A Discrete Choice Model of Yield Management

Kalyan Talluri *   Garrett van Ryzin †

September 15, 2000

Abstract

Customer choice behavior, such as "buy-up" and "buy-down", is an important phenomenon in a wide range of industries. Yet there are few models or methodologies available to exploit this phenomenon within yield management systems. We make some progress on filling this void. Specifically, we develop a model of yield management in which the buyers' behavior is modeled explicitly using a multi-nomial logit model of demand. The control problem is to decide which subset of fare classes to offer at each point in time. The set of open fare classes then affects the purchase probabilities for each class. We formulate a dynamic program to determine the optimal control policy and show that it reduces to a dynamic nested allocation policy. Thus, the optimal choice-based policy can easily be implemented in reservation systems that use nested allocation controls. We also develop an estimation procedure for our model based on the expectation-maximization (EM) method that jointly estimates arrival rates and choice model parameters when no-purchase outcomes are unobservable. Numerical results show that this combined optimization-estimation approach may significantly improve revenue performance relative to traditional leg-based models that do not account for choice behavior.

Key words. yield management, revenue management, discrete choice theory, airlines, dynamic programming

*Kalyan Talluri, Department of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain, email: talluri@upf.es

†Garrett van Ryzin, Graduate School of Business, Columbia University, New York, email:gjv1@columbia.edu

Introduction and Overview

Yield (or revenue) management is a practice that dates back to the deregulation of the U.S. airline industry in the late 1970's. It was developed as an outgrowth of the need to manage capacity sold at discounted fares, which were targeted to leisure travelers, while simultaneously minimizing the dilution of revenue from business travelers that are willing and able to pay full fares. Using statistical forecasting techniques and mathematical optimization methods, airlines developed automated systems to dynamically control the availability of the plethora of discounted fares that emerged in the post-deregulation era. The practice has since spread beyond airlines to the hospitality, rental-car, cruise-lines, railways, energy and broadcasting industries. Significant revenue benefits have been documented from such techniques - often an improvement of 2-8% in revenue over no revenue management or ad-hoc, manual controls [34].

Concurrent with the evolution of industry practice, a considerable amount of management-science literature on yield management has been published over the last twenty years. The earliest work on capacity control was Littlewood's [26] analysis of a simple, two-fare-class model of capacity allocations on a single flight leg. The problem with more than two fare classes (we define a fare class as a fare (rate, price) along with its associated set of restrictions to qualify for this fare) is considerably more complex, but Belobaba [4], [5], [6] developed two simple and effective heuristics for the single-leg problem based on the concept of expected marginal seat revenue (EMSR-a and EMSR-b) that are still in wide-spread use today.

On a theoretical level, single-leg models in which demand for each fare class is assumed to occur in non-overlapping periods have been developed and analyzed by Brumelle and McGill [11], Curry [14], Robinson [32] and Wollmer [42]. A key result of this work is that the optimal policy can be implemented using a set of so-called *nested allocations*. (See Brumelle and McGill [11] for a precise definition of nested allocations.) Lee and Hersh [24] introduced and analyzed a discrete-time model, Markov model that allows for an arbitrary order of arrivals. For further work on single-leg allocation problems, see Brumelle et al. [12], Kleywegt and Papastavrou [22], Lautenbacher and Stidham [23], Liang [25], Stone and Diamond [35], Subramanian et al. [36] and Zhao [45]. For analysis of multiple-leg (network) allocation problems, see Cooper [13], Curry [14], Dror et al. [17], Glover et al. [21], Simpson [33], Talluri [37], Talluri and van Ryzin [38], [39] and Williamson [40], [41]. A recent survey of yield management research is provided by McGill and van Ryzin [29]; Talluri and Barnhart [2] provide an overview of yield management and other airline operation research areas.

Despite the success of this body of work, most of the above-mentioned models make a common, simplifying - and potentially problematic - assumption; namely, that consumer demand for each of the fare classes is completely independent of the controls being applied by the seller. That is, the problem is modeled as one of determining which exogenously arriving requests to accept or reject, and it is assumed that the likelihood of receiving a request for any given fare class does not depend on which other fares are available at the time of the request. However, casual observation - and a brief reflection on one's own buying behavior as a consumer - suggests that this is not the case in reality. The likelihood of selling a full fare ticket may very well depend on whether a discount fare is available at that time; the likelihood that a customer buys at all may depend on the lowest available fare, etc. Clearly, such behavior could have important revenue management consequences and should be considered when making control decisions.

We lay no claim to uncovering this deficiency. Indeed, many researches have tried to address "buy-up" (buying a higher fare when lower fares are closed) and "buy-down" (substituting a lower fare for a high fare when discounts are open) effects in the context of traditional models. Phillips [31] proposed a "state-contingent" approach to yield management that adjusts controls based on forecasts that depend on the controls in effect (the system "state") at any point in time, though no recommendations on how to obtain such forecasts are given. Belobaba [4] proposed a correction to the EMSR heuristics that introduces a probability of buying a higher fare when a low fare is closed. While conceptually appealing for a two-fare-class model, such pair-wise "buy-up" probabilities are problematic in a multiple-fare-class setting. The probability of buying a given high fare should depend on which other high fares are also available. Also, one cannot directly observe "buy-up", so how does one separate "original" sales from "buy-up" sales? How are the probabilities (forecasts) adjusted when there are price changes? etc. Despite these difficulties, we are aware of several airlines that have experimented with or regularly use this approach. Indeed, Andersson [1] describes an significant research and development effort by SAS to apply logit choice models to estimate buy-up and dilution factors, which were then incorporated in various buy-up/dilution heuristics at one of SAS's hubs.

Another stream of work on understanding choice behavior is the passenger origin and destination simulator (PODS) studies of Belobaba and Hopperstad. (See [7].) PODS is a detailed simulation model of passenger purchase behavior developed by Hopperstad at Boeing. It includes factors for airline preference, time preference, path preference and price sensitivity. While it is a very detailed simulation model, the focus of the PODS studies is to test the performance of traditional forecasting and optimization methods under conditions of complex passenger choice behavior rather than to develop new estimation and optimization methods. Nevertheless, the PODS studies have provided many useful insights and clearly demonstrate the significant impact that choice behavior has on the performance of yield management systems.

The only theoretical models and methods that partially address choice behavior issues are dynamic pricing models, such as those studied by Bitran et al. [10], Feng and Gallego [18] and Gallego and van Ryzin [19], [20]. While these models allow demand to depend on the current price (the control in this case), they assume only one product is sold at one price at any point in time. Thus, customers face a binary choice; to buy or not to buy. In reality, airlines offer many fares simultaneously and customers choose among them based on price together with their preferences for non-price factors, such as refundability and whether or not they can meet various restrictions (e.g. Saturday night stay, minimum-stay and maximum-stay). The above dynamic pricing models do not capture this complexity, whereas in our model these restrictions are attributes whose values determine the customer choices.

In summary, while many attempts have been made to understand the impact of choice behavior on traditional yield management methods and to develop simple heuristics that partially capture buy-up and buy-down behavior, to date there is no methodology that directly and completely addresses the problem. In this paper, we develop a methodology that we believe fills this void. We analyze a single-leg yield management problem in which we explicitly model consumer choice behavior using a multinomial logit (MNL) model, which is a form of random utility model. The MNL is both a theoretically sound and empirically well-tested model of consumer choice behavior. (See Ben-Akiva and Lerman [8].) As mentioned above, Andersson [1] and colleagues at SAS have also applied the MNL to estimate airline passenger choice with encouraging results. Thus, it is a natural candidate for a choice-based optimization model.

Given this MNL model of consumers, we then formulate the single-leg, multiple-fare-class yield management problem as one of selecting a subset of fare classes to offer at each point in time. We derive optimality conditions for the resulting dynamic program. While the policy might appear to be potentially complex under this model, we show that a simple, nested allocation policy is optimal at each point in time - a policy no more complex than that of Lee and Hersh [24]. This means that the policy produced by our method can easily be implemented in current, single-leg systems that rely on nested allocations or booking limits.

We also develop a practical estimation procedure for our model. One major difficulty in estimating choice models in the yield management setting is that one typically cannot observe no-purchase decisions. In many industries where yield management is actively practiced, the sale transactions are conducted remotely and the only available data are purchase transactions. For example, consider a airline ticket purchase: a customer or agency can look at the availability and fares on a computerized reservation system and, given the available fares, the customer may decide not to take any flight; an airline has no way of recording this no-purchase event. Thus, it is impossible to distinguish between periods with no arrival and periods in which there was an arrival and the arriving customer decided not to purchase. We overcome this incomplete data problem by applying the expectation-maximization (EM) method of Dempster et al. [15] to the traditional maximum-likelihood discrete-choice parameter estimation. The method allows us to simultaneously estimate both the parameters of the choice model and the arrival rates using only transaction data on sales. Together, our estimation procedure and optimization model provide a theoretically sound, consistent, practical and complete approach to the problem.

The remainder of the paper is organized as follows: In Section 1 we define the choice-based model of the problem. In Section 2 we formulate a dynamic program and analyze the resulting optimal policy. Section 3 describes our EM-based estimation procedure. Finally, some brief numerical examples are given in Section 4 and our conclusions are given in Section 5.

1 Model

Time is discrete and indexed by t, and the indices run backwards in time (e.g. smaller values of t represent later points in time). In each period there is at most one arrival. The probability of arrival is denoted by λ, which we assume is the same for all time periods t. (Extending the results to time-varying arrival probabilities is straightforward but cumbersome; we omit the details to simplify the exposition.) There are n fare classes and $N = \{1, ..., n\}$ denotes the entire set of fare classes. Each fare class $j \in N$ has an associated revenue r_j, and without loss of generality we index fare classes so that $r_1 \geq r_2 \geq ...r_n \geq 0$.

In each period t, the airline must choose a subset $S_t \subseteq N$ of fare classes to offer. When the fares S_t are offered, the probability that a customer chooses class $j \in S_t$ is denoted $P_j(S_t)$. We let $j = 0$ denote the no-purchase choice; that is, the event that the customer does not purchase any of the fares offered in S_t. $P_0(S_t)$ denotes the no-purchase probability. The probability that a sale of class j is made in period t is therefore $\lambda P_j(S_t)$, and the probability that no sale is made is $\lambda P_0(S_t) + (1 - \lambda)$. (Note this last expression reflects the fact that having no sales in a period could be due either to no arrival at all or an arrival that does not purchase; as mentioned, this leads to an incomplete data problem in practice.)

The choice probabilities $P_j(S)$ are assumed to follow a MNL model. (See Ben-Akiva and Lerman [8] for a very readable, comprehensive reference on discrete choice models including the MNL.) In the MNL, consumers are utility maximizers and the utility of each choice is a random variable. Modeling utility as random can reflect either heterogeneity in preference among individual consumers or the presence of unobservable explanatory variables in the utility. In either case, as a result of uncertainty in the utilities, the choice outcome of any given consumer is uncertain.

Formally, the utility of each alternative j is assumed to be of the form

$$U_j = u_j + \xi_j$$

where u_j is the mean utility of choice j and ξ_j is an i.i.d., Gumbel random noise term with mean zero and scale parameter one for all j. Because utility is an ordinal measure, the assumption of zero mean and a scale parameter of one are without loss of generality; (see Ben-Akiva and Lerman [8].) Similarly, the no-purchase utility is assumed to be

$$U_0 = u_0 + \xi_0,$$

where and ξ_0 is also Gumbal with mean zero and scale parameter one. Again, since utility is ordinal, without loss of generality we can assume $u_0 = 0$. To estimate the mean utility, it is common to model it as a linear function of several know attributes (e.g., price, indicator variables for product restrictions, etc.), much as one would do in a linear regression model. Thus, we assume $u_j = \beta^T x_j$ where x_j is a vector of known attributes of choice i and β is a vector of weights on these variables. The weights β can be estimated from historical data on choice outcomes via maximum likelihood estimation. (See Section 3.)

Under this utility model, one can show (See Ben-Akiva and Lerman [8] for a derivation.) that the choice probabilities are given by

$$P_j(S_t) = \frac{e^{u_j}}{\sum_{i \in S_t} e^{u_i} + e^{u_o}}, \quad j \in S \ \text{ or } \ j = 0. \tag{1}$$

For notational convenience, we define $v_j = e^{u_j}, j = 0, 1..., n$, so that the choice probabilities can be expressed as

$$P_j(S) = \frac{v_j}{\sum_{i \in S} v_i + 1}, \quad j \in S \ \text{ or } \ j = 0. \tag{2}$$

(Recall, $u_0 = 0$ so $v_0 = 1$.) Note that since e^x is monotone increasing in x, higher values of u_j imply higher values of v_j.

2 Optimization

We next formulate a single-leg problem based on this choice model. Let C denote the aircraft capacity, T denote the number of time periods, t denote the number of remaining periods (recall time is indexed backwards) and x denotes the number of remaining seats. Define the value function $V_t(x)$ as the maximum expected revenue obtainable from periods $t, t-1, ..., 1$ given that there are x seats remaining at time t. Then the Bellman equation for $V_t(x)$ is

$$\begin{aligned} V_t(x) &= max_{S_t \subseteq N} \left\{ \sum_{j \in S_t} \lambda P_j(S_t)(r_j + V_{t-1}(x-1)) + (\lambda P_0(S_t) + 1 - \lambda) V_{t-1}(x) \right\} \\ &= max_{S_t \subseteq N} \left\{ \sum_{j \in S_t} \lambda P_j(S_t)(r_j - \Delta V_{t-1}(x)) \right\} + V_{t-1}(x), \end{aligned} \tag{3}$$

where $\Delta V_{t-1}(x) = V_{t-1}(x) - V_{t-1}(x-1)$ denotes the marginal cost of capacity, and we have used the fact that for all S,

$$\sum_{j \in S} P_j(S) + P_0(S) = 1.$$

The boundary conditions are

$$V_t(0) = 0, \quad t = 1, ..., T \quad \text{and} \quad V_0(x) = 0, \quad x = 1, ..., C. \tag{4}$$

A sequence of sets $S_t^*(x)$ achieving the maximum in (3) forms an optimal Markovian policy (*cf.* Bellman [3] and Bertsekas [9]).

Potentially, each optimization on the right hand side of (3) could require an evaluation of all 2^n subsets. However, we show next that the search can be reduced to an evaluation of only n sets. This result will also be used to show that the optimal policy has a form similar to the "nested allocation" schemes of traditional single-leg models.

Theorem 1 *Let $A_k = \{1, 2, .., k\}$ denote the set of the k highest fare classes. Then there is an optimal policy that consists of only opening fares in the sets A_k, $k = 1, ..., n$.*

Proof

Note that using (2) the optimization on the right hand side of (3) can be expressed as a nonlinear, integer (binary) program

$$
\begin{aligned}
IP) \qquad & \max \frac{\sum_{j=1}^{n} z_j v_j (r_j - c)}{\sum_{j=1}^{n} z_j v_j + 1} \\
& \text{s.t} \\
& z_j \in \{0, 1\}, \quad j = 1, ..., n,
\end{aligned}
$$

where to simplify notation we let $c = \Delta V_{t-1}(x)$ denote the marginal value of capacity at the next time period. Consider the continuous relaxation of this problem

$$
\begin{aligned}
NLP) \qquad & \max \frac{\sum_{j=1}^{n} y_j (r_j - c)}{\sum_{j=1}^{n} y_j + 1} \\
& \text{s.t} \\
& 0 \cdot \ y \cdot \ v,
\end{aligned}
$$

where $y = (y_1, ..., y_n)$ and $v = (v_1, ..., v_n)$. We will show this continuous relaxation NLP has an optimal solution of the form $y^* = (v_1, ..., v_k, 0, ..., 0)$. Since this solution is equivalent to the feasible solution $z_1 = z_2 = ... = z_k = 1, z_{k+1} = z_{k+2} = ... = z_n = 0$ of IP, this will imply that the relaxation is tight and that A_k is an optimal solution to (3).

To begin, for a real number w, satisfying $0 \cdot \ w \cdot \ \sum_{j=1}^{n} v_j$, consider the subspace $H(w) = \{y | \sum_{j=1}^{n} y_j = w\}$. Note on $H(w)$, the denominator in the objective function of (5) is constant. Thus, since the numerator is linear, the problem is reduced to a continuous knapsack problem on the subspace $H(w)$, and the optimal solution is the greedy solution: assign as much of the "mass" w as possible to the highest value class (Class 1 with value $r_1 - c$) until the constraint v_1 is reached, then move on to next highest value class (Class 2 with value $r_2 - c$) until the constraint v_2 is reached, etc., stopping when the full mass w is completely allocated. Since an optimal solution, y^*, to (5) corresponds to a particular value $w^* = \sum_{j=1}^{n} y_j^*$, it follows that the optimal solution is also of this form. Namely,

$$
y^* = (v_1, ..., v_k, \alpha^*, 0, ..., 0),
$$

where $\alpha^* \doteq w^* - \sum_{j=1}^{k} v_j < v_{k+1}$. If $\alpha^* = 0$, the relaxation (5) is tight as claimed and A_k is the optimal solution in the DP recursion (3).

To analyze α^*, consider the function

$$
g_k(\alpha) = \frac{a_k \alpha + b_k}{\alpha + d_k} \tag{5}
$$

where $a_k = r_{k+1} - c$, $b_k = \sum_{j=1}^{k} v_j (r_j - c)$ and $d_k = \sum_{j=1}^{k} v_j + 1$. From (5), one can see that $g_k(\alpha)$ is the value of the solution $y = (v_1, ..., v_k, \alpha, 0, ..., 0)$ to the relaxation NLP. It therefore follows that α^* must maximize the function $g_k(\alpha)$ on $[0, v_{k+1}]$, else y^* would not be optimal for NLP. Note however that the derivative of $g_k(\cdot)$

$$g_k'(\alpha) = \frac{a_k d_k - b_k}{(\alpha + d_k)^2}, \tag{6}$$

has a sign that is independent of α. Thus, $g_k(\alpha)$ is monotone. If $a_k d_k - b_k > 0$, then $g_k(\alpha)$ is an increasing function on $[0, v_{k+1}]$ and is therefore maximized at the right end point v_{k+1}, which contradicts the assumption $\alpha^* < v_{k+1}$. If $a_k d_k - b_k < 0$, then $g_k(\alpha)$ is a decreasing function and is maximized at the left end point 0, which implies $\alpha^* = 0$. If $a_k d_k - b_k = 0$, then the function $g_k(\alpha)$ is constant on $[0, v_{k+1}]$, so any value of α is optimal. In particular, $\alpha = 0$ is optimal. Therefore, there always exists an optimal solution with $\alpha^* = 0$, so that $y^* = (v_1, ..., v_k, 0, ..., 0)$ is an optimal solution to the relaxation (5) for some value k. Since this solution is also feasible for the original integer program (5), the theorem is proven. $\square$

Note that as a result of Theorem 1, the DP recursion can be simplified to

$$V_t(x) = max_{k \in N} \left\{ \sum_{j=1}^{k} \lambda P_j(A_k)(r_j - \Delta V_{t-1}(x)) \right\} + V_{t-1}(x),. \tag{7}$$

This is solvable in $O(nCT)$ time, where recall T is the number of time periods and C is the capacity. For example, a problem with $n = 10$ fare classes, $C = 100$ seats and $T = 1,000$ time periods would require 10^6 steps, each of which requires only a simple comparison of real numbers if properly implemented (e.g. by storing the probabilities $P_j(A_k)$ rather than computing them at each step).

Let $k^*(x)$ denote the smallest index of a set A_k that optimizes the DP recursion (7). It is useful to know how the optimal index $k^*(x)$ changes as a function of the marginal value $\Delta V_{t-1}(x)$:

Theorem 2 *The optimal index $k^*(x)$ is decreasing in $\Delta V_{t-1}(x)$.*

Proof

Define $g_k(\alpha)$ by (5) as in the proof of Theorem 1 and recall that $g_k(\alpha)$ corresponds to the value of the solution $y = (v_1, ..., v_k, \alpha, 0, ..., 0)$ to (5). Note that we can write

$$g_k'(\alpha) = \frac{r_{k+1} - \sum_{j=1}^{k}(r_j - r_{k+1})v_j - c}{(\alpha + \sum_{j=1}^{k} v_j + 1)^2}. \tag{8}$$

where $c = \Delta V_{t-1}(x)$ as before. As before, since the sign of $g_k'(\alpha)$ does not depend on α, this implies that $g_k(\alpha)$ is monotone on the interval $[0, v_{k+1}]$. We next analyze the sign of $g_k'(\alpha)$. Since the denominator of $g_k'(\alpha)$ is always positive, its sign is determined by the sign of the numerator.

Note that the numerator of $g_k'(\alpha)$ is decreasing in both k and c. That it is decreasing in c is trivial. That it is decreasing in k follows since: i) r_{k+1} is decreasing in k, ii) each term,

$(r_j - r_{k+1})v_j$, in the sum (which is subtracted) is increasing in k, and iii) these terms are nonnegative ($r_j \geq r_{k+1}$ for $j \cdot\ k$) and the number of terms increases with k.

From the proof of Theorem 1, a necessary condition for k^* to be the smallest optimal index is that $g'_{k^*}(0) \cdot\ 0$ (else, k^*+1 is optimal) and $g'_{k^*-1}(0) > 0$ (else k^*-1 is the smallest optimal index). Since $g'_k(0)$ is decreasing in k, it therefore follows that $g'_k(0) > 0$ for all $k = 1, ..., k^* - 1$ and $g'_k(0) \cdot\ 0$ for all $k \geq k^*$. Thus $k^*(x)$ can be defined as

$$k^*(x) = \min\{k : g'_k(0) \cdot\ 0\}.$$

Now, since $g'_k(\alpha)$ is also decreasing in c for all k, it follows that k^*(x) is decreasing in c as well. □

We show next that the marginal values are decreasing in the remaining capacity x. Combined with Theorem 2, this implies that the optimal index $k^*(x)$ is decreasing in x as well. Therefore, at any point in time, it is optimal to sequentially close down fare classes from lowest to highest as capacity is consumed. Because of this property, the optimal policy can be implemented as a nested allocation policy at each point in time, where Class 1 is always open, Class 1 and 2 are open as long as x exceeds a "protection level" θ_1, Classes 1,2 and 3 are open as long as x exceeds a second protection level $\theta_2 \geq \theta_1$, etc. (Alternatively, one could construct "booking limits" $B_j = C - \theta_j$ instead of protection levels as described in Belobaba [5]and Brumelle and McGill [11].) Unlike the static models studied by Belobaba [5], Brumelle and McGill [11], Curry [14] and Robinson [32], however, the protection levels in our policy will change with time t. This is identical to the behavior one encounters in other dynamic models, for example Lee and Hersh [24]. Nevertheless, it is likely that protection levels will not change drastically with t, in which case static booking levels would be optimal (or approximately optimal) over short periods of time. Moreover, in practice protection levels (or allocations) are updated frequently even when using static models to reflect changes in forecasts. Thus, we believe that the choice-based optimal policy can rather easily be incorporated within the nested allocation control methods currently used by many airlines.

We next provide the formal proof of the decreasing marginal value property:

Theorem 3 $\Delta V_t(x) \cdot\ \Delta V_t(x-1), \quad t = 1, ..., T, \quad x = 1, ..., C$

<u>Proof</u>
The proof is by induction on t. First, the statement is trivially true for $t = 0$ by the boundary conditions (4). Assume it is true for period $t - 1$. From 3 and the definition of $k^*(x)$,

$$\begin{aligned} \Delta V_t(x) - \Delta V_t(x-1) &= (\Delta V_{t-1}(x) - \Delta V_{t-1}(x-1)) \\ &\quad + \sum_{j=1}^{k^*(x)} \lambda P_j(A_{k^*(x)})(r_j - \Delta V_{t-1}(x)) \end{aligned}$$

$$
\begin{aligned}
&- \sum_{j=1}^{k^*(x-1)} \lambda P_j(A_{k^*(x-1)})(r_j - \Delta V_{t-1}(x-1)) \\
&- \sum_{j=1}^{k^*(x-1)} \lambda P_j(A_{k^*(x-1)})(r_j - \Delta V_{t-1}(x-1)) \\
&+ \sum_{j=1}^{k^*(x-2)} \lambda P_j(A_{k^*(x-2)})(r_j - \Delta V_{t-1}(x-2)) \qquad (9)
\end{aligned}
$$

From the optimality of the set defined by $k^*(\cdot)$, the following inequalities hold:

$$
\sum_{j=1}^{k^*(x-1)} \lambda P_j(A_{k^*(x-1)})(r_j - \Delta V_{t-1}(x-1)) \geq \sum_{j=1}^{k^*(x)} \lambda P_j(A_{k^*(x)})(r_j - \Delta V_{t-1}(x-1))
$$

and

$$
\sum_{j=1}^{k^*(x-1)} \lambda P_j(A_{k^*(x-1)})(r_j - \Delta V_{t-1}(x-1)) \geq \sum_{j=1}^{k^*(x-2)} \lambda P_j(A_{k^*(x-2)})(r_j - \Delta V_{t-1}(x-1))
$$

Substituting into (9) we obtain

$$
\begin{aligned}
\Delta V_t(x) - \Delta V_t(x-1) \quad \cdot \quad & (\Delta V_{t-1}(x) - \Delta V_{t-1}(x-1)) \\
&+ \sum_{j=1}^{k^*(x)} \lambda P_j(A_{k^*(x)})(r_j - \Delta V_{t-1}(x)) \\
&- \sum_{j=1}^{k^*(x)} \lambda P_j(A_{k^*(x)})(r_j - \Delta V_{t-1}(x-1)) \\
&- \sum_{j=1}^{k^*(x-2)} \lambda P_j(A_{k^*(x-2)})(r_j - \Delta V_{t-1}(x-1)) \\
&+ \sum_{j=1}^{k^*(x-2)} \lambda P_j(A_{k^*(x-2)})(r_j - \Delta V_{t-1}(x-2))
\end{aligned}
$$

Rearranging and canceling terms yields

$$
\begin{aligned}
\Delta V_t(x) - \Delta V_t(x-1) \quad \cdot \quad & (1 - \sum_{j=1}^{k^*(x)} \lambda P_j(A_{k^*(x)}))(\Delta V_{t-1}(x) - \Delta V_{t-1}(x-1)) \\
&+ \sum_{j=1}^{k^*(x-2)} \lambda P_j(A_{k^*(x-2)})(\Delta V_{t-1}(x-1) - \Delta V_{t-1}(x-2))
\end{aligned}
$$

By induction, $\Delta V_{t-1}(x) - \Delta V_{t-1}(x-1) \cdot\ 0$ and $\Delta V_{t-1}(x-1) - \Delta V_{t-1}(x-2) \cdot\ 0$. Therefore, $\Delta V_t(x) - \Delta V_t(x-1) \cdot\ 0$. □

Note that the result of Theorem 3 is independent of the MNL assumptions; it holds for any choice probabilities and whenever the optimization and control are performed optimally according to the DP recursion 3. Thus, marginal values are decreasing under any model of choice behavior, which is intuitive. (Nonmonotonicty of the marginal values could occur if there is demand for multiple seats (group requests); See for example Kleywegt and Papastavrou [22], Lee and Hersh [24] and Young and Van Slyke [44]). The nested allocation structure of Theorems 1 and 2, however, requires the MNL assumptions.

2.1 Additional comments on the optimality conditions

By reexamining the optimality condition (6) we can gain some additional insight into the choice-based policy. Note from the proof of Theorem 1, that it is optimal to open fare class $k+1$ if and only if $g_k'(0) \geq 0$. By rearranging (6) and using (1), we can express this condition as

$$r_{k+1} - \Delta V_{t-1}(x) \geq \sum_{j=1}^{k} P_j(A_k)(r_j - \Delta V_{t-1}(x)). \tag{10}$$

This expression is intuitive: The left hand side is the "net gain" from selling class $k+1$; that is, the revenue we get from class $k+1$ minus the opportunity cost, $\Delta V_{t-1}(x)$, of using a unit of capacity. The right hand side is the expected net gain from offering only classes $A_k = \{1, ..., k\}$ (e.g. the sum over all classes in A_k of the probability that a customer chooses i from A_k times the net gain from selling i.). The condition (10) simply says that if the net gain from selling $k+1$ is more profitable than the gamble of offering only A_k, then it pays to open $k+1$; else, $k+1$ should be closed.

The expression (10) should be compared to the optimality condition for traditional yield management models (See Lee and Hersh [24].); namely, it is optimal to open class $k+1$ if and only if

$$r_{k+1} - \Delta V_{t-1}(x) \geq 0$$

Note the right hand side above is zero while the right hand side of (10) is positive. This happens because in the traditional model if we close class $k+1$, we lose all demand for that class. Therefore it is optimal to accept fare $k+1$ whenever r_{k+1} exceeds the opportunity cost $\Delta V_{t-1}(x)$. However, in the choice-based model, if we close class $k+1$ customers chose from among the other classes that are offered (e.g. from A_k). Hence, the threshold on the right hand side of (10) is greater than zero. This difference reflects the fact that customers may "buy-up" to a higher class.

3 Estimation

We next consider the problem of estimating the model parameters β and λ from historical data. Estimation of the MNL model given a complete set of choice data is a well-studied problem. In particular, the maximum likelihood estimate (MLE) has good computational properties (Its log is jointly concave in most cases; See McFadden [27].), and the method has proved robust in practice. (See Ben-Akiva and Lerman [8] for further discussion and case examples.)

In our case, we have an arrival probability as well as choice parameters to estimate. However, given complete observations, estimation for our model is only a slight modification of the MNL case. In particular, let D denote a set of intervals, indexed by t, in which independent arrival events and choice decisions have been observed. The set D could combine intervals from many flight departures and, deviating somewhat from our notational convention thus far, t here does not necessarily represent the time remaining for a particular flight. For each period $t \in D$, let

$$a_t = \begin{cases} 1 & \text{if customer arrives in period } t \\ 0 & \text{otherwise} \end{cases}$$

Let A denotes the set of periods t with arrivals ($a_t = 1$) and $\bar{A} = D - A$ denote the periods with no arrivals. If $t \in A$, let $j(t)$ denote the choice made by the arriving customer. (For $t \in \bar{A}$ define $j(t)$ arbitrarily.) Finally, as before let S_t denote the set of open fare classes in interval t. The likelihood function is then

$$\prod_{t \in D} \left[\lambda \frac{e^{\beta^T x_{j(t)}}}{\sum_{i \in S_t} e^{\beta^T x_i} + 1} \right]^{a_t} (1-\lambda)^{(1-a_t)}$$

Taking logs, we obtain the log-likelihood function

$$\mathcal{L} = \sum_{t \in D} \left[a_t \left(\beta^T x_{j(t)} - \ln(\sum_{i \in S_t} e^{\beta^T x_i} + 1) \right) + a_t \ln(\lambda) + (1-a_t) \ln(1-\lambda) \right]. \tag{11}$$

Note that $\mathcal{L}$ is separable in β and λ. Maximizing $\mathcal{L}$ with respect to λ, we obtain the estimate

$$\hat{\lambda} = \frac{1}{|D|} \sum_{t \in D} a_t = \frac{|A|}{|D|},$$

where $|D|$ (resp. $|A|$) denotes the cardinality of D (resp. $|A|$). The MLE, $\hat{\beta}$, is then determined by solving

$$\max_{\beta} \sum_{t \in A} \left(\beta^T x_{j(t)} - \ln(\sum_{i \in S_t} e^{\beta^T x_i} + 1) \right) \tag{12}$$

This is simply the usual maximum likelihood problem for the MNL applied to those periods with customer arrivals. Combining these two estimates gives the MLE for our model with complete data.

As mentioned, the difficulty with this approach in practice is that one rarely observes all arrivals. Typically, only purchase transaction data are available. Thus, it is impossible to distinguish a period without an arrival, from a period in which there was an arrival but the arriving customer did not purchase. With this incompleteness in the data, the above MLE procedure cannot be used.

One can write down the ML formula for estimating the discrete-choice parameters with incomplete data, but as often happens in such cases, the function becomes very complex (and non-concave) and difficult to maximize. To overcome this problem, we propose using the expectation-maximization (EM) method of Dempster et al. [15]. The method works by starting with arbitrary initial estimates, $\hat{\beta}$ and $\hat{\lambda}$. These estimates are then used to compute the conditional expected value of $\mathcal{L}$: $E[\mathcal{L}|\hat{\beta},\hat{\lambda}]$ (the expectation step). The resulting expected log-likelihood function is then maximized to generate new estimates $\hat{\beta}$ and $\hat{\lambda}$ (the maximization step) and the procedure is repeated until it converges. While it is true that technical convergence problems can arise, in practice the EM method is a robust and efficient way to compute maximum likelihood estimates for incomplete data. (See McLachlan and Krishnan [30] for a comprehensive reference on the EM method.) Moreover, it has also been used in other yield management contexts, in particular by McGill [28] to estimate multi-variate normal demand data with censoring.

To apply the EM method in our case, let P denote the set of periods in which customers purchase and $\bar{P} = D - P$ denote period in which there are no purchase transactions. We can then write the complete log-likelihood function as

$$\begin{aligned}\mathcal{L} &= \sum_{t\in P}\left[\ln(\lambda) + \beta^T x_{j(t)} - \ln(\sum_{i\in S_t} e^{\beta^T x_i} + 1)\right] \\ &\quad + \sum_{t\in\bar{P}}\left[a_t\left(\ln(\lambda) - \ln(\sum_{i\in S_t} e^{\beta^T x_i} + 1)\right) + (1-a_t)\ln(1-\lambda)\right]. \end{aligned} \tag{13}$$

The unknown data are the values $a_t,\ t\in\bar{P}$ in the second sum. However, given estimates $\hat{\beta}$ and $\hat{\lambda}$, we can determine their expected values (denoted $\hat{a}_t$) easily via Bayes's rule:

$$\begin{aligned}\hat{a}_t \doteq E[a_t|t\in\bar{P},\hat{\beta},\hat{\lambda}] &= P(a_t = 1|t\in\bar{P},\hat{\beta},\hat{\lambda}) \\ &= \frac{P(t\in\bar{P}|a_t=1,\hat{\beta},\hat{\lambda})P(a_t=1|\hat{\beta},\hat{\lambda})}{P(t\in\bar{P}|\hat{\beta},\hat{\lambda})} \\ &= \frac{\hat{\lambda}P_0(S_t|\hat{\beta})}{\hat{\lambda}P_0(S_t|\hat{\beta}) + (1-\hat{\lambda})}, \end{aligned} \tag{14}$$

where

$$P_0(S_t|\hat{\beta}) = \frac{1}{\sum_{i\in S_t} e^{\hat{\beta}^T x_j} + 1}$$

is the no-purchase probability for observation t given $\hat{\beta}$.

Substituting $\hat{a}_t$ into (13) we obtain the expected log-likelihood for the incomplete data

$$E[\mathcal{L}|\hat{\beta},\hat{\lambda}] = \sum_{t\in P}\left[\beta^T x_{j(t)} - \ln(\sum_{i\in S_t} e^{\beta^T x_i} + 1)\right] - \sum_{t\in\bar{P}} \hat{a}_t \ln(\sum_{i\in S_t} e^{\beta^T x_i} + 1)$$

$$+\sum_{t\in P}\ln(\lambda)+\sum_{t\in\bar{P}}\left(\hat{a}_t\ln(\lambda)+(1-\hat{a}_t)\ln(1-\lambda)\right). \tag{15}$$

As in the case of the complete log-likelihood function, this function is separable in β and λ. Maximizing with respect to λ we obtain the updated estimate

$$\lambda^* = \frac{|P|+\sum_{t\in\bar{P}}\hat{a}_t}{|P|+|\bar{P}|}. \tag{16}$$

This is intuitive; our estimate of lambda is the number of observed arrivals, $|P|$, plus the estimated number of arrivals from unobservable periods, $\sum_{t\in\bar{P}}\hat{a}_t$, divided by the total number of periods $|P|+|\bar{P}|=|D|$. We can then maximize the first two sums in (15) to obtain the updated estimate β^*. Note that this expression is of the same functional form as the complete data case (12). The entire procedure is then repeated.

Summarizing the algorithm:

Step 0: Initialize $\hat{\beta}$ and $\hat{\lambda}$.

Step 1: Expectation step

- For $t\in\bar{P}$, use the current estimates $\hat{\beta}$ and $\hat{\lambda}$ to compute $\hat{a}_t$ from (14).

Step 2: Maximization step

- Compute λ^* using (16).
- Compute β^* by solving

$$\max_{\beta}\left\{\sum_{t\in P}\left(\beta^T x_{j(t)}-\ln(\sum_{i\in S_t}e^{\beta^T x_i}+1)\right)-\sum_{t\in\bar{P}}\hat{a}_t\ln(\sum_{i\in S_t}e^{\beta^T x_i}+1)\right\}$$

Step 3: Convergence test

- IF $\|(\hat{\lambda},\hat{\beta})-(\lambda^*,\beta^*)\|<\epsilon$, THEN STOP;
- ELSE $\hat{\lambda}\leftarrow\lambda^*$, $\hat{\beta}\leftarrow\beta^*$ and GOTO Step 1.

Since the expected log-likelihood, $E[\mathcal{L}|\hat{\beta},\hat{\lambda}]$, given in (15) is continuous in both (β,λ) and $(\hat{\beta},\hat{\lambda})$, a result by Wu [43] shows that if the sequence of estimates converges, the resulting value will be a stationary point of the incomplete log-likelihood function. Whether the sequence diverges - or converges to something other than the global maximum - is more difficult to determine. In practice, the method has proved to be very robust in other contexts and this has been our experience in simulated experiments for our problem. (See McLachlan and Krishnan [30] for further discussion of convergence properties of the EM algorithm.)

4 Numerical Example

In this section we provide results of a small simulation study to compare our choice-based method to a traditional single-leg method. Both the estimation and optimization methods were tested. In the simulation study, we made the assumption that demand in fact coincides with the MNL model of this paper. This is admittedly a somewhat optimistic assumption and does not address the validity of the choice model itself. While validating the MNL model is certainly important both from a theoretical and implementation standpoint, it would require extensive empirical testing against alternate demand models, and such a study is beyond the scope of this paper. At the same time, testing the model under MNL behavior gives some indication of its potential; in particular, if such a test shows that the choice-based model performs significantly better than traditional methods, then one cannot reject the hypothesis that it is a potentially superior approach. A more positive conclusion than this will require extensive empirical and simulation testing and actual implementation experience.

The traditional model we tested against was Belobaba's EMSR-b heuristic [6]. This is one of the most common seat protection heuristics used in practice. It is a *fixed protection level* policy, which sets a static set of protection levels for fare classes 1 through $n-1$ given by the vector $\theta = (\theta_1, \ldots, \theta_{n-1})$, where $\theta_1 \cdot \;\; \theta_2 \cdot \;\; \ldots \theta_{n-1}$. (There is no protection level for the lowest fare class, n.) Protection levels are nested in the sense that θ_j represents the number of seats to reserve (protect) for all of fare classes $1, 2, \ldots, j$. Reservations for fare class $j+1$ are accepted if and only if the number of seats remaining is strictly greater than the protection limit θ_j. (Such policies are optimal when low fare classes book strictly before higher fare classes and fare class demands are mutually independent [11].)

EMSR-b sets the protection levels θ_j as follows: Given estimates of the means, $\hat{\mu}_j$, and standard deviations, $\hat{\sigma}_j$, for each fare class j, the EMSR-b heuristic sets θ_j so that

$$r_{j+1} = \bar{r}_j P(\bar{X}_j > \theta_j),$$

where $\bar{X}_j$ is a normal random variable with mean $\sum_{i=1}^{j} \hat{\mu}_i$ and variance $\sum_{i=1}^{j} \hat{\sigma}_i^2$, and $\bar{r}_j$ is a weighted average revenue, given by

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

The idea behind this approximation is to reduce the complexity of the fully nested problem by aggregating fare classes $1, 2, .., j$ into a single fare class. Then, one treats the problem as a simple, two-fare-class problem and applies Littlewood's rule [26].

Converting the choice model parameters into inputs for the EMSR-b method is somewhat tricky. While there are several options for doing this, current industry practice, by and large, is to forecast demand in each fare-class independently (using, for example, time-series methods), and if a particular fare-class is closed to unconstrain (uncensor) the demand by using either the EM algorithm or other methods. One interpretation of this practice is that

j	1	2	3	4	5	6	7	8	9	10
r_j	600	550	475	400	300	280	240	200	185	175
v_j^L	0.407	0.438	0.490	0.549	0.638	0.657	0.698	0.741	0.758	0.769
v_j^H	0.050	0.064	0.093	0.135	0.223	0.247	0.301	0.368	0.397	0.417

Table 1: Fares and v_j values for numerical example

it is an attempt to forecast demand for each fare-class when all other fare-classes are open. We adopted this interpretation lacking any compelling alternative. Therefore, the demand inputs we used in our simulations is the demand for class i given that all other classes are open. Specifically, $\mu_j = \lambda T P_j(N)$ and $\sigma_j^2 = T\lambda P_j(N)(1 - \lambda P_j(N))$. (Demand is binomial with T trials and probability $\lambda P_j(N)$ of success.)

To test the two methods, we simulated arrivals and applied each method to control fare class availability. The capacity was $C = 185$ seats and there were $n = 10$ fare-classes with fares as shown in Table 1. For simplicity, we assumed the random utility had only one attribute, x, which was simply the price. The co-efficient, β, on the price attribute was taken to be either $\beta^L = -0.0015$ (low price sensitivity, denoted L) or $\beta^H = -0.005$ (high price sensitivity, denoted H). The values $v_j^L = e^{\beta^L x}$ and $v_j^H = e^{\beta^H x}$ are shown in Table 1 as well.

Arrivals over the booking period were generated by simulating a homogeneous Poisson process with mean 205. (Thus, if the booking period is broken up into intervals of size Δ, then $\lambda = 205\Delta$ in the choice-based DP.) The choice parameter as estimated using the EM method as described in Section 3. The training set consisted of 50 simulated days during which the available classes were controlled using EMSR-b as described above. The EM method produced an estimates of $\hat{\beta}^L = -0.0014$ and $\hat{\beta}^H = -0.0048$, which are very close to the actual values of $\beta^L = -0.0015$, $\beta^H = -0.005$. To mimic the real-world combination of forecasting and optimization, these estimated value was used in the choice DP algorithm.

Bookings were generated for a booking period of 100 simulated days and the controls from each method were applied to simulate accept/deny decisions. The results of 15 simulated flights are shown in Table 2. In the case where the price sensitivity is low, the choice DP has significantly higher revenues; 65,693 versus 53,543 for EMSR. This represents a 22.7% improvement in revenue, which is very large when compared to the typical 1-2% differences in revenues that one finds when comparing optimization methods. In the case of high price sensitivity, the revenue difference between the two methods is essentially identical within the simulation error of our test. This is not unsurprising since buy-up in particular can best be exploited when customers are not very price sensitive.

	Low Price Sensitivity		High Price Sensitivity	
	Choice DP	EMSR	Choice DP	EMSR
Avg. Revenue	65,693	53,543	36,615	36,745
Load Factor	71%	93%	66%	78%

Table 2: Simulation Results

In terms of qualitative behavior, the choice DP frequently closes lower-fare classes to force consumers to buy-up to higher fares. EMSR-b, in contrast, opens up many more fare classes and allows customers to buy at lower prices. This difference shows up in the load factors, which are lower for the choice-based DP. While this drop in load-factors may at first be worrisome, it is not unexpected given that the choice-based DP is deliberately increasing the probability that customers will not purchase by restricting discounts. However, the revenue increases from the higher fares more than compensate for these lower volume of sales in the low-price-sensitivity case and produce essentially the same revenues in the high-price-sensitivity case. While the magnitude of these results are specific to this particular set of numbers and choice probabilities, the results do show how it may be possible to increase revenue by exploiting choice behavior.

5 Conclusion

We believe that the choice-based DP in combination with the EM estimation procedure provides an appealing alternative to traditional yield management methods. The MNL logit is a well-developed model that is conceptually sound and has worked well in other application contexts. The DP model itself is no more complex than traditional models and the policy can be implemented as a nested allocation policy, so no major changes in control structure are needed. Finally, the EM estimation procedure is no more complicated than traditional, censored forecasting methods and can be applied to currently available data.

As for additional work, we see several topic worthy of further study. One is to generalize the method to other choice models, for example the nested logit model or the generalized extreme value (GEV) models [8], which might be more useful in certain settings. The EM method can probably be extended easily, but the simple nesting property may not hold for more general choice models. Another worthwhile extension would be to model choice among a set of flights. This was one of the topics investigated by Andersson at SAS [1] and it would be interesting to see if the estimation and optimization methods could be extended to model control of a set of related flights. Similar but even more complex would be to extend the model to networks. In both these cases, exact dynamic programming will most likely be impractical, so it would be interesting to see what approximation methods could be developed.

Finally, it would be worthwhile to devise both tests and estimation procedures for a heterogeneous mix of customer segments (e.g.: business and leisure) with distinct choice parameters for each segment. Theoretically, it would be interesting to extend the control results when the demand is composed of such a mix.

References

[1] Andersson, S.E. 1998. "Passenger Choice Analysis for Seat Capacity Control: A Pilot Project in Scandinavian Airlines," *Intl. Trans. Opl. Res.* , **5**, 471-486.

[2] Barnhart, C. and Talluri, K.T 1996. "Airline Operations Research," to appear in *Systems for Civil and Environmental Engineering: An Advanced Text Book* (ed. Charles ReVelle and Arthur McGarrity), John Wiley and Sons.

[3] Bellman, R. 1957. *Applied Dynamic Programming*, Princeton University Press, Princeton, N.J.

[4] Belobaba, P.P. 1987. "Air Travel Demand and Airline Seat Inventory Management," Ph.D. thesis, MIT, Cambridge, Mass.

[5] Belobaba, P.P. 1987. "Airline yield Management: An Overview of Seat Inventory Control," *Trans. Sci.* **21**, 63-73.

[6] Belobaba, P.P. 1989. "Application of a Probabilistic Decision Model to Airline Seat Inventory Control," *Oper. Res.* **37**, 183-197.

[7] Belobaba, P.P. and C. Hopperstad 1999. "Boeing/MIT Simulation Study: PODS Results Update," 1999 AGIFORS Reservations and Yield Management Study Group Symposium, April 27-30, London.

[8] Ben-Akiva, M. and S.R. Lerman, *Discrete Choice Analysis*, The MIT Press, Cambridge, Massachusetts.

[9] Bertsekas, D.P. 1995. *Dynamic Programming and Optimal Control, Volume I.*, Athena Scientific, Belmont, Mass.

[10] Bitran, G. Caldentey, R. and Mondschein, S. 1998. "Coordinating Clearance Markdown Sales of Seasonal Products in Retail Chains," *Oper. Res.* , **46**, 609-624.

[11] Brumelle, S.L., and McGill, J.I. 1993. "Airline Seat Allocation With Multiple Nested Fare Classes," *Oper. Res.* **41**, 127-137.

[12] Brumelle, S.L., McGill, J.I., Oum, T.H., Sawaki,K., Trethwway, M.W. 1990. "Allocation of Airline Seats Between Stochastically Dependent Demand," *Trans. Sci.* **24**, 183-192.

[13] Cooper, W. 2000. "Asymptotic Behavior of Some Revenue Management Policies," Univ. of Minnesota Working Paper, Minneapolis, Minn.

[14] Curry, R.E. 1989. Optimal Airline Seat Allocation with Fare Classes Nested by Origins and Destinations. *Trans. Sci.* **24**, 193-204.

[15] Dempster, A.P., N.M. Laird and D.B. Rubin 1977. "Maximum Likelihood From Incomplete Data via the EM Algorithm," *J. of the Royal Stat. Society*, B, **39**, 1-38.

[16] Diamond, M. and Stone, R (1991). Dynamic Yield Management on a Single-Leg, Unpublished Manuscript, Northwest Airlines.

[17] Dror, M., Trudeau, P. and Ladany, S.P. 1988. Network Models for Seat Allocation of Flights. *Tran. Res.* **22B**, 239-250.

[18] Feng, Y. and Gallego, G. 1995. Optimal Stopping Times for Promotional Fares and Optimal Starting Times for End-of-Season Sales. To appear in *Mgmt. Sci.*

[19] Gallego, G. and G. van Ryzin 1994. "Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons," *Mgmt. Sci.*, **40**, 999-1020

[20] Gallego, G. and G. van Ryzin 1997. "A Multi-Product, Multi-Resource Pricing Problem and Its Applications to Network Yield Management," *Oper. Res.*, **45**, 24-41.

[21] Glover, F., Glover, R., Lorenzo, J. and McMillan, C. 1982. "The Passenger-Mix Problem in the Scheduled Airlines," *Interf.* **12**, 73-79.

[22] Kleywegt, A.J. and J.D. Papastavrou 1998. "The Dynamic and Stochastic Knapsack Problem," *Operations Res.*, **46**, 17-35.

[23] Lautenbacher, C. J. and S. J. Stidham, "The Underlying Markov Decision Process in the Single-Leg Airline Yield Management Problem," *Trans. Sci.*, **34**, 136-146.

[24] Lee, T. C., and Hersh M. 1993. A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. *Trans. Sci.*, **27**, 252-265.

[25] Liang, Y. 1999. "Solution to the Continuous Time Dynamic Yield Management Model," *Trans. Sci.*, **33**, 117-123.

[26] Littlewood, K. 1972. "Forecasting and Control of Passengers," *12th AGIFORS Symposium Proceedings*, 95-128.

[27] McFadden, D. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, P. Zarembka (ed.), Academic Press, New York, 105-142.

[28] McGill, J.I. 1995. "Censored Regression Analysis of Multiclass Demand Data Subject to Joint Capacity Constraints," *Ann. Oper. Res.*, **60**, 209-240.

[29] McGill, J. and G.J. van Ryzin 1999. "Revenue Management: Research Overview and Prospects," *Trans. Sci.*, **33**, 233-256.

[30] McLachlan, G. and T. Krishnan 1996. "The EM Algorithm and Extensions," John Wiley & Sons, New York.

[31] Phillips, R. 1994. "State-Contingent Airline Yield Management," Presentation in Session TC33.4, INFORMS Detroit 1994.

[32] Robinson, L.W. (1991), "Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes," *Oper. Res.*, **43**, 252-263.

[33] Simpson, R.W. (1989), "Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control," MIT Flight Transportation Laboratory Memorandum M89-1, Cambridge, Massachusetts.

[34] Smith, B., Leimkuhler, J., Darrow, R., and J. Samuels 1992. "Yield Management at American Airlines," *Interfaces*, 22, 8-31.

[35] Stone, R. and M. Diamond 1992. "Optimal Inventory Control for a Single Flight Leg", working paper, Northwest Airlines, Operations Research Division, Minneapolis.

[36] Subramanian, C. J. Lautenbacher, and S. J. Stidham 1999. "Yield Management with Overbooking, Cancellations and No Shows," *Trans. Sci.*, **34**, 147-167.

[37] Talluri, K.T. 1996. "Revenue Management on a Line Network," Working Paper, UPF, Barcelona, Spain.

[38] Talluri, K.T. and G.J. van Ryzin 1998. "An Analysis of Bid-Price Controls for Network Revenue Management," *Mgmt. Sci.*, **. 44**, 1577-1593.

[39] Talluri, K. and G.J. van Ryzin 1999. "A Randomized Linear Programming Method for Computing Network Bid Prices," *Trans. Sci.*,**33**, 207-216.

[40] Williamson, E. L. 1988. "Comparison of Optimization Techniques for Origin-Destination Seat Inventory Control," Master's thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.

[41] Williamson, E.L. 1992. "Airline Network Seat Inventory Control: Methodologies and Revenue Impacts," Ph.D. thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.

[42] Wollmer, R.D. 1992. "An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First," *Oper. Res.* **40**, 26-37.

[43] Wu, C.F.J. 1983. "On the Convergence Properties of the EM Algorithm," *Annals of Stat.*, **11**, 95-103.

[44] Young, Y. and R. Van Slyke 1994. "Stochastic Knapsack Models of Yield Management," Technical Report, 94-76, Brooklyn Polytechnic University, 1994.

[45] Zhao, W. 1999. "Dynamic and Static Yield Management Models", Ph.D. thesis, The Wharton School, Operations and Information Management Department, University of Pennsylvania, Philadelphia.