Survey Paper

Airline Yield Management An Overview of Seat Inventory Control

PETER P. BELOBABA

Massachusetts Institute of Technology, Cambridge, Massachusetts

The seat inventory control component of airline yield management is examined, with an emphasis on the practical aspects of the problem. A survey of current airline practice indicates that seat inventory control is dependent on human judgment rather than systematic analysis. Past work on the development of mathematical methods in this area has focused on large-scale optimization models and simplified representations of the problem. There remains a need for practical solution approaches that incorporate quantitative decision tools.

It has been called the "adjustable-rate" air fare: "Tell us what you can afford and we'll send you a ticket."[1] Many airlines offer a wide range of fares for travel in a single city-pair market, including deeply discounted fares directed at the price-sensitive traveler, as well as higher-priced coach, business and first class fares. This pricing strategy has developed since deregulation, at least in part a response by major carriers to price competition from low-cost new entrant airlines. By offering a limited number of seats at these lowest fares, established airlines can at least appear to be competitive in price with the new entrants, and might even be able to fill otherwise empty seats with stimulated demand.

Seat inventory control is the practice of balancing the number of discount and full-fare reservations accepted for a flight so as to maximize total passenger revenues and/or load factors. Load factors can increase when more seats are made available at discounted fares. Selling too many seats at low fares, however, can cause per-passenger revenues (yields) to decrease, and diversion of high-fare passengers to more readily available low fares can lead to lower total revenues. Preventing such revenue dilution requires effective *yield management*, which includes both pricing and seat inventory control.

Although pricing is an important component of airline yield management, the fare levels to be offered on a flight are in most situations determined by pressures to match competitors' fares in the same city-pair market. Seat inventory control enables the airline to influence yields and total revenues on a flight-by-flight basis, within a given price structure. Controlling the mix of fares sold for a particular flight is thus viewed by some as by far the most important aspect of fare competition, more important than the actual prices charged.[2] Effective seat inventory management can permit pre-deregulation airlines to respond more rationally to price cutting by individual competitors and/or in specific markets. And, the payoff from effective seat inventory control can be substantial: Delta Airlines has estimated that selling just one seat per flight at a full fare rather than a discounted fare can add over $50 million to its annual revenues.[3]

The realization that effective yield management can increase revenues dramatically has prompted most airlines to consider improvements to virtually every aspect of the seat inventory control process. In what has become a race to find better ways to manage the sale of their seat inventories, airlines are expanding and reorganizing the departments responsible, upgrading reservations systems and developing sophisticated decision support systems.

Yield management in the airline industry is in a transitional phase, evolving from an *art* that relies almost exclusively on human expertise to a *science* that employs more systematic analysis and decision techniques. It is the purpose of this article to provide an overview of airline activities in seat inventory control—current and future. The scope of the problem

Transportation Science
Vol 21, No 2, May 1987

0041-1655/87/2102-0063 $01.25

is outlined first, followed by a description of current airline practices. The development of mathematical models to solve the seat inventory control problem is then reviewed in some detail. Finally, directions for further research and development of more sophisticated seat inventory control techniques are suggested.

THE SEAT INVENTORY CONTROL PROBLEM

THE SEAT inventory control problem for a particular flight is defined to a great extent by the capacity allocation and equipment utilization decisions made in the airline schedule planning process. From the outset, the decision to operate a particular type of aircraft on a flight routing has implications for seat inventory management, given that anticipated demand and aircraft size are unlikely to match exactly. A significant excess or shortage of seats relative to total demand for a particular flight departure will reduce or increase the need, respectively, for strict control of discount fare sales. Seat inventory control in essence provides an opportunity to adjust for imperfections in the airline's schedule design.

The simplest approach to controlling seat inventories is to deal with each flight leg independently, rather than trying to optimize seat allocations over the airline's entire network of flights. Even when a single flight leg is isolated for seat inventory control purposes, the problem still involves passengers with many different origin-destination (O-D) itineraries on the same aircraft, all of whom generate different amounts of revenue. The single-leg seat inventory problem is thus not simply one of allocating seats to, for example, four fare classes sharing a common cabin on the aircraft. It can also require decisions as to the desirability (in revenue terms) of selling a seat in a higher-yield fare class to a single-leg passenger as opposed to selling that seat in a lower-yield fare class, but at a higher total revenue, to a multi-leg or connecting passenger.

The complexity of the seat inventory control problem, even at the single-leg level of analysis, has increased tremendously with the development of hub-and-spoke route networks by most large airlines. A large air carrier can operate over 1000 flights per day, serve several thousand O-D markets and offer five fare classes in each market. American Airlines, for example, now serves over 2700 O-D markets, compared with about 800 before deregulation.[2] For any one flight departure to its Dallas-Fort Worth hub, passengers typically can be booked into one of at least five fare classes to one of more than 50 destinations. There can thus be over 250 possible fare class/destination combinations for each available seat on such a flight leg, each of which will have different levels of desirability to the airline, in terms of revenue, yield and aircraft loads. With reservations for future flights being accepted up to 11 months in advance, the size of the seat inventory control problem can become unmanageable.

Clearly, no airline is in a position to make separate seat inventory control decisions about each of the tens of thousands of price/product combinations it offers each day. Much of the battle in the development of an effective seat inventory management process involves balancing the aggregation of O-D markets and/or fare classes offered on a flight, necessary to keep the size of the problem manageable, against the disaggregation necessary to enable the airline to respond to competitors' actions with changes to specific fare class inventories in specific markets.

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 procedures. The aircraft assigned to a particular flight departure is a given and, in turn, the number of seats to be shared among the fare classes in the coach cabin can be regarded as a fixed quantity. Seat inventory decisions must also take into account the overbooking practices employed by the airline. Because the problem involves the management of available reservations "spaces" as opposed to physical seats on an aircraft, the interaction between the fare class mix of passengers booked and the number that will ultimately show up for a flight can have significant revenue implications. Determining the extent to which passengers booked in different fare classes exhibit different "no-show" rates is thus an important component of the seat inventory management process.

Managing the inventory of available reservations spaces on a future flight leg is therefore a process which occurs in the context of a predetermined departure time and aircraft type, and which is generally subordinate to capacity decisions involving the distribution of onboard space among physical compartments and the targeted limits for overbooking the flight. Furthermore, in most instances, the fare types (and in turn the fare classes) as well as their respective prices can be assumed to be given and constant throughout the period in which most of the bookings for a flight will be received (starting about 6 weeks prior to departure).

Seat inventory control for a future flight departure can be as simple as setting booking limits on discount fare classes once at the start of the reservations process for that flight, and taking no further action as reservations are accepted. A more sophisticated approach takes into account the information provided by actual reservations, through monitoring of booking trends and adjustment of discount fare class limits to maximize total flight revenues or flight loads. A comprehensive effort to manage seat inventories and increase revenues is a dynamic one in which traffic and reservations histories are used to set initial fare class booking limits, actual bookings in each fare class are monitored relative to these limits, and adjustments are made on the basis of an analysis of past data, current bookings, and forecasts of future bookings for the flight.

The seat inventory control system employed by a particular airline will be influenced by the characteristics of the airline's route network and by its fare structure. An airline offering a single fare level for all seats on flights serving point-to-point markets on a nonstop basis clearly need not be concerned about sophisticated seat inventory control techniques. At the other extreme, a carrier with multiple fare classes on flights into and out of large connecting hub complexes can benefit immensely from improved seat inventory control. Most established (i.e., prederegulation) airlines are closer to the latter extreme, and as such are very interested in all aspects of yield management. Among such airlines, there is a range of effort devoted to, and a range of sophistication achieved in, seat inventory control. As described in the following section, current practice in this area is evolving rapidly, yet the emphasis on human expertise in making seat inventory control decisions remains.

A SURVEY OF SEAT INVENTORY CONTROL PRACTICES

REPRESENTATIVES of eight large North American airlines were interviewed between August and December 1985 to determine the present and future status of seat inventory control at each carrier. The airline representatives were understandably reluctant to provide specifics as to the booking limits and other criteria used in managing seat inventories. It was possible, however, to develop an understanding of the process adopted by each carrier, including the organizational structures involved, reservations system and data retrieval capabilities, as well as the methods used to monitor bookings and control the sale of seats at discount fare levels.[4]

Seat inventory control and yield management are closely related to a range of other functions in the airline corporate structure, including pricing, marketing, sales, reservations, overbooking and payload control. It is difficult for an airline to combine all the functions critical to yield management into a single unit, although several of the carriers surveyed have moved in this direction. Pricing and overbooking control were the most frequently named functions to have been incorporated into the yield management unit, which in turn was most commonly found in the airline's marketing or market planning department. Coordination with the remaining related functions that, for various reasons, cannot be included in the same department, poses a problem for most of the carriers surveyed.

The personnel responsible for actually setting, monitoring and adjusting booking limits—inventory control agents—make up the largest group in most yield management units. There were substantial differences among the carriers surveyed in the use of inventory control agents, both with respect to the number of agents employed and the degree to which the agents are made responsible for specific markets.

For those carriers with relatively few agents working on seat inventory control, booking limit monitoring and adjustment is almost entirely an ad hoc process, perhaps targeted at selected markets and flight legs. The carriers with proportionately more agents generally take a more systematic approach in which teams of agents are responsible for groups of markets and/or flight legs. The carrier with the most agents relative to its daily departures has taken the notion of specialization to the extreme, making each seat inventory control agent responsible for all flight legs that traverse a particular route or set of routes. These agents are then held accountable for the traffic mixes and revenue levels achieved on their own routes.

Seat inventory control is highly dependent on the capabilities of the airline's reservations system. Yet, many of the airline reservations systems in place today are deficient in several areas of relevance to seat inventory management. The way in which fare classes that share a common seat inventory are structured in the reservations system differs among carriers, affecting the way in which booking limits on each class can be calculated and applied. Strictly independent ("stand-alone") fare class "buckets" represent the simplest reservations system structure. Most airlines have reservations systems in which discount fare classes are "nested" within the highest fare class and share the same seat inventory. These subclasses of the coach cabin inventory can themselves be independent of each other or nested serially in descending order of yield.

A major limitation of most airline reservations systems is the number of fare class buckets into which bookings can be logged. At least one bucket is required for each of the physical compartments on an aircraft. Any remaining buckets are used by most carriers as subunits of a shared coach compartment. Currently, most systems limit fare class bookings by flight leg and some have the capability of limiting sales to local passengers in favor of through and connecting passengers that generate more total revenue. The large number of fare product/O-D market combinations possible on a single flight leg makes it desirable for the airline to be able to take reservations in many fare classes and to limit sales in specific markets when necessary.

Airline attempts to make use of more fare class buckets to control their seat inventories by passenger O-D itinerary have been impeded by the existing standardized distribution system. The major reservations systems used by travel agents display seat availability by flight leg and accept bookings in a maximum of five buckets. Given the need to exchange availability and booking information with other reservations systems, airlines are constrained in improving their own systems by the need to maintain this standardization.

Several carriers are nonetheless upgrading their reservations systems to accept bookings in up to 40 reservations buckets. In one example of what will be the "new generation" of reservations systems, eight fare classes will have up to five subordinate buckets for controlling discount fare bookings in particular O-D markets. The expanded reservations systems will allow airlines, for example, to stop sales of extremely low-priced seats in selected markets on a connecting flight leg without closing down the entire fare class to additional bookings. Existing reservations systems do not permit the airline to distinguish between discount-fare passengers with very different itineraries and total fare levels, as these passengers are booked into the same fare class.

The development of such O-D based reservations systems is a top priority in the area of yield management for several of the airlines surveyed. American Airlines have acknowledged that they in fact have such a system already in place, and that the focus of their seat inventory control process has changed to one of "selling the system."[5] That is, it is their goal to manage fare class inventories with respect to the revenues generated by the passengers on local, through, and connecting itineraries, all on the same flight leg. The decision models required to achieve this goal are far more complex than those required for simple leg-based seat inventory control, and are still in the developmental stages.

An airline's reservations system also plays an important role in providing decision support for seat inventory control. Decisions must be made with respect to initial fare class booking limits, which will potentially require revision on the basis of actual reservations and forecast demand. Both the initial limits and demand forecasts must be derived at least in part from historical booking patterns and traffic data for the same or similar flights. The capability to retrieve and summarize relevant historical data is thus crucial to seat inventory management. Some of the carriers surveyed are developing their own decision support tools for use in conjunction with their reservations systems. An option being considered by the rest is the purchase of separate programs and even computer systems for decision support purposes.

Several software companies are marketing yield management "packages" designed to extract data from airline reservations systems and provide decision support for seat inventory control. For example, the Control Data Corporation's (CDC) "MARKSMAN" airline yield management system "boils down the vast mass of data to manageable proportions and presents it to the agent in a form which allows decisions to be made."[3] While it also offers an ability to monitor actual bookings relative to historical patterns, the CDC system is essentially a statistical data management software package developed for seat inventory control applications. This and other "packages" on the market are not designed to determine optimal booking limits, as they do not have the capability to forecast demand and revenue levels, nor do they make use of revenue or traffic optimization routines.

The carriers surveyed all recognize the importance of efficient and usable decision support systems for seat inventory management. Whether such a system is developed in-house or simply purchased as a "package," airline managements are eager to invest in this aspect of yield management. In fact, it is safe to say that the areas currently receiving the greatest amount of resources and effort from airlines interested in improving their seat inventory control process are reservations and decision support systems.

The organizational structure of an airline's yield management unit, together with the decision support tools available to it, provide a foundation for the tasks of setting, monitoring and adjusting fare class booking limits so as to maximize flight revenues. It is in this component of the yield management process that differences among airlines in terms of sophistication are most apparent.

At the simplest level, setting the booking limits for discount fare classes can be done on an "across-the-board" or default value basis. The use of default values for all flights operated in certain types of markets or with particular aircraft types requires little in the way of resources, but does not take into account important differences in passenger mixes and booking patterns between markets or even between flights in the same market. More important, the use of default booking limits set once at the start of the booking process ignores the information provided by actual bookings for the flight. All of the carriers surveyed have progressed beyond this basic level of seat inventory control, although several have done so only in the recent past.

The emphasis in current seat inventory control practice seems to be on *monitoring* actual bookings relative to fare class booking limits. The monitoring function can easily be automated through the airline's reservations system. Even the simplest of systems can be programmed to generate reports listing the future flights for which the number of accepted reservations approaches fare class booking limits. All of the carriers surveyed have reporting systems that can perform this function, although not all of them have a seat inventory control process in place that makes full use of the monitoring reports available.

The airline reservations systems surveyed all monitor actual bookings relative to preset and static fare class limits, with varying degrees of sophistication. An improvement to the monitoring process is provided by the CDC yield management package, among others. On the basis of historical booking trends for a flight or group of flights, the CDC package generates "booking threshold curves" which show the expected range for cumulative bookings at any point before departure. The flights for which actual bookings stray outside the range of these curves are then flagged by the system and listed in periodic reports.

The least advanced aspect of yield management and seat inventory control at all of the airlines surveyed is that of booking limit adjustment to maximize flight revenues. This task is the most important component of seat inventory control, yet it remains dependent on human judgment rather than systematic analysis. When an airline's reservations monitoring system flags a flight for which actual bookings approach any one of the limits or the threshold set for that flight, a decision must be made either to increase the availability of seats in the relevant fare class or to allow the system to close it down to additional reservations. This decision is currently being made by individuals or groups of individuals on the basis of experience and judgment at every airline surveyed, although the ongoing development of decision support tools is designed to reduce the amount of guesswork involved.

At least two carriers are hoping to improve this adjustment process by developing algorithms for finding the fare class booking limits that will maximize expected flight revenues. Such algorithms will not eliminate the need for human judgment in seat inventory control entirely, since any optimal solution would be probabilistic in nature and would be derived from forecasts based on historical data. There will always be variables that cannot be accounted for in such algorithms, including rapid changes in the competitive environment of airline markets and the occurrence of unexpected events that affect flight bookings. The objective in developing optimization models for seat inventory control is to allow yield management agents to focus their efforts on these variables by making routine tasks more systematic.

MATHEMATICAL MODELS FOR SEAT INVENTORY CONTROL

THE LOW level of sophistication in the tasks of setting and adjusting fare class booking limits as practiced by airlines is in large part attributable to a lack of practical models for making optimal decisions. The theoretical and empirical research necessary to develop optimization models for determining the number of seats to allocate to each fare class simply did not keep pace with the rapidly changing competitive conditions in the airline industry since deregulation. In fact, the need for such models was not even acknowledged by most airlines until recently.

This section provides an overview of the mathematical concepts, models and solution methods relevant to airline seat inventory control. The discussion centers on past work, which started in the early 1970s with the introduction of advance purchase excursion (APEX) fares in international markets. The use of expected marginal revenues by fare class in finding optimal seat allocations is discussed first. Solution algorithms that make use of the expected marginal revenue approach to optimize booking limits in reservations systems with independent buckets are then reviewed.

The airline seat inventory management problem has both probabilistic and dynamic elements to it. The problem is probabilistic because 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 problem is dynamic because the total number of reservations requests accepted for a flight will change from day to day, potentially affecting estimates of requests still to come and, in turn, the optimal allocation of remaining seats among fare classes.

The notion of probabilistic demand is central to the airline seat inventory control problem, as the expected number of requests for each fare class must be estimated from historical distributions of demand. Also important is the capacity constraint on the total number of seats available on a flight leg. The number of seats allocated to a particular fare class might not always exceed the number of requests for that fare class, resulting in rejected demand, or "spill."

Figure 1a illustrates these concepts graphically, for a simplified two-class seat inventory control example. We define $p_i(r_i)$ to be the probability density function

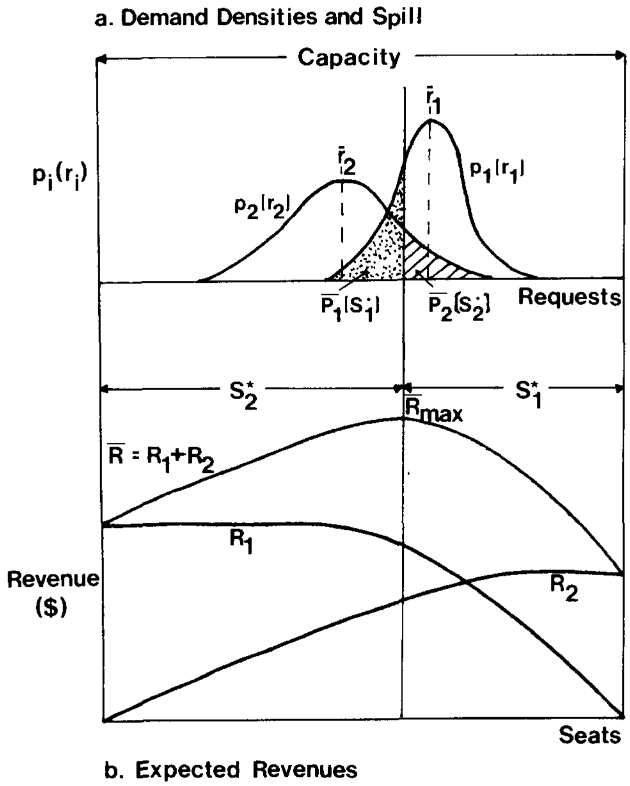

Fig. 1. Optimal seat allotment.

for the total number of requests for reservations, r_i, received by the airline for seats in class i by the close of the booking process for a scheduled flight leg departure. The probability of receiving $S_1{}^*$ or more requests for fare class 1 is therefore $\bar{P}_1(S_1{}^*)$, as indicated by the shaded area under $p_1(r_1)$.

The implications of these probabilistic concepts for mathematical modeling of the problem become apparent in the simplest of revenue maximization models. We can postulate a static model of the seat inventory control problem which can be used to determine the optimal allocation of seats between two independent (i.e., non-nested) fare classes, subject to the total capacity constraint. The total capacity of the cabin to be shared among i fare classes is C, such that:

$$C = \textstyle\sum_i S_i. \qquad (1)$$

Let f_i be the average fare or relative revenue received by the airline when a reservation request for fare class i is accepted, and $\bar{b}_i(S_i)$ the expected number of bookings in class i, given a seat allocation of S_i. We want to find the values of S_i that will maximize total expected revenue, $\bar{R}$, for a flight:

$$\begin{aligned} \bar{R}_i(S_i) &= f_i \cdot \bar{b}_i(S_i), \quad \text{for all} \quad i \\ \bar{R} &= \textstyle\sum_i \bar{R}_i \end{aligned} \qquad (2)$$

subject to the capacity constraint.

In the two-class example, we have fare classes 1 and 2 with relative fares f_1 and f_2. The seats allocated to each fare class (S_1 and S_2) must share a total capacity of C seats. To find the value of $S_1 = C - S_2$ that will maximize total expected revenues, $\bar{R}$, for the flight, we differentiate $\bar{R}$ with respect to S_1 and set to zero:

$$\begin{aligned} \bar{R} &= \bar{R}_1(S_1) + \bar{R}_2(C - S_1) \\ &= f_1 \cdot \bar{b}_1(S_1) + f_2 \cdot \bar{b}_2(C - S_1) \\ \partial\bar{R}/\partial S_1 &= 0. \end{aligned} \qquad (3)$$

That is, seats are allocated between the fare classes such that the marginal expected total revenue with respect to additional seats in each class is equal to zero. At optimality, total expected flight revenues cannot be increased by taking a seat from class 1 and allocating it to class 2 instead. The expected marginal seat revenue for each class, $\text{EMSR}_i = \partial\bar{R}_i/\partial S_i$, will also be equal across all relevant fare classes, but will not necessarily equal zero, due to the imposed capacity constraint.

The expected marginal seat revenue of the Sth seat in fare class i is simply the average fare level in class i, f_i, multiplied by the probability of selling S_i or more seats, $\bar{P}_i(S_i)$. The optimality conditions for the above example can therefore be expressed as follows:

$$\begin{aligned} f_1 \cdot \bar{P}_1(S_1{}^*) &= f_2 \cdot \bar{P}_2(S_2{}^*) \\ \frac{\bar{P}_1(S_1{}^*)}{\bar{P}_2(S_2{}^*)} &= \frac{f_2}{f_1}. \end{aligned} \qquad (4)$$

The optimal values of S_1 and S_2 will depend on the parameters of the probability densities of expected demand for each fare class, the relative fares or revenue levels, and the total capacity available. The relationships between $\bar{R}$, S_i and C are illustrated for the simple two-class model in Figure 1b. $\bar{R} = \bar{R}_1 + \bar{R}_2$ is maximized when $S_1{}^*$ out of a total of C seats are allocated to the higher-priced fare class 1.

The values of $S_i{}^*$ derived from this model represent the optimal allotments of the available seats to independent buckets based on expected demand levels for each fare class. The demand for each fare class is assumed to be independent of that for the other class, and the optimal seat allocation is made only once, at the beginning of the booking period for a flight. In reality, demand for different types of fares might not be independent, as high demand for one fare class could be associated with high demand for another. More important, this static seat inventory management model does not account for the dynamic nature of the reservations process in which actual bookings accepted for a flight might provide valuable additional information about the ultimate number of requests that can be expected.

The approach of equating marginal revenues in each of two fare classes to find the revenue-maximizing seat allotments for a flight leg was applied to a dynamic reservations context by LITTLEWOOD[6] in 1972. He suggested that total flight revenues would 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 of selling that same seat at a higher fare. That is, low-fare passengers paying f_2 should be accepted as long as:

$$f_2 \geq \bar{P}_1(S_1) \cdot f_1 \quad (5)$$

where

$\bar{P}_1(S_1)$ = probability of selling all remaining seats to high-fare passengers

f_1 = higher fare level

The smallest value of S_1 that satisfies the above condition is the revenue-maximizing booking limit on class 1.

Variations on Littlewood's simple model were proposed by Trans World Airlines analysts in 1973[7] and by RICHTER[8] of Lufthansa in 1982. In each case, however, the extensions proved to be essentially equivalent to the original model. Richter's model accounted for losses in total expected revenue when low-fare passengers ultimately deny space to higher-fare passengers. The differential revenue from allocating an additional seat to low-fare passengers was defined as the difference between the additional revenue realized from the low-fare sale and the revenue lost from prospective high-fare passengers, as follows:

$$\text{DR} = f_2 \cdot \bar{P}_2(S_2) - f_1 \cdot \bar{P}_1(C - S_2 + 1). \quad (6)$$

The expected marginal seat revenues for the two fare classes are equal when DR is set to zero.

The original formulation (6) included probability densities of requests for both low-fare and high-fare demand, which proved to be unnecessary in cases where either dynamic booking limit revision or nested fare classes ensure that high-fare requests will not be denied as long as seats are available. Richter demonstrated that the optimal limit on low-fare seats is *not* a function of low-fare demand in such cases, although the distribution of low-fare requests will influence the expected total revenue for the flight. The optimality condition in (6) was reduced to:

$$f_2/f_1 = \bar{P}_1(C - S_2 + 1), \quad (7)$$

which is in fact equivalent to Littlewood's "simple model" in (5).

Work has also been done to test the validity of the assumptions made in deriving this simple model for allocating seats between two fare classes. The most important of these assumptions include:

1. Low-fare passengers always book first;
2. There are no cancellations of bookings;
3. A rejected request is revenue lost by the airline.

A sensitivity analysis of the simple model under these assumptions was performed by MAYER[9] of El Al in 1976. This analysis showed that the greater the difference between f_1 and f_2, the more sensitive the total expected flight revenue will be to a non-optimal allocation of seats. The decrease in expected revenue will be smaller when too many seats are allocated to low fare passengers than when too few seats are offered.

The question of how the assumption that low-fare passengers book first affects the optimal seat allotments and expected revenue levels was also addressed by Mayer, as well as by TITZER and GRIESSHABER[10] in 1983. Mayer relaxed the "early-bird" assumption of the simple model, assuming instead that low-fare passengers book first in each of many *periods* before departure. Titzer and Griesshaber went a step further, comparing the "early-bird" assumption with a simulated parallel process of booking in which the rate of low-fare requests decreases and the rate of high-fare requests increases as departure day approaches. Both analyses showed that the booking behavior assumption should not have a significant impact on the optimal seat allotments as determined by the model, as long as demand for each fare class is assumed to be independent.

The expected marginal revenue approach was applied by BUHR[11] of Lufthansa in 1982 to the problem of allocating seats on a two-leg flight (A to B to C), although only one fare class was considered. He defined the expected "residual" revenue from allocating an additional seat to a passenger flying from A to C as:

$$E_{AC}(S_{AC}) = \bar{P}_{AC}(S_{AC}) \cdot f_{AC} \quad (8)$$

where:

f_{AC} = average fare for A to C passengers

$\bar{P}_{AC}(S_{AC})$ = probability of selling S_{AC} or more seats to passengers from A to C

With demand for each O-D market assumed to be independent, the problem involves allocating seats on each leg to either a through or local passenger. The booking limits derived thus involve independent "buckets." Buhr postulated that total flight revenues would be maximized when:

$$E_{AC}(S_{AC}) = E_{AB}(S_{AB}) + E_{BC}(S_{BC}) \quad (9)$$

subject to the capacity constraint. An iterative solution method was used to find the optimal values of S_{AC} and $S_{AB} = S_{BC}$.

With respect to the allotment of seats among fare classes, Buhr suggested a two-step approach in which the optimal booking limits are determined for each flight leg first. An allocation of seats among fare classes would then be performed, given the flight leg limits. Buhr acknowledged that varying the allotment of low-fare seats for different passenger itineraries could change the expected revenue levels on each leg of the flight, but did not address this problem.

Buhr's formulation was extended to more complicated situations by WANG[12] of Cathay Pacific in 1983. He set out to develop a model for determining the optimal seat allotments on multiple leg flights with multiple fare classes. He suggested that optimal load targets for each fare class in each O-D city-pair served by the flight could be found by allocating each marginal seat to the O-D and fare class combination with the highest total expected revenue. He thus assumed independence of both O-D markets and fare class buckets. Given a flight that serves j O-D pairs and offers k fare classes in each O-D market, each O-D and fare class combination i has an expected marginal revenue, EMR_i, of:

$$\mathrm{EMR}_i = \sum f_{jk} \cdot P[r_{jk} > b_{jk}], \quad \text{for feasible } jk \text{ pairs} \tag{10}$$

where:

f_{jk} = fare yield for fare class k in O-D market j

$P[r_{jk} > b_{jk}]$ = probability that another request for (j, k) will be received given b_{jk} bookings accepted

The terms of the right hand side are summed over all feasible combinations of sequential (j, k) pairs for the flight. For example, on a flight operating A to B to C, a feasible sequence would be to allocate the marginal seat to a low-fare class from A to B and to save that same seat for a high-fare passenger from B to C. In total, there would be 12 feasible combinations for a one-stop flight with three fare classes. Wang's approach requires the feasible combinations to be ranked in terms of expected marginal revenue as each seat is allocated incrementally.

Equating the expected marginal revenues of incremental seats allocated to all fare class/itinerary combinations will maximize expected revenues in cases where the inventory buckets are completely independent. In the simple two-class revenue maximization model presented at the start of this section, it was possible to find the optimal seat allotments through a straightforward application of differential calculus. When the problem is expanded to multiple fare classes and passenger itineraries, it becomes more difficult to find the optimal points analytically. The tools of mathematical programming and network flow analysis have been applied to such problems to find the optimal seat allotments through more efficient solution methods.

Glover et al[13] postulated that there is "some number of passengers at each fare class on each flight segment that will optimize revenue" for the airline, and developed a network-based seat allocation model for Frontier Airlines. The model was designed to find the mix of passenger itineraries flowing over the airline's network in independent fare classes that would maximize total daily revenues. One set of arcs in the network formulation represented flight leg loads (with capacity constraints), another set of arcs represented the various passenger itineraries (PIs) by fare class. The model accommodated up to 600 daily flights, 30,000 PIs and five fare classes. Its one major shortcoming was that the demand estimates used as input for each PI/fare class combination were entirely deterministic.

Analysts at both Boeing Aircraft[14] and McDonnell-Douglas[15] have developed mathematical programming techniques that take into account probabilistic demand for each fare class. Because the expected revenue objective function for each class is nonlinear given a constant fare and stochastic demand, a simple linear programming approach is inadequate. McDonnell-Douglas analysts proposed a formulation of the single-leg seat allocation problem that makes use of binary decision variables in an integer programming framework. Each variable, X_{ik}, represents the combination of fare class i and seat k on a flight leg. A 150-seat aircraft, 4 fare class formulation would thus require 600 such decision variables. Associated with each X_{ik} is the marginal expected revenue of selling the kth seat in class i, denoted $m_i(k)$, derived by multiplying the average fare level in class i by the probability of selling k or more seats in that class.

With an available total capacity of n seats, the formulation of this integer program is as follows:

$$\text{MAX } \bar{R}(n) = \textstyle\sum_i \sum_k X_{ik} m_i(k) \tag{11}$$

subject to:

$$\textstyle\sum_i \sum_k X_{ik} \le n$$

$$0 \le X_{ik} \le 1$$

The solution to this integer program will have the X_{ik} values corresponding to the n largest values of $m_i(k)$ equal to 1, with all other $X_{ik} = 0$.

The problem of the nonlinear objective function was also addressed at Boeing Aircraft by D'SYLVA,[16] who used a piece-wise linear approximation of the expected revenue curve in an integer programming formulation. He found that 5 to 10 binary decision variables could replace the 200 used previously for a

200-seat aircraft in approximating the revenue function of each fare class. D'Sylva used this approach to extend Glover's algorithm to include stochastic demand. An arc was added to the network representation for each of the straight-line approximations used in the total expected revenue objective function of the integer programming formulation. A comparison of solutions to the probabilistic and deterministic formulations showed that the latter overestimated expected revenues by about 12%. Furthermore, the best variable demand solution produced a five percent higher expected revenue than the deterministic solution.

At McDonnell-Douglas, WOLLMER[17] also pursued the network approach to the multiple fare class/multiple flight leg problem, based on the integer programming formulation in equation (11). In his network formulation, each X_{ik} represents a passenger itinerary/fare class combination on a particular flight leg, and is associated with an arc of length $m_i(k)$. Total network revenues are maximized by finding the n longest arc combinations, or paths, from the network origin node to the final destination node.

As this network formulation was expanded to include multiple flight legs and passenger itineraries involving connecting flights, the number of binary decision variables required increased rapidly. Wollmer suggested that while the complete network formulation will contain a large number of arcs, only a few of the arcs need to be included at any one time for consideration under the longest-path criterion. Specifically, at most two arcs for each fare class/passenger itinerary (PI) combination need to be considered at any one time—the lowest revenue (shortest) arc with an existing flow of one, and the highest revenue (longest) arc with flow of zero.

A solution algorithm for this network problem therefore need only solve a series of longest path problems for a relatively small network. At each iteration, the expected marginal revenue of each fare class/PI combination in the current reduced network must be calculated, the largest value identified, and a seat allocated to that path. The marginal expected revenues (lengths) associated with the arcs on the longest paths would then be revised, and the procedure repeated.

Wollmer and others at McDonnell-Douglas have continued work on the seat allocation problem, extending the mathematical programming formulations to an airline connecting hub operation. At the same time, the network characteristics of these formulations and the expected marginal revenue approach to seat allocation are being used to develop solution algorithms more efficient than the integer programming methods described above. Furthermore, work is under way on dynamic applications of these optimization models to the reservations process, to revise booking limits as reservations are accepted and flight departure day approaches.

The major shortcomings of all of the mathematical programming and network formulations outlined above involve the assumed independence of fare class inventories and the size of the formulations required to model stochastic demand accurately. First, the optimal seat allotments as derived by these solution methods will not necessarily maximize revenues in a nested fare class reservations structure. In a nested system, it might well be impossible to find optimal booking limits analytically. Second, the size of the seat inventory control problem and the volatility of the airline competitive environment dictate that any optimization model be both efficient and adaptable to changing conditions. Decision rules, like Littlewood's "simple" model, that can be used dynamically to limit bookings on flight legs or in specific markets might be a more practical approach to improving seat inventory control than large optimization models.

None of the mathematical models reviewed here address some practical considerations that can play an important role in determining the optimal booking levels for different fare classes. The relationship between overbooking and seat inventory control has been overlooked. Furthermore, refused requests are not necessarily lost to the refusing airline. The refused passenger may be accommodated in the requested fare class on another flight of the same airline, or might in fact agree to accept a reservation in a higher fare class on the originally requested flight. There is also the possibility of significant correlation of demand levels among fare classes. All of these considerations can have a considerable impact on the optimal fare class booking limits, and may complicate the derivation of these limits. Nonetheless, a realistic solution approach must take these factors into account.

PROSPECTS FOR IMPROVING SEAT INVENTORY CONTROL

SEAT INVENTORY control in the airline industry is at an intermediate level of sophistication, although there exist substantial differences between individual carriers. The increased importance to airline profitability of effective yield management has prompted many carriers to invest in improvements to the decision support tools required, and some are exploring the possibilities of making the process more systematic with the help of mathematical optimization and forecasting techniques.

The airlines that have undertaken to improve their yield management methods have found that seat inventory control not only involves several interrelated components, it is also closely related to other functions within the airline. Within the airline's yield management unit, improvements to decision support tools are of little value without corresponding improvements to the data inputs used by those tools, as well as the ways in which the outputs are used to make seat inventory control decisions. Sophisticated tools and mathematical models are of little value without adequate numbers of trained analysts to use them. Above all, any effort to improve such tools is wasted if a lack of coordination with related functions serves to undermine the revenue-maximizing objectives of seat inventory control. A coordinated systems approach with upper management support is therefore required to make the most of any investment in yield management by an airline.

Improvements to decision support systems can be realized in a relatively short time. An airline can purchase or develop in-house a decision support system to make use of the enormous amount of historical reservations and traffic data stored by most reservations systems. In addition to storing historical data, reservations systems are the operational center of the seat inventory control process. In contrast to the development of decision support tools, however, changing reservations system capabilities and relating them more to the needs of yield management is a medium-term objective for most airlines.

The most complex and longest-term objective for improving seat inventory control methods is the development of mathematical decision models to forecast future bookings and to optimize planeload revenues. The tools of decision analysis and mathematical programming are suited well to the revenue maximization problem. The challenge is to adapt these tools to the seat inventory control context and to develop practical versions which can be calibrated with reasonable validity on the basis of historical data.

Optimization models for seat inventory control will inevitably require input data in the form of forecast demand levels for future flights and estimates of expected revenues. The emphasis in past research has been on the development of optimization routines, while the problem of demand forecasting has received little attention. Given the inherent variability in air travel demand and the volatility of airline markets, the development of generalizable yet accurate forecasting models will be difficult. The challenge in this case is to apply statistical and econometric methods to reduce the uncertainty associated with future demand levels.

The optimization models summarized in this article assume that both the demand levels for each fare class in the same market and the fare class reservations "buckets" in the airline reservations system are independent. The former assumption is open to question and empirical analysis, while the latter is clearly not valid given that most airline reservations systems have nested or interdependent fare class structures.

Nevertheless, the work described above effectively incorporates probabilistic demand levels into the seat inventory control problem, and the extensions to more complicated situations will provide some direction for further research. The development of simplified, more efficient solution algorithms and their application to the dynamic seat inventory control problem would be of greatest practical benefit to the airline industry. The emphasis in any further model development, however, must be on matching the solution to the practical constraints of the problem, including reservations system capabilities, data availability, and the nature of airline competitive practices. Above all, efforts must be made to incorporate more realistic demand assumptions, to improve the accuracy of the input data required by such models, and to recognize the modeling problems posed by reservations systems with interdependent or nested fare classes.

The quest for better yield management methods by airlines can involve substantial investments in personnel and computer support. Potentially the largest gains from more sophisticated seat inventory control techniques might not be realized for several years, as better optimization models are developed and implemented. The returns, however, can be significant. Each increase in yield (cents/RPM) of half a cent can mean $40 million in additional revenue annually to airlines the size of Republic or Western, and $140 million to airlines like Eastern, Delta and TWA. With restrictions on discount fares, more restrained matching of competitors' low fares and limited sales of seats at low fares, pre-deregulation airlines have begun to adapt successfully to the increased competitiveness of a deregulated industry. Effective seat inventory control is critical for continued or even increased profitability for these carriers.

REFERENCES

1. K. LABICH, "Fare Wars: Have the Big Airlines Learned to Win?" *Fortune* 24–28 (October 29, 1984).
2. R. CRANDALL, Chairman—American Airlines, quoted in "Life Under Deregulation The American Way," *Travel Weekly*, May 31, 1985.
3. "Yield Managers Now Control Tactical Marketing," *Lloyd's Aviation Economist* 12–13 (May 1985).
4. P. BELOBABA, Airline Yield Management: An Overview of the Problem and Current Practice, FTL Memorandum M86-1, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, Mass., 1986.
5. T. G. PLASKETT, Senior Vice-President Marketing—

American Airlines, Presentation to Shearson, Lehman Brothers Airline Industry Seminar, Key Largo, Fla., February 14, 1986.

6. K. LITTLEWOOD, "Forecasting and Control of Passenger Bookings," *AGIFORS Symp. Proc.* **12,** 95–117 (1972).
7. A. V. BHATIA AND S. C. PAREKH, Optimal Allocation of Seats by Fare, Presentation by Trans World Airlines to AGIFORS Reservations Study Group, 1973.
8. H. RICHTER, "The Differential Revenue Method to Determine Optimal Seat Allotments By Fare Type," *AGIFORS Symp. Proc.* **22,** 339–362 (1982).
9. M. MAYER, "Seat Allocation, or A Simple Model of Seat Allocation via Sophisticated Ones," *AGIFORS Symp. Proc.* **16,** 103–125 (1976).
10. B. TITZE AND R. GRIESSHABER, "Realistic Passenger Booking Behavior and the Simple Low-Fare/High-Fare Seat Allotment Model," *AGIFORS Symp. Proc.* **23,** 197–223 (1983).
11. J. BUHR, "Optimal Sales Limits for 2-Sector Flights," *AGIFORS Symp. Proc.* **22,** 291–303 (1982).
12. K. WANG, "Optimum Seat Allocation for Multi-Leg Flights with Multiple Fare Types," *AGIFORS Symp. Proc.* **23,** 225–237 (1983).
13. F. GLOVER, R. GLOVER, J. LORENZO AND C. MCMILLAN, "The Passenger-Mix Problem in the Scheduled Airlines," *Interfaces* **12,** 73–79 (1982).
14. The Optimal Partitioning of An Airplane's Seating Capacity, Unpublished Internal Report, Boeing Commercial Aircraft Company, Seattle, Wash., 1982.
15. R. D. WOLLMER, An Airline Reservation Model for Opening and Closing Fare Classes, Unpublished Internal Report, McDonnell-Douglas Corporation, Long Beach, Calif., 1985.
16. E. D'SYLVA, O and D Seat Assignment to Maximize Expected Revenue, Unpublished Internal Report, Boeing Commercial Aircraft Company, Seattle, Wash., 1982.
17. R. D. WOLLMER, A Hub-Spoke Seat Management Model, Unpublished Internal Report, McDonnell-Douglas Corporation, Long Beach, Calif., 1986.