APPLICATION OF A PROBABILISTIC DECISION MODEL TO AIRLINE SEAT INVENTORY CONTROL

PETER P. BELOBABA

Massachusetts Institute of Technology, Cambridge, Massachusetts

(Received August 1987; revisions received April, September 1988; accepted November 1988)

The application of booking limits on the number of seats available at different prices on the same flight allows airlines to increase revenues. Effective seat inventory control by an airline depends on forecasts of future bookings, the revenue values associated with each fare type, and an ability to make systematic tradeoffs between booking requests so as to maximize total flight revenues. This article describes the implementation of a computerized system for making these tradeoffs and setting booking limits on future flights at Western Airlines in early 1987. The Expected Marginal Seat Revenue (EMSR) decision model developed for this application takes account of the uncertainty associated with estimates of future demand as well as the nested structure of booking limits in airline reservations systems. The Automated Booking Limit System implemented at Western made use of the EMSR model to set and revise booking limits periodically prior to flight departure. Although the system did not take into account several important components of the seat inventory control problem, a revenue impact test on a sample of actual flights showed a significant revenue improvement over the judgmental methods used previously.

Airlines charge different prices for identical seats on the same aircraft flight by defining different *fare products* through the application of restrictions on ticket purchase and travel. In airline reservations systems, limits are placed on the number of seats available in each fare class or booking class, which can contain several fare products. Effective application of fare class booking limits allows airlines to generate incremental revenues without incurring corresponding increases in operating costs. Controlling the mix of fare products sold can translate into revenue increases of $200 to $500 million for carriers with total revenues of $1 to $5 billion (Fuchs 1987), and many airlines are in the midst of developing methods to tap this potential.

This article examines the application of a probabilistic decision model to the problem of seat inventory control, implemented as part of an automated system for setting booking limits at Western Airlines. First, the problem from the airline's perspective is described, followed by a brief review of past applications of mathematical approaches to the problem. The Expected Marginal Seat Revenue (EMSR) model developed for this application is then presented, including an overview of extensions developed by the author but not implemented. The characteristics of the Automated Booking Limit System (ABLS) and a revenue impact test conducted in early 1987 are described. Finally, the test results and their impacts on the airlines seat inventory management process are discussed. The lessons for further system development conclude this article.

1. Seat Inventory Control

Whether an airline calls it *yield* management or, more appropriately, *revenue* management, efforts to manage the revenue mix of passengers carried involve both pricing and seat inventory control. Although pricing has a direct impact on revenues, an airline can seldom impose price changes without taking the reactions of its competitors into account. Seat inventory control, on the other hand, is a tactical component of revenue management that is entirely under the control of each individual airline. Seat inventory control has the potential of increasing total revenues on a departure-by-departure basis, something that would be far more difficult through pricing actions.

Given a commitment to operate a scheduled flight with a fixed operating cost, and acknowledging the

Subject classification: Decision analysis, applications: Airline seat inventory control.

0030-364X/89/3702-0183 $01.25

very low marginal costs of carrying additional passengers on that flight, an airline that maximizes total flight revenues will, in fact, maximize operating profits. The seat inventory control problem, then, is to determine the number of seats to make available to each fare class from a common or shared inventory (i.e., the coach cabin of the aircraft) so as to *maximize total expected revenues* for a scheduled future flight leg departure. The flight leg approach to seat inventory control, currently used by most airlines, represents an effort to maximize flight leg revenues, not necessarily total system revenues.

At the level of the individual flight leg, seat inventory decisions must be made within the constraints imposed by the airline's network, schedule and reservations system capabilities. The aircraft type to be used for a particular flight departure is known and, in turn, the total number of seats available in the coach cabin can be regarded as fixed. Furthermore, in most instances, the fare products (and thus the fare classes) as well as their respective prices can be assumed to be given and constant throughout the booking period for the flight.

Finding the optimal limits on the number of bookings that may be accepted in a particular fare class on a future flight leg requires estimates to be made of both the expected demand for each fare class and the average revenue associated with each class. The demand for a fare class is considered to be the total number of passengers that will request a seat on a particular flight leg and fare class as part of their origin-destination itinerary. Whether these estimates are based entirely on historical patterns or derived from a forecasting model, data from past flights are required. For a forecast of leg demand, information on booking levels prior to departure and actual boardings by fare class, flight leg, and day of the week must be extracted from the reservations system and stored for seat inventory control decision support purposes.

The seat inventory control process for a future flight departure involves setting initial booking limits on each fare class that must share a common inventory of seats, monitoring actual bookings relative to these initial limits, and then adjusting fare class limits as bookings are accepted. While monitoring is currently the most automated step, the tasks of setting and adjusting booking limits by fare class remain dependent on ad hoc human judgment rather than systematic analysis and decision making. As described in this article, it is possible to make this process more systematic with the application of quantitative decision tools.

2. Mathematical Approaches

Although some theoretical work dealing with airline seat inventory control has been published, the development and implementation of practical models for determining the number of seats to make available in each fare class on a future flight simply did not keep pace with the rapid changes in airline marketing and pricing practices that have transpired since deregulation. This section summarizes briefly the mathematical approaches that have been proposed in the literature. A more detailed review of past work was presented by Belobaba (1987a).

The relationship between fare class inventories or reservations *buckets* in an airline's reservations system affects the way in which the seat inventory control problem is represented mathematically and, in turn, the solution methods that are most appropriate. The simplest reservations system structure involves distinct and separate inventories for each fare class. The booking limits on each bucket must sum to the total capacity of the shared cabin in such systems. When overbooking is involved, the booking limits on each bucket must sum to the overall limit on reservations for the shared cabin. In contrast, a *nested* reservations system is one in which the fare class inventories are structured such that a high fare request will not be refused as long as any seats remain available in lower fare classes. A nested reservations system is thus binding in its limits on lower fare classes, but its limits are "transparent" from above (for higher fare classes).

The distinction between separate and nested fare class inventories is important both to the way in which the seat inventory control problem is represented and to the mathematical methods that are used to determine optimal seat allocations or fare class booking limits. Equally important is the distinction between the static problem, in which fare class booking limits are applied at the start of the booking process for a future flight, and the dynamic problem, in which booking limits may be revised as actual bookings are accepted. In the latter case, the length of the interval between revisions determines the degree to which the differences between distinct and nested fare class inventories will affect optimal booking limits. Request-by-request revision of limits on distinct inventories eliminates the expected revenue advantages associated with nested inventories. The longer the interval between revisions, however, the greater the importance of finding the optimal booking limits that apply to nested fare classes.

The tools of differential calculus, Lagrangian multipliers, mathematical programming and network optimization have been applied to the allocation of a given number of seats among two or more distinct fare class inventories, depending on the complexity of the particular problem representation involved. When the multiple fare class problem is formulated as a constrained revenue maximization problem, the optimality conditions are

$$\frac{\partial \bar{R}}{\partial S_i} = \frac{\partial \bar{R}}{\partial S_j} = \lambda \quad \text{for all fare classes } i \neq j \tag{1}$$

where λ, the Lagrangian multiplier, equals the expected marginal revenue for the last seat, S, allocated to each fare class. Seats are allocated among fare classes such that the total expected marginal revenue with respect to seats allocated is equal across all relevant fare classes.

Further expansion of the problem to multiple fare classes and multiple-leg flights or even connecting flight operations requires the application of mathematical programming and related techniques to find the optimal fare class seat allotments numerically. These methods use the principle of equating marginal revenues of the last seats allocated to each inventory in determining the revenue-maximizing seat allotments.

The potential application of mathematical programming techniques to the seat allocation problem was considered by Mathaisel and de Lamotte (1983) at MIT. The network flow approach was pursued by Glover et al. (1982), who developed a network based seat allocation model for Frontier Airlines. Other analysts have addressed the problem of incorporating probabilistic demand into mathematical programming formulations of the seat allocation problem. McDonnell-Douglas analysts proposed a formulation of the single-leg seat allocation problem that makes use of binary decision variables in a linear integer programming framework (Wollmer 1985). D'Sylva (1982) of Boeing Aircraft used a piecewise linear approximation of the expected revenue curve in a linear programming formulation to extend Glover's algorithm to include stochastic demand.

Making seat inventory decisions dynamically with the help of mathematical programming techniques requires an assessment of the value of accepting a current reservation request relative to the decrease in expected total revenue associated with removing one seat from the available inventory on the flight leg(s) requested. This comparison of the expected revenue differential for each incremental unbooked seat with a "certain" revenue from the current reservations request is extendable in conceptual terms to the most complex network formulations. Given that in practice most airlines only make revisions to fare class booking limits periodically, the solution to the static seat allocation problem must apply over the entire interval between revisions. For applications in which this interval is substantial, mathematical programming approaches will not necessarily give the true revenue-maximizing seat allotments for a *nested* reservations system.

Finding the optimal booking limits for a nested reservations system involves applying the optimizing principle of equating marginal seat revenues to the problem of dynamic booking limit revision and incorporating it into iterative solution approaches. Littlewood (1972) applied the marginal seat model to a dynamic reservations context for the two-class, single flight leg seat inventory control problem. He suggested that revenues could be maximized by "closing down" the low fare class to additional bookings when the certain revenue from selling another low fare seat is exceeded by the *expected* revenue from saving that seat for a potential high fare passenger.

Applications of this *marginal seat* principle succeeded in incorporating probabilistic demand explicitly into the seat inventory revenue maximization problem for a flight leg. The simple decision rule presented by Littlewood, as well as by Bhatia and Parekh (1973) and Richter (1982), determined optimal fare class limits for two fare classes on a single flight leg. The same marginal seat principle can be applied to the static problem when *nested* fare class inventories are involved.

3. Expected Marginal Seat Revenue Model

Determining the booking limit on each fare class that will maximize total revenues for a future scheduled flight departure is a dynamic process. The static problem is to establish these fare class limits at the start of the booking process, taking into account the uncertainty associated with expected bookings by fare class, to the extent possible. The dynamic problem is to revise these initial limits on the basis of the additional information provided by actual bookings as departure day approaches. The Expected Marginal Seat Revenue (EMSR) model recommends fare class booking limits, taking into account the probabilistic nature of future demand for a flight. The EMSR formulations developed for application at Western Airlines are presented briefly in this section. A more detailed description can be found in the author's doctoral dissertation (Belobaba 1987b).

The nature of price sensitive travel and the application of advance purchase restrictions on the lowest priced fare products mean that the lowest fare classes tend to book up first, well before the majority of requests for the highest fare class are received. The seat inventory control problem therefore is to determine how many seats *not to sell* in the lowest fare classes and to retain for *possible* sale in higher fare classes closer to departure day. A decision model therefore must find the *protection levels* for higher fare classes which can be converted into *booking limits* on lower fare classes. In a nested fare class reservations system, each booking limit is the *maximum* number of seats that may be sold to a fare class (including all lower fare classes with their own, smaller booking limits). The booking limit on the highest fare class is, thus, the total capacity of the shared cabin. The protection level for the highest fare class is the difference between its booking limit and the booking limit of the next lowest class.

The seat allocation approaches in past works assumed independent fare class demand densities to correspond with the distinct fare class inventory assumption. This assumption of no relationship between demand levels for different fare classes is retained. Furthermore, we assume initially that a consumer denied a flight/fare class request represents a booking loss to the airline. On the other hand, an accepted booking represents certain revenue for the airline, as we assume that no booking cancellations or passenger no-shows occur.

There exists uncertainty about the ultimate number of requests that an airline will receive for seats on a future flight and, more specifically, for the different fare classes offered on that flight. The total demand for a particular flight, on average, fluctuates systematically in cycles described by day of the week and season of the year. There also will be stochastic variation in demand around the expected values, among similar flights sampled consistently over a homogeneous period of time. This stochastic demand for a future flight departure can be represented by a probability density function. Past analyses generally have assumed a Gaussian (normal) distribution of total demand for a flight, with means and variances that depend on the market being studied and on the nature of its traffic (Belobaba 1985). We define $p_i(r_i)$ to be the probability density function for the total number of requests for reservations, r_i, received by the airline for seats in fare class i by the close of the booking process for a scheduled flight leg departure.

The number of seats allocated to a particular fare class, S_i, might not exceed the number of actual requests for that fare class, resulting in rejected demand, or *spill.* Thus, we can define a cumulative probability that all requests for a fare class will be accepted as a continuous function of S_i:

$$P_i(S_i) = P[r_i \leq S_i] = \int_0^{S_i} p_i(r_i)\, dr_i. \tag{2}$$

Conversely

$$\begin{aligned} P[r_i > S_i] &= \int_{S_i}^{\infty} p_i(r_i)\, dr_i \\ &= 1 - P_i(S_i) = \bar{P}_i(S_i). \end{aligned} \tag{3}$$

The probability of receiving more than S_i requests for fare class i, or the probability of spill occurring, is therefore $\bar{P}_i(S_i)$.

We define EMSR_i to be the expected marginal seat revenue for class i when the number of seats available to that class is increased by one. The expected marginal seat revenue of the S_ith seat in fare class i, $\text{EMSR}_i(S_i)$, is simply the average fare level in that class multiplied by the probability of selling S_i or more seats:

$$\text{EMSR}_i(S_i) = f_i \cdot \bar{P}_i(S_i). \tag{4}$$

Note that $\text{EMSR}_i(S_i)$ depends directly on $\bar{P}_i(S_i)$, the probability that the S_ith seat made available to class i will be sold.

Consider a single-leg flight for which bookings will be accepted in two nested fare classes, 1 and 2, having average fare levels f_1 and f_2, respectively. In order to maximize total expected flight revenues, the reservations process should give priority to class 1 passengers. Class 1 will have the total available capacity of the shared cabin, C, as its booking limit, BL_1. The seats protected from class 2 and available exclusively to class 1 will be denoted S_2^1. The optimal protection level S_2^1 for class 1 is the value of S_2^1 that satisfies the condition

$$\text{EMSR}_1(S_2^1) = f_2. \tag{5}$$

Graphically, the optimal value of S_2^1 is the point at which the $\text{EMSR}_1(S_1)$ curve intersects f_2, as shown in Figure 1. The optimal booking limit on class 2 is BL_2, the difference between the capacity of the shared cabin, C, and the optimal protection level, S_2^1.

This solution will maximize *expected* revenues in cases where the booking limit is set at the start of the reservations process (static seat inventory control). A class 2 request will be rejected only when BL_2 is reached, at which point, the expected revenue for all remaining seats will be greater than the class 2 average fare. If class 2 requests never reach BL_2, the unsold seats will be available for unexpectedly high class 1

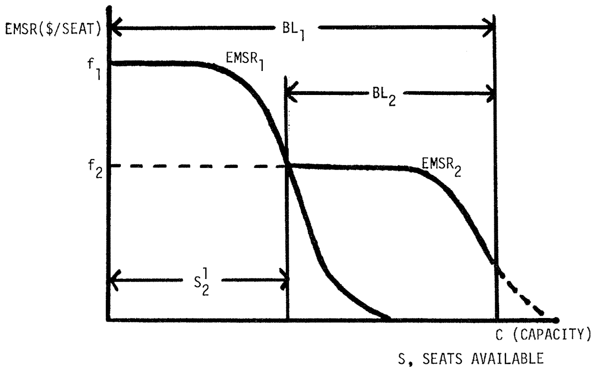

Figure 1. Maximizing expected revenues for the 2-class example.

demand. In any event, the *expected* revenue per seat from class 1 requests in excess of S_2^1 is below f_2, in the absence of additional information on actual bookings for the future flight being managed.

Extending this approach to multiple fare classes on a single flight leg simply requires that more comparisons of expected marginal revenues be made among the relevant classes. In the general case of k fare classes offered on a flight leg, the optimal values of S_j^i must satisfy

$$\mathrm{EMSR}_i(S_j^i) = f_j, \quad i < j, \quad j = 1, \ldots, k. \tag{6}$$

The total number of comparisons required for k nested fare classes is given by

$$\frac{k(k-1)}{2}. \tag{7}$$

These protection levels, in turn, determine the booking limits on each fare class j:

$$\mathrm{BL}_j = C - \sum_{i<j} S_j^i. \tag{8}$$

That is, all seats with an expected marginal revenue greater than f_j should be held back from sale to class j. Otherwise, any request for a class j seat may be accepted. It is possible that one or more values of BL_j derived from these equations might be negative, in which case, class j should not be offered at all if expected revenues are to be maximized. In such a case

$$\mathrm{BL}_j = \max\left[0,\ C - \sum_{i<j} S_j^i\right]. \tag{9}$$

The incremental number of seats protected for class j is the *nested protection level* for class j, denoted NP_j. The nested protection level for class j thus is given by

$$\mathrm{NP}_j = \mathrm{BL}_j - \mathrm{BL}_{j+1}. \tag{10}$$

In the EMSR framework, the lowest fare class does not have a protection level per se, but rather a booking limit equal to the number of seats that remain after all upper classes have been protected. With a capacity of C seats and k fare classes, then, the values of NP_j must satisfy

$$C = \sum_{j<k} \mathrm{NP}_j + \mathrm{BL}_k. \tag{11}$$

The EMSR protection levels S_j^i, nested protection levels NP_j, and booking limits BL_j are shown graphically for a three-class example in Figure 2.

EMSR booking limits can be generated for any number of nested fare classes with no change to the model's basic structure. It is important to recognize that not all reservations systems are nested in this way, or they might not be nested at all. Nesting is preferable in seat inventory management because there is no difference in the physical seats or the on-board service being sold to different fare classes. Airlines without nested fare class systems are denying themselves the flexibility of accommodating unexpectedly high demand levels in high-fare classes and, in turn, are losing potential revenues.

4. Dynamic Application of the EMSR Model

The EMSR decision framework can be applied to a seat inventory control context in which booking limits may be revised on a regular basis as the flight departure day nears. In such a situation, additional information is available in the form of actual bookings already accepted for the future flight. Because an actual booking in any fare class will translate (barring cancellations and no-shows) into a revenue passenger occupying a seat, incorporating actual bookings into the

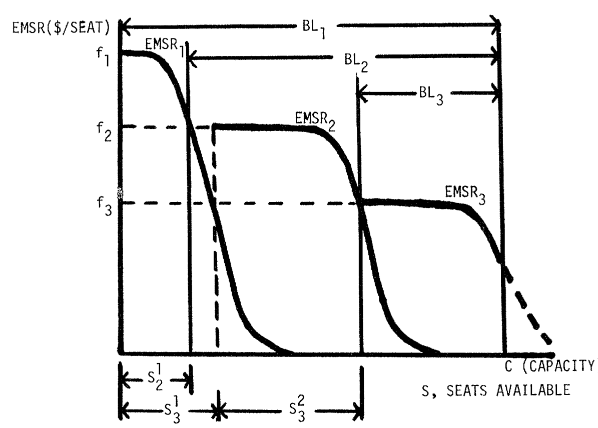

Figure 2. EMSR solution for the nested 3-class example.

EMSR decision framework can reduce the uncertainty associated with the estimates of expected demand used as input.

In the static problem, the EMSR model requires an estimate of the *total* expected requests by class. In the dynamic case, estimates of future requests at various times before departure are required to calculate optimal protection levels for the unbooked seats still available for flight. Dynamic application of the EMSR framework involves repetitive use of the static model described in the previous section, but with revised input data. The objective is to determine the optimal fare class limits for the time period remaining to departure, irrespective of the (non) optimality of the booking decisions already made. Thus, each EMSR calculation in the dynamic case is based on a static assessment of expected fare class revenues from that point in time, based on the most recent available demand information for the flight leg.

For the dynamic problem, the estimates of bookings to come must be generated from densities of demand by fare class from an historical sample of requests made between day t before departure and the day of departure. We define r_i^t to be the number of requests made for class i between days t and 0 before departure, meaning:

$$r_i^t \leq r_i \tag{12}$$

by definition. The probability density of requests from day t onward is $p_i(r_i^t)$, and the probability of receiving S or more requests for class i in the time remaining to departure is $\bar{P}_i^t(S)$.

On any day t prior to flight departure, the inputs required by the EMSR model are the average fare or revenue levels for each fare class, f_i, which may or may not remain constant over the booking period, and the estimates of $\bar{P}_i^t(S)$ for all relevant values of S, derived from the $p_i(r_i^t)$ densities. The optimal seat protection level for class 1 relative to class 2 for the period remaining before departure is $S_2^1(t)$, such that

$$\text{EMSR}_1^t[S_2^1(t)] = f_1 \cdot \bar{P}_1^t(S_2^1) = f_2. \tag{13}$$

This protection level for day t can be used to find the revised optimal booking limit on class 2, as

$$\text{BL}_2(t) = C - b_1^t - S_2^1(t) \tag{14}$$

where b_1^t is the number of bookings already accepted in class 1 up to day t before departure. The maximum number of seats still available is $C - b_1^t$, and $S_2^1(t)$ of these seats are protected for class 1. In essence, actual bookings are *protected* along with the additional seats required to accommodate expected bookings to come in a revenue-maximizing manner.

For more than two fare classes, comparisons between the EMSR(S) values of all upper classes relative to the average fare levels of lower classes are required, as before. These comparisons involve the demand densities of future requests from the current day t. Each comparison of a higher fare class i with a lower fare class j generates an optimal value of $S_j^i(t)$ that satisfies

$$\text{EMSR}_i[S_j^i(t)] = f_i \cdot \bar{P}_i^t(S_j^i) = f_j. \tag{15}$$

The revised booking limits for day t take actual bookings into account:

$$\text{BL}_j(t) = C - \sum_{i<j} S_j^i(t) - \sum_{i<j} b_i^t. \tag{16}$$

The nested protection levels, $\text{NP}_j(t)$, for successively lower fare classes are derived as in the initial case, but with actual bookings included:

$$\begin{aligned}\text{NP}_j(t) &= \text{BL}_j(t) - \text{BL}_{j+1}(t)\\ &= \sum_{i\leq j} S_{j+1}^i(t) - \sum_{i<j} S_j^i(t) + b_j^t.\end{aligned} \tag{17}$$

As before, $\text{BL}_j(t)$ is constrained to be greater than or equal to zero. It also is constrained in this case to be no lower than the actual number of bookings already accepted in class j and all lower classes up to day t. Because requests for lower fare classes are generally received earlier in the booking process than full-fare requests, it is possible that the revised $\text{BL}_j(t)$ derived from the EMSR calculations will be lower than the number of bookings already on hand in classes j and lower. With no possibility of cancellations or no-shows assumed, the revised booking limit on class j then becomes

$$\text{BL}_j(t) = \max\left[C - \sum_{i<j} S_j^i(t) - \sum_{i<j} b_i^t, \sum_{k\geq j} b_k^t, 0\right]. \tag{18}$$

The EMSR framework thus can be used to determine optimal protection levels and recommended booking limits in a nested multiple fare class reservations system for a single, future flight leg. Initial booking limits may be derived on the basis of estimates of total expected requests for a future flight, before the reservations process begins. These limits may then be revised dynamically during the booking process, taking into account both actual bookings and estimates of future requests by fare class.

This dynamic application of the EMSR decision model retains the simplistic assumptions with respect to demand densities, refused requests, cancellations and no-shows made in previous works. Furthermore, we added the assumption that there is no relationship between the booking rates of different fare classes or among time periods before departure. Forecasting models that include such a relationship can be used to generate estimates of future demand as a function of actual bookings, should that be necessary.

5. Passenger No-Shows and Choice Behavior

The development of the EMSR decision framework has overlooked, to this point, the possibility that a reservation might not translate into a passenger being carried on the flight in question. Furthermore, the assumption that the demand for each fare class is distinct does not account for passengers deciding to accept a reservation in a higher fare class when the desired lower fare class is sold out. This section deals briefly with each of these complications to the simple demand patterns assumed above. Neither of the enhancements to the EMSR model outlined below were implemented at Western Airlines before the revenue impact test was performed.

With each accepted booking assumed to represent certain revenue, the basic model focuses on the problem of managing the physical seat inventory on the aircraft. There is some probability, however, that a booked passenger might not be carried on the flight for which the reservation was made. Whether the original booking is cancelled prior to departure or the passenger simply fails to appear at departure time, the outcome from the airline's perspective is the loss of a revenue passenger for that particular flight. This loss is the opportunity cost of having removed a seat from the available inventory and then not receiving any revenue for it.

The airline industry practice of controlled overbooking of flights above the physical capacity of the aircraft has evolved with the objective of minimizing such costs. Overbooking analysis is performed to determine the extent to which a future flight should be overbooked so as to minimize the sum of the lost revenues associated with empty seats and the costs of denying boarding to passengers with confirmed reservations (Rothstein 1968, 1985). For the purposes of seat inventory control, it is important to recognize the potential interaction between the fare class mix of passengers booked for a flight and overbooking, which could have significant revenue implications for the airline.

The EMSR decision framework can be extended to account for flight overbooking. The details of this extension are included in the author's doctoral dissertation. The objective here is to illustrate how overbooking proportions by flight and even by fare class could be incorporated into the EMSR framework for deriving revenue-maximizing fare class booking limits.

With the introduction of the possibility that a booking made in fare class i will not generate revenue f_i with certainty, the EMSR decision framework must be adjusted. The demand inputs required are still estimates of the densities of *requests* for each fare class. The difference here is that each accepted request (booking) in a fare class cannot be treated as if the revenue associated with that class will always be realized. The expected revenue associated with accepting a request will be lower than the actual fare level in that class because of the possibility of cancellation or no-show. The overbooking percentage for each fare class determines the extent to which the expected revenue from a booking is reduced due to this uncertainty.

We define this overbooking percentage to be the *overbooking factor*, OV, where OV $\geqslant 1.0$, and where OV applies at the time of the EMSR calculations. Given overbooking factors OV_1 and OV_2, the optimal protection level for class 1 from class 2, S_2^1, must satisfy

$$\bar{P}_1(S_2^1) \cdot f_1 \cdot \frac{1}{\mathrm{OV}_1} = f_2 \cdot \frac{1}{\mathrm{OV}_2}. \tag{19}$$

The revenue levels of the two classes, in essence, are deflated by the assumed overbooking factors. Thus, we are reducing the expected marginal revenue of each incremental seat for class 1 for which the possibility of protection is being evaluated.

The value of S_2^1 that satisfies the above condition is the protection level for class 1 from class 2, expressed in terms of the number of *reservations spaces* (as opposed to physical seats) that should be protected for exclusive use of class 1 passengers. Because the demand inputs to the EMSR framework for all fare classes will be in terms of expected requests, the calculated protection levels for all upper fare classes will be in terms of reservations spaces, with the relevant overbooking factors already incorporated. The generalized decision rule for the EMSR framework with overbooking factors is to protect S_j^i seats for class i from class j such that

$$\begin{aligned}\mathrm{EMSR}_i(S_j^i) \cdot \frac{1}{\mathrm{OV}_i} &= \bar{P}_i(S_j^i) \cdot f_i \cdot \frac{1}{\mathrm{OV}_i} \\ &= f_j \cdot \frac{1}{\mathrm{OV}_j}.\end{aligned} \tag{20}$$

The derivation of overbooking limits, BL_j^*, from this revised EMSR decision rule is complicated by the fact that the optimal protection levels are expressed in terms of reservations spaces (overbooking factors included), while the capacity of the aircraft is in terms of physical seats. The simplest case occurs when all the OV_i values are equal across fare classes, so that the OV factors drop out of the above formulation, reducing it to the original EMSR formulation. In the application of the EMSR approach at Western, a single overbooking factor was applied, such that the overbooking limit on the total capacity of the shared aircraft cabin, C^*, is simply

$$C^* = BL_1^* = OV \cdot C. \tag{21}$$

Booking limits for each subordinate fare class in the coach cabin may be derived after an overbooking target is established for the total capacity of the cabin. The protection levels generated by the EMSR model are for *reservations spaces* rather than physical seats, and the *overbooking* limits on each fare class BL_j^* are given by

$$BL_j^* = C^* - \sum_{i<j} S_j^i. \tag{22}$$

The net result is that each fare class may be overbooked by the same percentage, and C^* will be the same regardless of the fare class mix actually booked for any particular flight.

The fact that fare classes are designed to appeal to different air travel demand segments suggests that passengers booked in each fare class might exhibit different no-show behavior. It is plausible that passengers in lower fare classes will be more likely to show up for booked flights than those in higher fare classes for several reasons, the most important being the cancellation penalties associated with the lowest excursion fares. Few airlines currently have the detailed no-show data required to confirm this hypothesis empirically. Nonetheless, if it can be determined that different no-show behavior exists across fare classes, the incorporation of different overbooking factors can make the EMSR model responsive to changes in the fare class mix of bookings accepted.

A second characteristic of fare class demand not addressed by the basic model tested at Western is the disposition of refused reservations requests, which might not always result in a booking loss to the airline. Depending on the individual consumer's choice process, the unavailability of a desired flight and fare class can lead to:

1. a *vertical shift* to a higher fare class, same flight;
2. a *horizontal shift* to a different flight, same fare class and airline;
3. a *booking loss* to the refusing airline.

For the purposes of managing the seat inventory for a single flight leg, the probability of interest is that of a vertical shift in fare classes on the same flight. The probability that a passenger refused a request for fare class i will accept a booking in the next highest fare class $(i - 1)$ is $P_i(v)$. The EMSR formulation can be extended to include the possibility of vertical shift on the part of the refused passenger. Although the probability of horizontal shift, $P_i(h)$, is also important to the airline wishing to maximize system revenues, the focus of the EMSR model on individual flight legs allows us to assume that a horizontal choice shift is equivalent to a booking loss.

When a class 2 request is received by the airline and a booking is accepted, a revenue of f_2 is realized, in the absence of overbooking and no-shows. If the request is refused, the expected revenue associated with the denied passenger accepting a vertical shift in fare classes is

$$P_2(v) \cdot f_1. \tag{23}$$

We want to find the incremental protection level required for class 1 to take into account this potential class 1 revenue when class 2 is closed.

From the basic EMSR formulation, a protection level for class 1 from class 2 of S_2^1 still will be required. Additional seats protected for class 1, V_2^1, can be taken either by a refused class 2 or a class 1 passenger. The expected marginal revenue from a class 1 passenger in the V_2^1th additional protected seat is

$$EMSR_1(S_2^1 + V_2^1) = f_1 \cdot \bar{P}_1(S_2^1 + V_2^1). \tag{24}$$

If the upgrade probability, $P_2(v)$, is greater than zero, the incremental expected revenue associated with potential vertical shifts from class 2 may be realized if the seat is not purchased by a class 1 passenger.

The combined expected marginal seat revenue for the V_2^1th seat protected for class 1 thus is equal to f_1 multiplied by the probability that a class 1 request will be received for that seat *or* a vertical choice shift is accepted, given that BL_2 is reached. The optimal value of V_2^1 must therefore satisfy

$$EMSR_1(S_2^1 + V_2^1) \cdot [1 - P_2(v)] + P_2(v) \cdot f_1 = f_2 \tag{25}$$

where S_2^1 is the protection level for class 1 in the absence of vertical choice shifts, as before. The combined expected revenue from each additional seat protected for class 1 will be greater than or equal to f_2, given that BL_2 is reached. If BL_2 is not reached, this additional protection will have no impact in a nested reservations system.

The total protection level for class 1 from class 2 is then $(S_2^1 + V_2^1)$, and the booking limit on class 2 (without overbooking) is

$$BL_2 = C - S_2^1 - V_2^1. \tag{26}$$

Overbooking factors may be incorporated as before, in which case, the total EMSR protection level $(S_2^1 + V_2^1)$ is treated the same as S_2^1 in the calculation of overbooking limits.

The impact of including one or more $P_i(v)$ values in the EMSR formulation will be an *increase* in the protection levels for each of the higher fare classes. Each lower fare class will see its booking limit decrease by the incremental protection level required to account for the possibility of vertical choice shifts to the next highest fare class. The magnitude of this decrease will depend on the relative magnitudes of the $P_i(v)$ values estimated or assumed, highlighting the importance of this probabilistic element to the EMSR framework. As mentioned, neither fare class specific overbooking factors nor upgrade probabilities were incorporated into the ABLS tested at Western.

6. Automated Booking Limit System

Under a research agreement between the Flight Transportation Laboratory at MIT and Western Airlines, an Automated Booking Limit System (ABLS) was developed and implemented during 1986. The decision model programmed into the system to recommend fare class booking limits for future flight leg departures was based on the EMSR approach described above. By the end of 1986, the EMSR revenue maximization model and a dynamic booking limit adjustment routine had been programmed into ABLS. For a number of reasons, including the fact that Western Airlines would no longer be operating on its own after March 31, 1987, further system development was not possible at Western. Nonetheless, the revenue impact of a completely automated approach to setting fare class booking limits relative to the manual and ad hoc methods used previously was tested for a sample of actual flights during the first 3 months of 1987.

6.1. Data Requirements

The EMSR decision framework requires two types of input data to determine fare class limits for a particular, future flight leg departure—estimates of expected demand and of the average revenue associated with a passenger booking, by fare class. The demand inputs required for initial applications of the EMSR model are estimates of the total number of requests expected for a future flight leg departure, by fare class. Because the model takes into account stochastic variation in demand, an estimate of the variance in total requests around the expected value is also required. For dynamic applications of the EMSR framework, estimates of partial demand by fare class in the form of requests still to come from day t to departure are required, as well as estimates of the variation of this partial demand. Furthermore, the number of actual bookings already accepted for the particular, future flight leg being considered are also necessary demand inputs.

Currently, most airline reservations systems log total bookings by fare class for a future flight leg. Data base management systems have been developed by many airlines to generate extracts of the current booking levels and limits by fare class on a daily basis. These extracts become part of an evolving historical data base of total bookings by day before departure for all flight legs for which reservations are being accepted. Once the flight has departed, a complete booking history for that flight leg can be retrieved from the data base.

The EMSR model requires estimates of the mean and standard deviation of *requests* by fare class. These estimates may be derived from a sample of past operations of the same flight, or similar flights on the same leg. If this historical sample includes only departed flights for which no fare class booking limits were reached during the reservations process (implying that no requests were refused), net booking levels may be used directly in the estimation of requests by fare class.

It is not likely that all the flights in the historical sample will have booked up without reaching one or more of the fare class booking limits. In such cases, one or more requests for a particular fare class and flight are likely to have been refused by the airline, and the disposition of these refused requests cannot be determined from the available data. The net booking levels in the reservations system data base thus represent a *constrained* estimate of total requests for a particular flight leg and fare class. It is possible to use statistical methods to derive unconstrained estimates of requests by fare class, given the booking levels for each observation in the sample and knowledge of whether these booking levels were, in fact, constrained by a fare class booking limit (Boeing 1982).

For both the initial and dynamic applications of the EMSR decision model then, the necessary demand estimates can be derived from the existing data available to most airlines from their reservations systems.

In contrast to the volume of historical booking data that airlines can extract from their reservations systems, the availability of detailed revenue data is limited at many airlines. As an input to the EMSR decision model, the revenue associated with bookings in different fare classes is just as important as the estimates of fare class demand. Whether the data are provided by a comprehensive revenue data base or a sporadic sample of ticket coupons, some estimate of the relative revenue value of passengers booked in each fare class for the future flight under consideration is required.

The demand and revenue estimates used as input by the EMSR decision model, thus, are affected by characteristics of the data being collected and stored by the airline. The greater the extent to which the available data do not correspond to the input needs of the EMSR model, the greater the need for estimation procedures and assumptions to generate the required inputs. For the airline wishing to realize rapid improvements to its seat inventory control process by implementing a decision approach like that of the EMSR model, however, working within the constraints of the available data and using estimation methods might be the only alternative.

6.2. ABLS Development and Implementation

These data availability issues had to be addressed by Western Airlines as it prepared to implement an ABLS for seat inventory management. Although several of the problems identified over the course of the year-long development effort were rectified, there remained numerous policy issues and data availability limitations that could not be resolved. The system developed was, to a large extent, built around existing procedures and capabilities.

The objective of developing ABLS was to make the process of setting and adjusting fare class booking limits for future flight departures more systematic and to automate it as much as possible. It was hoped that the implementation of an automated system would reduce the manual effort required on the part of a relatively small staff of seat inventory control analysts, allowing them to focus their analysis efforts on the small proportion of flights requiring closer attention. Inclusion of the EMSR decision model in the system was intended to provide the analysts with specific recommendations of what the fare class booking limits *should be*, based on a systematic evaluation of the input data.

The system was developed to consist of two parallel components or routines: batch and on-line. The batch routine was designed to set and periodically revise fare class booking limits for all future flight leg departures, based on booking data from a sample of recent departures of the same flight leg on the same day of the week. Revenue averages extracted from a prorated revenue data base for the most recent available sample period also were used as inputs.

The batch routine operated on a day-of-week rotation, calculating bookings limits for all flight legs scheduled to depart on, for example, all future Tuesdays up to 90 days out in a single run. Fare class limits were revised weekly thereafter, up to and including 6 days before departure. These revision runs not only incorporated the most recent input data available, they recalculated the optimal seat protection levels required for expected requests still to come by fare class, as estimated from historical build-up patterns for the same flight leg. Actual bookings were added to these protection levels, from which revised booking limits on each fare class were derived.

The on-line routine allowed the analysts to intervene and to run the EMSR decision model for an individual flight leg and day of the week. It also enabled the analysts to compare the model's recommended booking limits with current limits and actual bookings on hand for each future departure of the flight leg on the same day of the week. The EMSR recommended protection levels and booking limits simply could be overridden manually before being loaded into the reservations system.

ABLS was designed to allow user (analyst) intervention in the application of the EMSR decision model to future flights. This capability was especially important in light of the many imperfections in the system. As of the end of 1986, ABLS did not include a demand forecasting model or an ability to make seasonal adjustments to demand and revenue estimates. Furthermore, no upgrade probabilities were included in the derivation of fare class booking limits. These limitations required that analysts have the capability to override the recommended limits on flight legs where upgrade potential was thought to be significant, and in markets or during periods for which the recent historical data did not provide a valid estimate of the revenue or demand conditions expected for future departures of the same flight leg.

7. EMSR Revenue Impact Test

A performance evaluation of the ABLS at Western Airlines was conducted during the first three months of 1987. In this section, the testing methodology is described and its limitations are discussed. The revenue impacts of EMSR seat inventory control relative to strictly manual methods are then presented, and their implications for further system development are assessed.

7.1. Test Methodology

At the start of the ABLS performance evaluation, the system users (i.e., the seat control analysts) were informed that specific flight legs operating on specific days of the week would have their fare class booking limits set and revised automatically by ABLS. The group of 21 flight leg/day of week combinations for which ABLS was allowed to set and revise fare class booking limits without analyst intervention comprised the BATCH test group.

The analysts were allowed to use the on-line routine for each of the 21 flight legs in the test on another day of the week for each leg, different from the BATCH day of week. Ultimately, interactive use of ABLS was to be the norm, with the combination of automation and human intervention resulting in what should have been higher flight leg revenues than the automated system alone could generate. For the purposes of this test, however, this on-line subset of the test sample was not intended to provide a measure of the joint performance of the system and the analysts.

For each of the 21 flight legs in the test sample, five days of the week remained for which booking limits were set manually by the analysts, based entirely on their own judgment, as it had been for years. One of the remaining days of the week was selected to represent a CONTROL flight for each of the BATCH flights in the test, and was unknown to the analysts. For each flight leg, all departures on one day of the week had fare class booking limits set automatically by the ABLS batch routine, while all departures on another day of the week were managed manually by seat inventory control analysts, thereby providing the control group for comparison purposes. The primary focus of this evaluation of ABLS performance was on the traffic and revenue levels of the BATCH group relative to the CONTROL group.

Several problems with interpreting the results of this test were identified even before the test began. From the outset, the day of the week rotation programmed into ABLS dictated that a day of the week approach to comparative testing be used. Although the days of the week for each flight leg were selected originally on the basis of similar historical loads and booking patterns, unexpected consistent differences in demand for one day in the test could give a systematic advantage for one test group over the other on that flight leg.

A second interpretation problem involved the fact that the effects of different fare class booking limits only can be evaluated for flights that actually reached one or more of these limits in at least one of the test groups. If no fare class limits were reached for a set of flight legs, any difference in loads or revenues could only be attributed to variations in demand between the departed flights. Even though the flight legs in the test sample were selected on the basis of high load factors for historical periods similar to the test period, there was no guarantee that unexpectedly low demand would not be observed on some of the test flights due to changing market conditions.

Finally, the inability to subject the same flight departure to different seat inventory control methods under exactly the same demand conditions meant that much of the analysis of the test results would be speculative. Even with a complete history of the fare class limits and the booking levels by day before departure, we can only speculate about what *might have* occurred in the booking process for the same flight departure had different fare class limits been applied. This is the major problem faced by airline managers hoping to measure the effects of seat inventory control practices—there is no way to determine exactly what revenues or loads have been realized in the absence of seat inventory control or under different methods.

In light of these interpretation constraints, the results of this test had to be scrutinized on a departure-by-departure basis, to assess how different fare class booking limits *likely* affected actual bookings in each fare class, as well as total flight revenues. The objective in interpreting the test results was, therefore, to identify BATCH and CONTROL flight departures that had similar booking patterns in the absence of different fare class limits. For the purpose of making direct comparisons between flights withas similar a demand pattern as possible, BATCH-CONTROL flight pairs were identified from the same week of the test period.

7.2 Revenue Impact Results and Assessment

The results of the revenue impact test were evaluated in terms of the differences in fare class mixes of passengers carried, load factors, and total flight revenues between the BATCH and CONTROL test groups. The assessment of revenue impacts was based on a comparison of *flight pairs* (i.e., BATCH versus CONTROL) that departed during the same week on the same flight leg. Post-departure results and booking histories were collected for a total of 210 flight pairs (21 flight legs, 10 test weeks). The major task in analyzing the results, as introduced above, involved identifying the flight pairs that demonstrated differences in loads and/or revenues due to the application of different fare class booking limits, with all else being equal.

This identification process involved making judgments as to the similarity of the booking patterns of the two flights in each flight pair, based on a number of criteria. Final reservations and boarding totals by fare class as well as complete booking histories for each flight pair were examined. The general rule used in selecting valid flight pairs was that any difference in load factors and fare class mix of passengers between the flights had to be explained by differences in the fare class limits applied during the booking process. The objective of flight pair selection was to identify BATCH and CONTROL flights operating on the same flight leg during the same week that exhibited reasonably similar demand processes and demonstrated the impacts, both positive and negative, of the Automated Booking Limit System.

The final set of valid flight pair comparisons proved to be relatively small. Out of the 210 possible flight pairs in the test, approximately two-thirds were eliminated because of unexpectedly low demand. Winter flight cancellations and a change in aircraft type on one of the flight legs also contributed to the number of flight pairs eliminated from the outset. Systematic day of the week demand differences or simply different booking build-up processes eliminated about one-half of the remaining flight pairs, leaving 36 valid flight pair comparisons (72 flight departures) for an assessment of ABLS impacts.

Examples of the flight pairs showing positive and negative revenue impacts of ABLS are provided in Tables I and II, respectively. For each flight pair listed, the fare class mix of passengers as well as the flight load factors and the percentage difference in flight revenues (BATCH over CONTROL) are shown. The four fare classes shown—Y, M, B, and Q—represent a descending hierarchy of fare values, although booking discrepancies caused the Y-class average revenues to be consistently lower than M- and B-class revenues. In general, however, the M/B/Q relationships held. The passenger mix is the number of passengers actually boarded by fare class, and the load factor is the percentage of seats on the aircraft filled for that specific flight departure. The specific flights and markets have been omitted for reasons of data confidentiality.

The asterisks in both tables indicate that the associated fare class limit was reached by bookings for that flight during its reservations process, and that the closing of the fare class had a significant impact on how the flight proceeded to book up relative to the other flight in the pair. Especially for flights that booked out very close to departure, a 100% load factor might be indicated without asterisks for one or more of the fare classes. The asterisks imply a *significant* constraint on the booking process for the relative fare classes.

Table I provides five examples of the 25 flight pairs in which ABLS was judged to have a positive impact on total flight revenues under similar demand conditions. The flight pairs in this table illustrate cases in which the revenue benefit of selling more low priced seats outweighed that of closing down the lowest fare classes in the hope that denied requests result in fare class upgrades. Note, however, that in the fourth pair listed, lower Q-class limits on the CONTROL flight did increase the number of B passengers carried, reflecting a propensity on the part of denied Q passengers to upgrade to a higher B fare.

The higher total revenues for the BATCH flights in Table I can be explained in the majority of flight pairs listed by higher load factors stemming from more liberal booking limits on lower fare classes. Specific comments for some of the examples shown are as follows.

- *Week* 2, *AAA/BBB*: Although the number of passengers carried in B and Q classes were similar, the number in M-class differed substantially due to a lower M-class limit for the CONTROL flight.
- *Week* 4, *CCC/DDD*: In the only flight pair for which the BATCH flight showed a lower load factor than the CONTROL flight, higher total revenues were realized on the former because of *lower* Q limits and substantial upgrades from Q to B.
- *Week* 7, *FFF/CCC*: In this short-haul, highly competitive market, very little upgrade activity was observed when Q-class was constrained, especially for relatively low load factor flights, like this one.

The average impact per flight pair for the 25 flight pairs in the positive impact group amounted to a 14.3% higher revenue for the BATCH over the CONTROL flight. When weighted by the revenue levels and loads associated with each flight pair, the aggregate positive impact on total revenues amounts to a 12% advantage for the BATCH flights. The average flight load factors for all flights in this positive impact group were also 12.3 percentage points higher for the BATCH than for the CONTROL flights, due to the application of ABLS-generated fare class booking limits.

Table I
ABLS Positive Revenue Impact

Test Week	Flight Leg	Test Group	*Y*	*M*	*B*	*Q*	L.F. (%)	Revenue Impact (%)
2	AAA/BBB	Batch	5	58	22	19	97	+17.1
		Control	9	41*	27	16	87	
4	CCC/DDD	Batch	17	3	56*	28*	97	+10.2
		Control	4	13	38*	51*	99	
5	EEE/AAA	Batch	9	7	9	97*	90	+15.1
		Control	11	14	11	61*	71	
7	BBB/AAA	Batch	3	26	15	69	83	+5.5
		Control	2	30	14	56*	75	
7	FFF/CCC	Batch	4	1	8	62	63	+22.6
		Control	5	4	9	41*	49	

To summarize the results of the positive impact flight pairs, the BATCH limits on lower fare classes were substantially higher than those for CONTROL flights, with a few exceptions. The CONTROL limits on the lowest fare classes were generally *too low* to maximize total flight revenues. On the other hand, there is some evidence that the BATCH limits on the lower fare classes might have been too high, giving passengers willing to purchase higher priced seats access to lower fare classes.

The relatively high booking limits on the lower fare classes set by ABLS for the BATCH flights can be attributed to at least two factors. First, the prorated leg revenue averages used as inputs by ABLS did not differ radically among the fare classes for many of the flight legs in the test. This was especially true for flight legs in highly competitive, short-haul markets, in which the price levels of the fare products sold in adjacent fare classes differed by as little as $10.

The second factor is related to the first, and involves the lack of any measure of upgrade potential in the EMSR formulation used in this test. Whereas the seat control analysts set fare class limits on CONTROL flights with some expectation of upgrade potential, the BATCH limits reflected an assumption of zero upgrade potential. As a result, there is evidence in many of the flight pairs falling into the positive impact group of upgrade behavior for the CONTROL flights with constrained lower fare class demands. For the flight legs operating in markets in which the price levels of fare products in adjacent fare classes were similar, or in which little effective competition was present, the actual upgrade potential appeared to be significant.

Examples of the flight pairs for which the inability of ABLS to take into account fare class upgrade potential led to a negative revenue impact are shown in Table II. All the flight pair examples show a BATCH flight load factor greater than or equal to that of the CONTROL flight, but a total flight revenue that was lower for the BATCH flight in each use. The revenue advantage for the CONTROL flights stemmed from lower booking limits on Q-class and substantial upgrade movements to higher revenue fare classes. The asterisks in Table II demonstrate this phenomenon

Table II
ABLS Negative Revenue Impact

Test Week	Flight Leg	Test Group	*Y*	*M*	*B*	*Q*	L.F. (%)	Revenue Impact (%)
1	DDD/GGG	Batch	5	3	12	85	98	−4.0
		Control	8	28	12	54*	95	
3	AAA/EEE	Batch	2	6	13	114	99	−4.0
		Control	4	25	51	43*	90	
5	DDD/HHH	Batch	4	5	9*	89*	100	−10.7
		Control	4	14	39*	43*	93	
7	JJJ/AAA	Batch	7	3	53*	73*	100	−2.8
		Control	13	23	38*	61*	99	
8	AAA/BBB	Batch	3	36	21*	46*	99	−15.3
		Control	22	52*	22*	10*	99	

clearly. Flight pairs that merit further comment include:

- *Week* 1, *DDD/GGG*: The upgrade movement here involved a vertical shift from Q to M-class, likely because the B-class fare product offered in the local market served by this flight leg was not as attractive an upgrade alternative as the M-class product (i.e., same price, more heavily restricted).
- *Week* 7, *JJJ/AAA*: The BATCH flight reached its limits in both B and Q classes, but these limits were higher than for the CONTROL flight in the same classes.
- *Week* 8, *AAA/BBB*: An example of upgrade movement from M to Y-class is provided by the CONTROL flight in this pair.

For the 11 flight pairs that showed negative revenue impact, the average revenue difference per flight pair amounted to a 6% shortfall for the BATCH relative to the CONTROL flight. The weighted difference in aggregate revenues for this group amounted to a 5.9% shortfall for the BATCH flights. The difference in average flight load factors was 5.7 percentage points in favor of the BATCH flights, many of which departed with no empty coach seats.

The overall conclusion that can be drawn from the negative impact flight pairs is that the BATCH flights had booking limits on the lower fare classes that were *too high* relative to the CONTROL flights. Upgrade activity contributed to higher total revenues for the CONTROL flights in this group, in spite of higher load factors for the BATCH flights. Most of the flight pairs in this group, however, belonged to one of three flight legs in the test, suggesting that these flight legs experience a significant upgrade potential that might not be realized on other flight legs in different markets.

The aggregate revenue and load factor results for the 36 flight pairs considered to reflect a valid comparison of ABLS versus the manual method of seat inventory control are summarized in Table III. An overall positive revenue impact of 6.2% was realized for the BATCH over the CONTROL flights in this comparison. The average flight load factor advantage for BATCH flights amounted to 10.3 percentage points. This discrepancy between the magnitude of the increase in load factor and total revenues is explained by the lower overall yield realized on the BATCH flights. The shortfall in overall yield was outweighed, however, by the increase in total revenues, providing evidence that yield maximization does not necessarily mean revenue maximization.

The potential positive impact on total flight revenues of an approach to seat inventory control more systematic than the application of booking limits based on analysts' judgments was clearly demonstrated by this initial test of a relatively basic version of the ABLS. The test helped to reinforce management's perception of the importance of focusing on total revenues rather than yields in seat inventory control.

Table III
Summary of ABLS Impact

ABLS Impact	Positive (25 flight pairs)	Negative (11 flight pairs)	Total (36 flight pairs)
Traffic			
Total passengers	+15.1%	+5.9%	+12.0%
Load Factors			
Average per flight	+12.3 pts.	+5.7 pts.	+10.3 pts.
Revenues			
Total flight revenues	+12.0%	−5.9%	+6.2%
Aggregate yield (cents)	−0.22	−1.18	−0.52

8. Lessons for Future System Development

The implementation experience at Western demonstrated that any airline seat inventory control system must be customized to fit the requirements, capabilities, and policies of the airline for which it is designed. The ABLS approach will provide a foundation for further development at Delta Air Lines, which purchased Western and inherited the lessons of the ABLS experience. Any further development will require major revisions and adjustments to ensure consistency with Delta's current practices and data availability.

The availability of valid and accurate data inputs to the EMSR model proved to be as important to ABLS implementation as the formulations of the model itself. Airline data is rarely available in the format or at the level of detail required for reliable estimates of future demand and revenues. Future development efforts, thus, will depend heavily on the airline's ability to collect and retrieve data from reservations histories and ticket revenue samples.

The objective of ABLS development at Western was to automate the seat inventory control process as much as possible, given a very small number of human analysts. The presence of many more experienced analysts at other airlines can reduce the emphasis on complete automation. The Western experience highlighted the fact that no seat inventory control system can be totally automated. Skilled analysts are required to examine input data and consider the recommended booking limits in light of exogenous variables and information that cannot be programmed into an automated system.

The shortcomings of the system tested and of the test itself will also guide future developmental efforts. The system tested included only the basic EMSR formulations for calculating initial fare class limits and revising them periodically in light of changes to the input data and actual bookings for a future flight. The EMSR model extensions to include upgrade probabilities and to take into account overbooking factors were not incorporated into ABLS. The revenue inputs reflected substantial data pollution in many cases, while the demand inputs were simple historical averages from recent operations of the same flight leg. No seasonal adjustment or growth trend forecasting was performed.

The relatively small proportion of flights for which an impact was observed demonstrated the importance of integrating the EMSR decision model with an efficient reservations monitoring system. The EMSR decision model is not applied to all flight legs and markets at all times. It is most effective in relatively high demand situations. High load factor flights generate a disproportionately high percentage of total airline revenues, and it is the high load factor flights that require the greatest amount of seat inventory control attention. A monitoring system that identifies flights for which high demand is expected would provide an important complement to the EMSR model.

The implementation and testing of ABLS at Western demonstrated the potential benefits of a systematic application of a quantitative decision model to seat inventory control. These benefits could be even greater with further improvements to the estimation methods and decision models employed. The EMSR decision framework provided the quantitative approach used by the automated system to derive fare class booking limits. The specific formulations used in ABLS, however, did not incorporate fare class upgrade probabilities. The test results provided numerous examples of the importance of passenger upgrade potential to revenue maximization. Furthermore, the demand and revenue inputs to the EMSR framework had shortcomings that could be overcome through the development of better estimation techniques.

The greatest challenge in the development and testing of the ABLS was to provide the managements of the airlines involved with proof of the revenue benefits of the system. This task was complicated by the inherent difficulties of measuring the impact of *any* seat inventory control policies, given the numerous variables that can contribute to the demand and revenue on a particular flight. Furthermore, the notion of maximizing *expected* revenues seemed difficult for management to comprehend, particularly when shown results for individual flights that clearly were not optimal. The probabilistic nature of demand and the mathematical formulations can be counterintuitive to the results oriented airline manager. Still, the promise of positive results like those obtained during the ABLS test at Western prompted Delta to continue the developmental effort.

References

BELOBABA, P. P. 1985. TWA Reservations Analysis: Demand Distribution Patterns. Flight Transportation Laboratory, MIT, Cambridge, Mass.

BELOBABA, P. P. 1987a. Airline Yield Management—An Overview of Seat Inventory Control. *Trans. Sci.* **21,** 63–73.

BELOBABA, P. P. 1987b. Air Travel Demand and Airline Seat Inventory Management. Ph.D. Dissertation, MIT, Cambridge, Mass.

BHATIA, A. V., AND S. C. PAREKH. 1973. Optimal Allocation of Seats by Fare. Presentation by Trans World Airlines to AGIFORS Reservations Study Group.

Boeing Commercial Airplane Company. 1982. Boeing Promotional Fare Management System, Analysis and Research Findings, pp. 7–16. Boeing, Seattle (July).

D'SYLVA, E. 1982. O and D Seat Assignment to Maximize Expected Revenue. Sales and Technology Department, Boeing Commercial Airplane Company, Seattle.

FUCHS, S. 1987. Managing the Seat Auction. *Airline Bus.* 40–44 (July).

GLOVER, F., R. GLOVER, J. LORENZO AND C. MCMILLAN. 1982. The Passenger Mix Problem in the Scheduled Airline. *Interfaces* **12** (No. 3), 73–79.

LITTLEWOOD, K. 1972. Forecasting and Control of Passenger Bookings. *AGIFORS Symp. Proc.*, 95–117.

MATHAISEL, D. F. X., AND H. DUCHESNE DE LAMOTTE. 1983. Optimization of Air Transport Yields and Capacity. Memorandum M83-1, Flight Transportation Laboratory, MIT, Cambridge, Mass.

RICHTER, H. 1982. The Differential Revenue Method to Determine Optimal Seat Allotments by Fare Type. *AGIFORS Symp. Proc.*, 339–362.

ROTHSTEIN, M. 1968. Stochastic Models for Airline Booking Policies. Ph.D. thesis, New York University, New York.

ROTHSTEIN, M. 1985. OR and the Airline Overbooking Problem. *Opns. Res.* **33,** 237–248.

WOLLMER, R. D. 1985. An Airline Reservation Model for Opening and Closing Fare Classes. Unpublished Internal Report, McDonnell-Douglas Corp., Long Beach, Calif.