\gamma_n^2 + C \sum_{j=1}^i \gamma_n \sqrt{E(|\theta_j^n - \theta_j^*|^2)} \leq (1 - 2\gamma_n \delta) E(Z_n) + C\gamma_n^{1+\beta/2^i}.$$

Applying Lemma 4 to the above inequality, we conclude that there exist constants $\lambda < +\infty, \beta > 0$, and $n_0 < +\infty$ such that the sequence $u_n = \lambda \gamma_n^{\beta/2^i}$ satisfies

$$u_{n+1} \geq (1 - 2\gamma_n \delta) u_n + C\gamma_n^{1+\beta/2^i},$$

for all $n > n_0$. Taking λ sufficiently large so that $EZ_{n_0} \leq \lambda \gamma_{n_0}^{\beta/2^i}$ and applying induction on n we have $EZ_n \leq \lambda \gamma_n^{\beta/2^i}$. Therefore (10) also holds for $i + 1$, and the induction is complete. \square

Some comments on Assumptions A2 and A3 are in order. A2 requires that the distribution of demand not be “too flat” in the neighborhood of θ^* . For example, one can show that A2 is satisfied if the distributions have a density that is bounded below by a strictly positive constant in the neighborhood of θ^* . Intuitively, A2 is needed because if the distribution is too flat near θ^* , the algorithm could “stall” before reaching θ^* . The Lipschitz condition A3 is satisfied if the individual demand distributions are Lipschitz smooth. For

example, A3 holds if the demand distributions have uniformly bounded densities. Both A2 and A3 are not overly restrictive.

Finally, note that (10) suggests that the convergence rate decreases with i , the index of the fare class. That is, lower fare classes have slower convergence than higher fare classes. Our numerical results in §5 illustrate this behavior.

4. Practical Modifications to the Basic Algorithm

In this section, we discuss four modifications to the basic algorithm (4) that address problems mentioned earlier.

4.1. Capacity Constraint

Let c be the leg capacity or maximum number of bookings allowed. The recursion (4) can then be modified to

$$\theta_i^{n+1} = \min\{c, (\theta_i^n - \gamma_n H_i(\theta_i^n, X^n))\}, \quad i = 1, \dots, k, \quad (13)$$

which corresponds to projecting θ^{n+1} onto the constraint set $[0, c]^n$. With this modification, it is not difficult to show that the algorithm converges to a point $\hat{\theta}$, where $\hat{\theta}_i = \min\{\theta_i^*, c\}$, $i = 1, \dots, n$.

4.2. Integrality

In practice, passenger demand and seat allocations are integral so adjustments of less than one seat are not feasible. However, without fractional adjustments the algorithm could become “stuck” at a nonoptimal point. One solution to this dilemma is to randomize the choice of protection levels.

For illustration, we only consider the case $k = 1$ (two fare classes). Let θ be a continuous parameter and let p be the actual protection level used. Let U be a uniform $[0, 1]$ random variable. Then at iteration n , we use protection level p_n where

$$p_n = \begin{cases} \lfloor \theta \rfloor & U \leq \theta - \lfloor \theta \rfloor \\ \lceil \theta \rceil & U > \theta - \lfloor \theta \rfloor \end{cases}. \quad (14)$$

This corresponds to randomizing the selection of $\lfloor \theta \rfloor$ and $\lceil \theta \rceil$ based on the value of θ . We then redefine the event A_1 at iteration n to be $A_1(\theta, X) = \{X_1 > p_n\}$.

Under this scheme, $h_1(\theta)$ is continuous and, provided mild conditions on the discrete demand distribution are met,² it satisfies the conditions of Theorem 1.

Note in this case that θ^n converges to a θ^* satisfying $P\{X_1 > \kappa^* + 1\{U > \theta^* - \kappa^*\}\} = r_2$, for some integer κ^* , where κ^* satisfies $P\{X_1 > \kappa^*\} \leq r_2$ and $P\{X_1 > \kappa^* + 1\} > r_2$. Therefore, κ^* is the optimal integer protection level. The policy, however, randomizes between a protection level κ^* and $\kappa^* + 1$, which results in some deviation from optimality. Nevertheless, the simplicity of this randomization scheme is attractive.

4.3. Booking Lead Times

Many airlines open bookings for flights 10 months or more prior to flight departure. However, most booking activity occurs over a shorter time span, typically 30 to 60 days before departure. It is this shorter, effective booking lead time that is relevant to setting seat protection levels. Standard airline practice is to fix one set of protection levels at the beginning of the full booking period and then delay further adjustments until effective booking begins. Thereafter, reading of booking levels and (possible) adjustment of protection levels occurs with increasing frequency as departure time approaches. A total of 15 readings and adjustments across the full booking period is common in practice. These multiple adjustments are designed to accommodate the nonstationarity of demand and to incorporate recent demand data into current booking levels.

In this subsection we discuss modifications to the basic adaptive scheme necessary to accommodate the effective booking period. For purposes of this discussion, consider a flight that departs every week and that receives most bookings in the 10 weeks prior to each departure.

A direct, but unsatisfactory, implementation of algorithm (4) would involve dividing the flight sequence into 10 interleaved, independent sequences, each of which is updated once every 10 weeks. In the first of these sequences, fill events from the flight that departs in Week 0 would be used to set protection levels for the flight that will be departing in Week 10,

² Namely, that the probability mass function is bounded away from zero in the neighborhood of the stationary point.

and that departure will be used to set levels for the Week 20 departure, and so on. The second through tenth sequences develop in a similar way from the Week 1 through 9 departures. This implementation is unsatisfactory for at least two reasons: First, data (on demand, which is assumed stationary) are not shared across the separate flight sequences, and second, the protection levels within each sequence are adjusted very slowly over time (once every 10 weeks).

Fortunately, a simple modification of the adaptive algorithm can be used to produce a single sequence of protection levels that use information from all flights and have the same convergence properties as those determined by the original algorithm. To see this, let k denote the number of time units (e.g., weeks) in the effective booking lead time, and assume a new flight departs every unit of time, where time units are indexed by n as before. In this case, bookings for the first flight begin (effectively) at $n = 1$, the first flight departs at $n = k$, and the first complete observations of fill events are not available until time $k + 1$. Flights $1, 2, \dots, k$ must use protection levels based on initial guesses or information external to the algorithm. We assume that the initial k protection levels are identical, viz

$$\theta^1 = \theta^2 = \dots = \theta^k. \tag{15}$$

Let X^n be the demand for the flight, labelled n , that begins booking at time n , and departs at time $n + k$. Flight n uses protection levels θ^n given by (15) for $n \leq k$, and for $n > k$ it uses

$$\theta^{n+1} = \theta^n - \gamma_{n-k} H(X^{n-k}, \theta^{n-k}). \tag{16}$$

This adjustment may seem strange—an adjustment direction away from an old protection vector θ^{n-k} is being applied to the current vector θ^n —but, it turns out that this is simply a delayed version of the original recursion (4). Indeed, summing the original recursion (4) we obtain

$$\theta^{n+1} = \theta^1 - \sum_{i=1}^n \gamma_i H(X^i, \theta^i).$$

Similarly, by summing the modified recursion (16) and using the fact that $\theta^1 = \theta^k$ from the initial condition (15) we obtain

$$\theta^{n+1} = \theta^1 - \sum_{i=1}^{n-k} \gamma_i H(X^i, \theta^i).$$

Thus θ^n in the modified recursion (16) is equal to θ^{n-k} in the original recursion (4). It therefore follows that Theorem 1 holds for the modified sequence (16).

What about updating protection levels during the booking process itself? (e.g., a flight that starts booking in Week 1 for departure in Week 11 could possibly use different booking levels in each week prior to departure). This creates difficulty because fill events $A_i(\theta, X)$ no longer have the same interpretation. In particular, in this case the fact that the *bookings on hand* for a fare class at the time of departure are less than the booking limit does not mean that total *demand* was less than the booking limit. Such behavior could easily cause the iteration (4) to fail, but this question deserves further investigation.³

4.4. Nonstationary Demand

The stochastic approximation scheme follows a sequence of steps of decreasing size to converge to an optimal set of protection levels. Strictly speaking, this restricts the approach to series of flights for which the protection levels are not believed to vary. In reality, however, optimal protection levels drift over time as demand and pricing factors vary. Variations on stochastic approximation have been developed to deal with such nonstationary systems. In general, they track a moving set of optimal parameters with steps that do not tend to zero in size. See Benveniste et al. (1990).

We have conducted simulations with nonstationary demand using a variant of stochastic approximation that gives more weight to recent fill event frequencies in a manner similar to exponential smoothing. Not surprisingly, the method correctly tracks the optimal protection levels, but with a lag typical of (simple) exponential smoothing. We conjecture that this type of tracking system will perform reasonably well, but again this topic deserves further investigation.

³ It is important to recognize that this type of dynamic censorship is also a problem in conventional forecasting/optimization systems. A dissertation by Lee (1990) addresses this problem.

5. Numerical Examples

In this section we present numerical examples of the performance of our adaptive algorithm in the stationary demand case. For comparison, we solve the same examples with a representative procedure that combines censored forecasting with EMSR protection levels. We compare the convergence of the protection levels produced by the two methods and their revenue performance under various starting conditions, load factors, and demand distributions.

These comparisons are based on simulated data in an idealized stationary setting. While the examples are useful for illustrating some of the operating characteristics of the methods, they cannot lead to final conclusions about relative merits. Such conclusions can only come from trials in practice.

5.1. A Combined Forecasting-EMSR Scheme

The combined forecasting-EMSR scheme constructs a demand forecast based on a correction to the censored observations of demand in each fare class. For our tests, we corrected for censorship with an estimate of the survivor function $S(x) = P(X > x)$ based on life tables. Details can be found in Lawless (1982, §2.2).

The life table estimator works as follows: Let n denote the total number of observations (censored and uncensored). Let $t_1 < t_2 < \dots, t_m$ be m distinct intervals. (We call $[t_j, t_{j+1})$ interval j .) Let n_j be the number of observations with values t_j or more (the number of "at risk" observations at the start of interval j); let d_j be the number of uncensored observations that fall in interval j (the number of "deaths" in the interval j); and let w_j be the number of censored observations that fall in interval j (the number of "withdrawals" because of censoring in interval j). Define $n_0 = n$ and note that $n_j = n_{j-1} - d_j - w_j, j = 1, \dots, m$.

Then the *standard life table estimate* is given by

$$\hat{S}\left(\frac{t_j + t_{j+1}}{2}\right) = \prod_{i=1}^j \left(1 - \frac{d_i}{n_i - w_i/2}\right),$$

$j = 1, \dots, m.$

(Note the approximation of $S(t)$ is taken at the midpoint of interval j .) The idea here is that each term 1

$- (d_i/n_i - w_i/2)$ is an estimate of the conditional probability that demand exceeds t_{i+1} given that it exceeded t_i . The denominator, $n_i - w_i/2$, is an estimate of the number of samples at risk during period i , where $w_i/2$ is a correction term for the number of censored observations in period i (e.g., a censored observation during period i is assumed to be at risk for half the period on average).

In our implementation, we maintained 20 intervals ($m = 21$ for each fare class, with $0 = t_1 < t_2 \dots < t_m < t_{m+1} = +\infty$, chosen so that $P(X \in [t_j, t_{j+1})) = 0.05$ for all $j = 1, \dots, m$). While more intervals clearly result in a more accurate estimate of the survivor function, 20 provided adequate accuracy in our case, especially as the values t_1, \dots, t_m were chosen to match each distribution.⁴

We then used the life table estimator to estimate the mean and standard deviation of the distribution by linear regression. Specifically, let $\Phi(x)$ be the standard normal distribution and let $\Phi^{-1}(x)$ denote its inverse. Define $s_j = \Phi^{-1}(1 - \hat{S}((t_j + t_{j+1})/2))$. If demand is normally distributed, the points $(s_j, t_j) j = 1, \dots, n$ should lie approximately on a straight line, namely $s_j = at_j + b$. Using linear regression, we estimated the slope, \hat{a} , and the intercept, \hat{b} , and then constructed estimates of the mean, $\hat{\mu} = 1/\hat{a}$, and standard deviation, $\hat{\sigma} = -\hat{b}/\hat{a}$. This procedure results in potentially biased estimates (See Lawless 1982, §2.5.), but it is simple to implement and seemed to perform well in our tests. We updated the life table and linear regression estimates of the mean and standard deviation after each simulated flight departure.

This procedure may not always produce an unbounded estimate of the mean and standard deviation. In particular, if all n samples are censored, then the life table estimate is 1 for all values of j and the linear regression produces an estimate of $\hat{a} = 0$, which results in an unbounded estimate of the mean and standard deviation. This is quite normal behavior for censored demand estimators. Indeed, if all observations are censored, then any reasonable estimator (e.g.,

⁴Of course, in practice one would not be able to fine-tune the intervals of time t_1, \dots, t_m so precisely, since the demand distribution is unknown. Thus, a wider range—with more intervals—would be required to ensure that the data were adequately covered.

maximum likelihood) either will be uncomputable or produce an unbounded mean and/or standard deviation. In such cases, we ignored the forecast and simply maintained the current protection levels until the forecast produced bounded estimates.

For setting seat protection levels, we used a variation of the expected marginal seat revenue (EMSR) heuristic (Belobaba 1989), called EMSR-b. This is the most common seat protection heuristic used in practice. EMSR-b works as follows: Given estimates of the means, $\hat{\mu}_i$, and standard deviations, $\hat{\sigma}_i$, for each fare class i , the EMSR-b heuristic sets protection level θ_i so that $f_{i+1} = \bar{f}_i P(\bar{X}_i > \theta_i)$, where \bar{X}_i is a normal random variable with mean $\sum_{j=1}^i \hat{\mu}_j$ and variance $\sum_{j=1}^i \hat{\sigma}_j^2$, and \bar{f}_i is a weighted average revenue, given by

$$\bar{f}_i = \frac{\sum_{j=1}^i f_j \hat{\mu}_j}{\sum_{j=1}^i \hat{\mu}_j}.$$

The idea behind this approximation is to reduce the complexity of the fully nested problem by aggregating fare classes 1, 2, . . . , i into a single-fare class. Then, one treats the problem as a simple, 2-fare-class problem.

Again, we emphasize that this overall forecasting-EMSR scheme is not constructed to be the most sophisticated one possible. Rather, it is intended be representative of a basic yield management system. Both methods could no doubt be refined further. However, such refinements often involve idiosyncratic and/or ad hoc modifications that only serve to make performance comparisons more complex, controversial, and ultimately less insightful. Our intention, therefore, is to perform a transparent test of one simple approach (the adaptive algorithm) against another simple approach (basic censored forecasting and EMSR-b).

5.2. Test Problem Scenarios

Our test problem is a modification of Wollmer’s (1992) 5-fare-class example, in which we aggregated his Classes 3 and 4 to reduce the problem to 4 classes. (Classes 3 and 4 in Wollmer had very similar fares of \$534 and \$520, respectively.) With 4-fare-classes, there are $n = 3$ protection levels to determine.

Table 1 Fares, Demand Statistics, and Protection Levels for Numerical Examples

Class	Fare	Mean	Std. Dev.	θ -EMSR	θ^* -Normal	θ^* -Log N
1	\$1,050	17.3	5.8	16.7	16.7	15.9
2	\$567	45.1	15.0	51.5	44.6	45.7
3	\$527	73.6	17.4	131.4	134.0	130.0
4	\$350	19.8	6.6	n.a.	n.a.	n.a.

Table 2 Starting Values of Protection Levels for Numerical Examples

	θ_1	θ_2	θ_3
Low	0	15	65
High	35	110	210

The data along with optimal and EMSR-b protection levels are shown in Table 1. The protection level θ^* -Normal is the optimal level when demand is normally distributed, while θ^* -Log N is the optimal protection level when demand is log-normally distributed.

To test the convergence of the adaptive algorithm and the forecasting-optimization scheme, we purposely started with protection levels that were far from optimal, corresponding to high and low starting values (see Table 2). These somewhat extreme values were chosen to test the convergence properties of each algorithm. In practice, one may have some prior knowledge about demand that can be used to set better initial protection levels. At the same time, using extreme starting values provides a good robustness test. High starting values produce less initial censoring in the higher fare classes, so one has better observations of the actual demand distributions. However, revenues may be low due to high levels of rejected demand. Low starting values produce severe initial censoring, which may adversely affect forecast accuracy. In terms of revenue, low starting values are generally better if demand is low, but may produce poor revenue performance when demand is high because insufficient capacity is reserved for higher fare classes. Observing how the algorithms react to these various factors provides useful insights.

We also used two demand scenarios. The high

demand scenario has a starting inventory of 124 seats, corresponding to a 125% demand factor (ratio of expected total demand to capacity) and approximately a 95% load factor (ratio of average number of seats sold to capacity) under optimal protection levels. The low demand scenario starts with 164 seats, resulting in a demand factor of 95% and a load factor of approximately 90% under optimal protection levels. (Both scenarios are set with somewhat higher demand factors than normally encountered in practice to highlight the revenue impacts of the different methods.) For each scenario, we ran the two methods in parallel on a same sample path of 100 random departures. For the stochastic approximation procedure, we used the gain sequence $\gamma_n = 200/(10 + n)$, which appeared to provide good performance on a range of examples.

On the same sample path, we also ran the optimal policy; that is, a policy that applies the fixed protection level θ^* on each realization. This provided a benchmark for revenue performance. For simplicity, we also assumed observed demand was censored by the protection levels but not the capacity constraint.

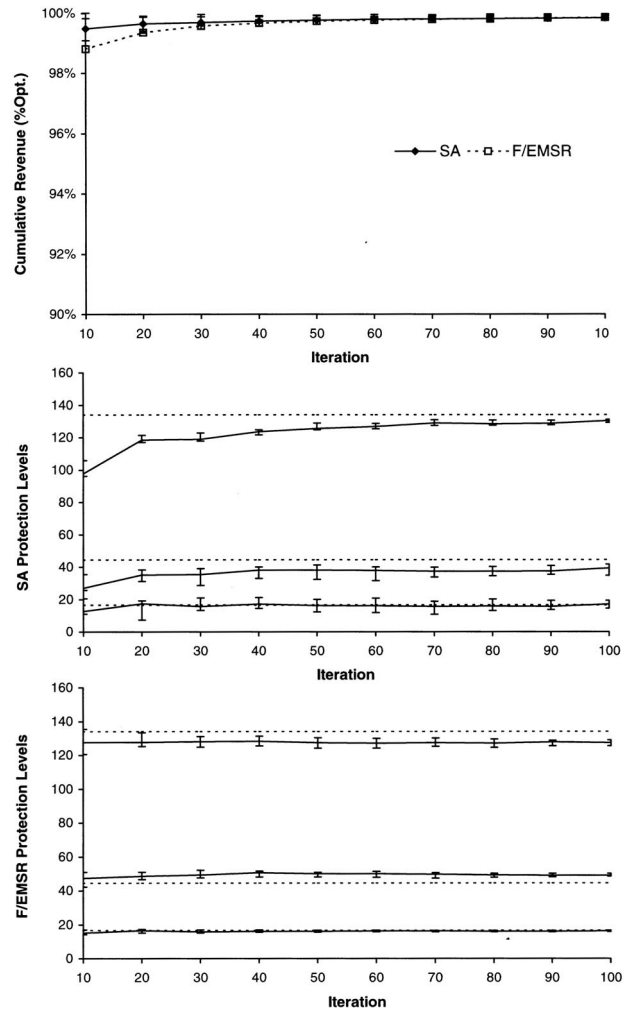
5.3. Numerical Results

The first set of simulations used normally generated demand with parameters given in Table 1. Thus, the actual demand distribution is consistent with the assumptions made in the forecasting and optimization procedures. Each method was tested against the same sample path (i.e., the simulations are coupled). We performed 64 simulations of each 100-flight sample path for each case and tracked the protection levels and revenue performance over time.

Figure 1 shows three graphs of the data for the first case of low demand and low starting values. To generate these graphs, we sampled revenue and protection level data at every tenth iteration during the progression of the algorithms (e.g., at values $n = 10, n = 20, \dots$). This sampling was performed for each of the 64 sample paths, and summary statistics were then computed to illustrate the typical evolution of revenues and protection levels over time.

The top graph of Figure 1 shows the average cumulative revenue as a percentage of the optimal revenue for the two methods as a function of the number of iterations (flights). The error bars show the 95% con-

Figure 1 Low Demand, Low Start, Normal Distribution



fidence intervals about these averages. The middle graph shows the average protection levels over time for the stochastic approximation (SA) procedure. The horizontal dotted lines are the optimal protection levels. The lowest line corresponds to θ_{1}^* , the middle line to θ_{2}^* and the top line to θ_{3}^* . The solid lines are the corresponding average protection levels produced by the stochastic approximation (SA) method. The error bars on the solid lines give the 25th percentile and 75th percentile values for each protection level at each iteration, which provides some sense of the variability in protection levels across sample paths. The bottom graph shows the identical plot of protection levels for the F/EMSR method.

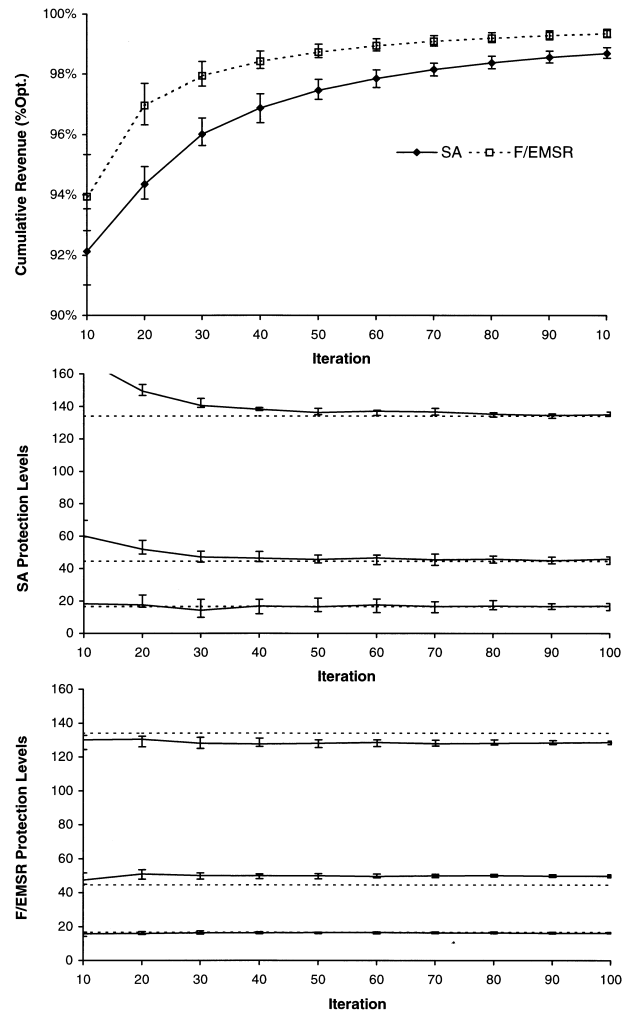
Note from Figure 1 that, in this first case, both procedures have nearly identical (and very close to optimal) cumulative revenue performance, although SA is slightly better early on.⁵ Note that the F/EMSR procedure quickly reaches stable protection levels; however, the second and third protection levels deviate from the optimal ones. This is consistent with the known nonoptimality of EMSR levels beyond the first level (see Table 1). The SA procedure takes longer to converge. In particular, the third protection level (the top line in the graph) is the slowest to converge. This behavior is consistent with the bounds on convergence rate developed in Theorem 1. Also, when the protection levels converge, there is minimal deviation from optimality, which is consistent with the theoretical results as well.

Figure 2 shows the results for the same low demand factor, but with starting protection levels that are all higher than the optimal levels (see Table 2). As in Figure 1, the F/EMSR procedure converges more quickly than the SA procedure. However, in this case the faster convergence of the F/EMSR has a more significant impact on the cumulative revenue performance: F/EMSR generates about 2%–3% higher revenue on average in the early iterations.

Note also that the absolute revenue performance of both procedures is considerably worse in this case compared with the previous case of low initial protection levels, especially in the early iterations. With low demand, overprotecting seats is worse than underprotecting them, and thus erring on the side of low initial starting protection levels is preferred.

The results are quite different in the high-demand factor case. Figure 3 shows the simulation results for low starting protection values and high load factor. Note as indicated by the error bars in the bottom graphs that the F/EMSR procedure is very volatile and somewhat slow to converge in the early iterations. The revenue effect of this behavior is quite significant, with F/EMSR generating cumulative revenues

Figure 2 Low Demand, High Start, Normal Distribution

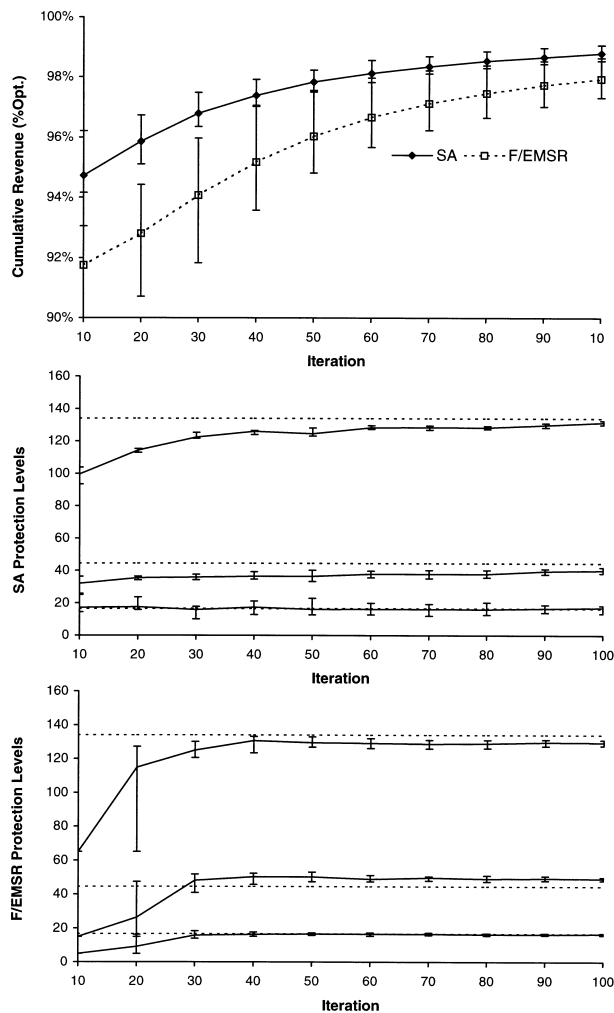


roughly 8% lower than optimal and 2%–3% lower than SA in the early iterations. However, the performance and protection levels of F/EMSR improve after about 30 iterations. In contrast, the SA procedure is considerably more stable and it converges faster in the early iterations, which accounts for its superior revenue performance.

F/EMSR performs badly in this case because the forecasting procedure suffers from the frequent censoring caused by a combination of low protection levels and high demand. As mentioned above, when all observations are censored the forecasting procedure produces unbounded estimates of the demand means and standard deviations. The F/EMSR proce-

⁵ The fact that F/EMSR has lower cumulative revenues at $n = 10$ despite the fact that the protection levels look close to optimal is due to the behavior of its protection levels in iterations 1 to 10, which are not shown in Figure 1. In particular, F/EMSR requires several iterations to produce a finite forecast.

Figure 3 High Demand, Low Start, Normal Distribution



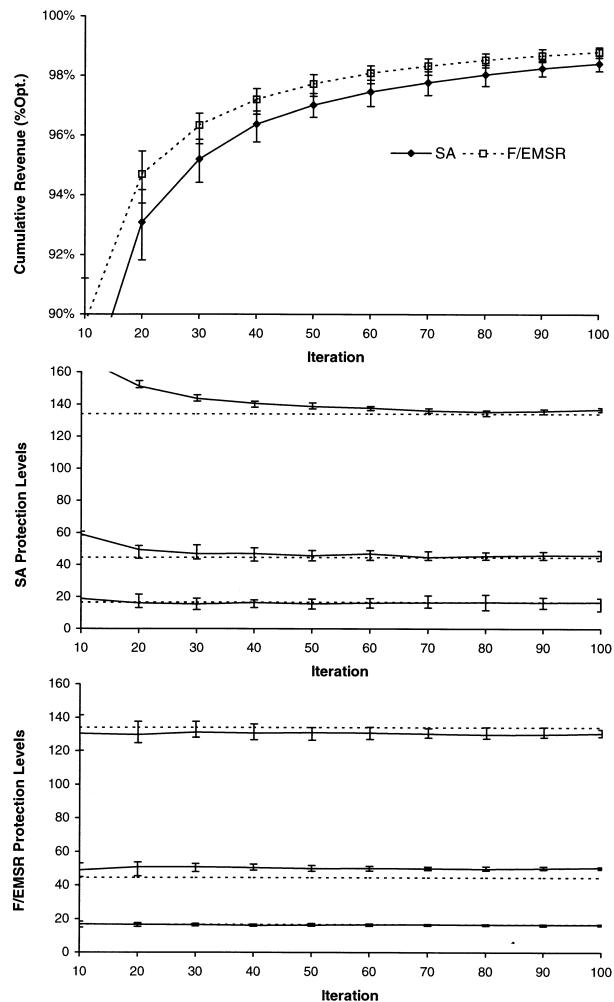
procedure therefore operates with the initial low protection levels for many iterations until a few uncensored observations of demand are obtained and bounded estimates can be computed.

This behavior suggests a potential application of the SA procedure. Namely, SA may prove useful as a means of automatically adjusting protection levels in the early life of new flights when very little demand information is available and forecasting is difficult due to a high degree of censoring. In such cases, the SA method not only provides a robust way to adjust protection levels, but it also serves to speed up the forecasting method itself by nudging protection levels

in the right direction, thereby reducing the amount of censoring.

When the starting protection levels are high and the demand factor is high, the situation is reversed; F/EMSR performs better than SA as shown in Figure 4, though, compared with Figure 3, the absolute performance of both methods is worse overall than the previous case (i.e., overprotecting is worse than underprotecting). In this case with high protection levels, there is little censoring of the higher fare classes and the forecasting procedure quickly produces good estimates of the means and standard deviations. Thus, by Iteration 10 the F/EMSR is able to achieve near

Figure 4 High Demand, High Start, Normal Distribution



optimal protection levels. The SA procedure, in contrast, suffers from slower relative convergence.

In summary, the SA procedure appears to have some distinct advantages in cases where there is a high degree of initial censoring. However, in cases where censoring is less of a problem, the slower convergence of the SA method is a weakness and causes it to underperform F/EMSR. These findings are certainly intuitive and give some sense of the relative strengths and weaknesses of each approach.

We also tested these same four cases with random demand drawn from a log-normal distribution with the same mean and variance. Because the adaptive algorithm requires no distributional assumption, our hypothesis was that SA would perform relatively better in this case. However, the actual simulation results were nearly identical to those for the normal demand case, with the exception that the F/EMSR method had slightly greater deviations from the optimal thresholds. The deviations, however, were not large enough to significantly affect the overall revenue performance. Of course, it is possible that a distribution that is more extreme than the log-normal (e.g., a bimodal distribution) might introduce significant errors in a forecasting and optimization procedure that assumes normality.

6. Conclusions

Adaptive procedures for yield management are attractive because they are simple and robust; however, our preliminary numerical studies indicate that the method has mixed performance—underperforming traditional forecasting and optimization methods when demand is not highly censored but outperforming traditional methods when demand is heavily censored. This behavior suggests that, for airlines with existing revenue management systems, adaptive algorithms may be most useful not to replace, but to augment, traditional forecasting and optimization approaches. Thus, for example, an adaptive approach could automate short-run updating of protection levels in cases where forecasts are highly unreliable or dramatic changes in the market are taking place. Adaptive methods may also be appropriate for small or start-up airlines that lack the resources required to

develop and maintain a full revenue management system.

The adaptive algorithm can also be used as a simple, simulation-based method for computing optimal protection levels within the optimization stage of a traditional forecasting and optimization system. This approach is quite similar to Robinson's (1991) Monte Carlo method for determining optimal protection levels, but offers the potential of greater data efficiency (though most likely slower convergence).

While our paper shows how to construct an adaptive algorithm and provides theoretical guarantees on long-run performance, more research is needed. For example, we have not tested the performance of this approach with large numbers of fare classes and have not generalized the method for nonstationary demand or protection levels that are modified during the booking process. Moreover, it may be that the short-run, transient performance of the method is more relevant in practice than its convergence and long-run performance.

Finally, we believe that our approach may prove useful for studying the *process* of forecasting and optimization over time. Indeed, one can view forecasting and optimization as methods of generating "directions" and "step sizes" for updating protection levels. Stochastic approximation theory may prove a useful theoretical framework for studying the convergence properties of a wide class of forecasting and optimization methods.⁶

⁶ Research supported in part by the Natural Sciences and Engineering Research Council of Canada NSERC OGP0138093.

References

- Belobaba, P. P. 1987a. Air travel demand and airline seat inventory management. Unpublished Ph.D. Dissertation. MIT, Cambridge, MA.
- . 1987b. Airline yield management: An overview of seat inventory control. *Trans. Sci.* **21** 63–73.
- . 1989. Application of a probabilistic decision model to airline seat inventory control. *Oper. Res.* **37** 183–197.
- Benveniste, A., M. Metivier, P. Priouret. 1990. *Adaptive Algorithms and Stochastic Approximation*. Springer-Verlag, Berlin, Germany.
- Blum, J. 1954. Multivariate stochastic approximation methods. *Ann. Math. Statist.* **25** 737–744.
- Brumelle, S. L., J. I. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Oper. Res.* **41** 127–137.

- , ——, T. H. Oum, K. Sawaki, M. W. Tretheway. 1990. Allocation of airline seats between stochastically dependent demand. *Trans. Sci.* **24** 183–192.
- Curry, R. E. 1989. Optimal airline seat allocation with fare classes nested by origins and destinations. *Trans. Sci.* **24** 193–204.
- Etschmaier, M. M., M. Rothstein. 1974. Operations research in the management of the airlines. *Omega* **2** 160–175.
- Lawless, J. F. 1982. *Statistical Models and Methods for Lifetime Data*. John Wiley and Sons, New York.
- Lee, A. O. 1990. Airline reservations forecasting: Probabilistic and statistical models of the booking process. Unpublished Ph.D. Dissertation, MIT, Cambridge, MA.
- Littlewood, K. 1972. Forecasting and control of passengers. *Proc. 12th AGIFORS Sympos.*, American Airlines, New York, 95–117.
- Lukacs, E. 1975. *Stochastic Convergence*. Academic Press, New York.
- McGill, J. I., G. J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Trans. Sci.* **33** 233–256.
- Robbins, H., S. Monro. 1951. A stochastic approximation method. *Ann. Math. Statist.* **22** 400–407.
- , D. Siegmund. 1971. A convergence theorem for non-negative almost supermartingales and some applications. J. Rustagi, ed. *Optimizing Methods in Statistics*. Academic Press, New York, 235–257.
- Robinson, L. W. 1991. Optimal and approximate control policies for airline booking with sequential non-monotonic fare classes. *Oper. Res.* **43** 252–263.
- Smith, B. C., J. F. Leimkuhler, R. M. Darrow. 1992. Yield management at American Airlines. *Interfaces* **22** 8–31.
- Weatherford, L. R., S. E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking and pricing. *Oper. Res.* **40** 831–844.
- Wollmer, R. D. 1992. An airline seat management model for a single leg route when lower fare classes book first. *Oper. Res.* **40** 26–37.

Accepted by Linda V. Green; received October 1997. This paper was with the authors 9½ months for 2 revisions.