OPERATIONS RESEARCH
Vol. 53, No. 3, May–June 2005, pp. 415–431
ISSN 0030-364X | EISSN 1526-5463 | 05 | 5303 | 0415

DOI 10.1287/opre.1050.0194

Revenue Management for Parallel Flights with Customer-Choice Behavior

Dan Zhang, William L. Cooper
Department of Mechanical Engineering, University of Minnesota, 111 Church Street S.E., Minneapolis, Minnesota 55455
{zdan@me.umn.edu, billcoop@me.umn.edu}

We consider the simultaneous seat-inventory control of a set of parallel flights between a common origin and destination with dynamic customer choice among the flights. We formulate the problem as an extension of the classic multiperiod, single-flight "block demand" revenue management model. The resulting Markov decision process is quite complex, owing to its multidimensional state space and the fact that the airline's inventory controls do affect the distribution of demand. Using stochastic comparisons, consumer-choice models, and inventory-pooling ideas, we derive easily computable upper and lower bounds for the value function of our model. We propose simulation-based techniques for solving the stochastic optimization problem and also describe heuristics based upon an extension of a well-known linear programming formulation. We provide numerical examples.

Subject classifications: dynamic programming/optimal control: applications; transportation: yield management.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received April 2003; revisions received December 2003, May 2004; accepted May 2004.

1. Introduction

Airline revenue management involves dynamically controlling the availability and prices of many different classes of tickets to maximize revenues. Despite the vast technical literature on the subject of revenue management, relatively few papers explicitly model the interaction between the random demand for tickets and the seller's choice of booking policy. In fact, the majority of work assumes that the distribution of demand for the various ticket types is exogenously determined, and is therefore not affected by the booking policy (of course, all models do assume that the policy does affect which ticket requests are accepted). Among the papers that do attempt to capture behavioral effects, most are concerned with how the customers decide which type of ticket to buy for a single specified flight, given the availability of the various tickets for the particular flight in question.

In this paper, we consider a different issue, namely, how to select booking limits when the airline has many flights between a particular origin and destination within a short timespan. A check of airline schedules (using, e.g., online services such as Orbitz or Travelocity) shows that this scenario occurs frequently in the real world. We focus on the situation in which customers select among these "parallel" flights. In particular, we address what Belobaba (1989) calls "horizontal shifts" of passengers to different flights but the same fare class.

To see the importance of this scenario, consider an airline that has two flights per day between cities A and B. Suppose that the airline sells two classes of tickets (high fare and low fare) on each flight, that nearly all high-fare customers want to fly on the flight later in the day, and that they are inflexible regarding this preference. Suppose that, on the other hand, low-fare customers are essentially indifferent between the flights and are willing to fly on either one. In such a case, the airline should shut off low-fare ticket sales on the later flight while keeping many low-fare tickets available for the earlier one. This will allow the airline to sell tickets to low-fare customers on the early flight without displacing high-fare customers on the later flight. Note that this is a different issue than that mentioned above, whereby customers select among different ticket types on a particular flight. Alternatively, as described in Karaesman and van Ryzin (2004), the airline might oversell the earlier flight and bump (at a cost) customers to the later flight. We do not consider such methods in this paper.

Currently available models do not provide a formal method for addressing problems where customers choose among parallel flights. Our Markov decision process (MDP) modeling approach to such problems is an extension of that in which there are sequentially realized "blocks" of demand for each ticket class. We employ a vector-valued state variable (with one entry per flight), and for any given period (i.e., block) we allow the distribution of sales for each flight to depend upon the booking limits in effect for *all* flights. This setup allows for fairly general customer dynamics within a period.

Because there is quite a large body of work on revenue management, we limit the scope of our literature review to variants of the expected marginal seat revenue (EMSR) model, which employ the aforementioned block-demand setup. We will also describe some recent work that explicitly models customer choice behavior in various settings. For an overall survey of the field of revenue management, we refer the reader to McGill and van Ryzin (1999) or Boyd and Bilegan (2003).

1.1. Literature Review

In some of the earliest revenue-management work, Littlewood (1972) describes a simple technique for setting a booking limit on the number of low-fare tickets available for sale in a problem with two fare classes for a single-leg flight. Later, Belobaba (1989) considers the case with multiple (more than two) booking classes. Subsequently, a number of authors have developed a framework for determining booking limits for a single-leg flight with m mutually independent demand classes that arrive in sequential blocks. Hereafter, we will term models that employ variations of the sequential-block-of-demand assumption as *EMSR-type models*. Although the block-demand assumption is not entirely realistic, these models have, in fact, proved to be useful ingredients in real-world revenue management systems.

Li and Oum (2002) discuss the rough equivalence of the optimality conditions of Brumelle and McGill (1993), Curry (1990), and Wollmer (1992), who all derive optimal policies for EMSR-type models under slightly varying assumptions. Robinson (1995) extends the analysis to blocks of arrivals with fares that need not be monotonically increasing. Several papers consider the interaction between stochastically dependent booking classes. For instance, Brumelle et al. (1990) consider two possibly dependent booking classes, and derive optimality conditions under mild assumptions on the dependence structure. Cooper and Gupta (2005) employ stochastic order relations to investigate, among other things, the effect of stochastic dependence on total expected revenue.

Lautenbacher and Stidham (1999) model the booking control problem on a single-leg flight as an MDP. They establish the connection between EMSR-type models and other one-arrival-per-period MDPs. Their analysis confirms that for a wide class of seat management problems, a booking-limit policy is indeed optimal. (Implicit in some earlier work is the consideration of *only* booking-limit policies. Under more general models of the customer arrival process, there typically will be non-booking-limit policies that outperform any booking-limit policy—see, e.g., Chatwin (1998). Nevertheless, booking-limit policies continue to dominate airline practice because many distribution channels allow only these types of policies. Consequently, one could reasonably argue that the formulation of the problem should indeed require the selection of a booking-limit policy.) In the model described in this paper, we consider only booking-limit policies.

As mentioned earlier, relatively few papers consider customer-choice behavior in the revenue management literature. Talluri and van Ryzin (2004) analyze an MDP formulation of a single-leg problem with customer choice among the open fare classes. They show that an optimal policy can be found by searching over a relatively small class of policies, and provide conditions under which there is an optimal nested booking-limit policy for the problem with customer choice. Zhao and Zheng (2001) consider a two-class single-flight model with flexible customers, and formulate the problem as an optimal stopping model. Belobaba and Weatherford (1996) propose corrections to the EMSR heuristics in the presence of customer diversion. The above-mentioned paper of Brumelle et al. (1990), in which dependent demands are analyzed, also models the "sell-up" phenomenon, whereby closing Class-2 ticket sales may cause some would-be Class-2 customers to purchase Class-1 tickets. Hence, the booking policy does induce changes in demand distribution for Class 1. This portion of their study formalizes an earlier analysis of Belobaba (1989). Other papers that model diversion from one class to another include Pfeifer (1989) and Bodily and Weatherford (1995).

Several recent papers in the inventory control literature consider customer-choice behavior in retail situations. Mahajan and van Ryzin (2001b) analyze the problem of optimizing inventory levels of substitutable products under dynamic consumer substitution. Smith and Agrawal (2000) study the effect of substitution on demand and address the issue of jointly setting stocking levels under a Markovian choice model. In these models where dynamic customer behavior is taken into account, exact analysis has proved to be difficult, causing authors to seek effective approximation methods.

Shumsky and Zhang (2004) consider a multiperiod capacity allocation model with multiple classes and service upgrading. In their model, inventory is differentiated (in terms of product or service quality), and the demand of a lower-class customer may be satisfied by a product that is one level higher in each period. In contrast to our work, they assume that prices of different products are stationary across periods, and allow the allocation in each period to be made after observing demand in that period. Their main focus is how much unsatisfied lower-class demand should be upgraded (or alternatively, how much higher-class inventory can be allocated to lower-class demand).

Netessine and Shumsky (2004) formulate a game-theoretic version of the single-leg airline revenue management problem assuming two airlines are in competition, and the demand of one airline depends on the booking policy of the other airline. They provide conditions under which a Nash equilibrium (in booking limits) exists under assumptions on consumer substitution patterns. They also consider a centralized model where the two flights are assumed to be owned by one airline. Their centralized model is somewhat similar in spirit to the model considered in our paper, although their focus is to compare the system-optimal controls from the centralized solution with the competitive outcome.

Many MDP models, including the one considered here, are too complex for exact solution or even for storage of an optimal policy. Consequently, it is of particular interest to develop methods to solve MDPs approximately—and this is the approach we take in this paper. For an overview of approximation methods for MDPs, see Bertsekas and Tsitsiklis (1996). A slightly different idea is that of Müller (1997), who investigates how the value function of an MDP changes when the transition probability distributions are changed and all the other model parameters remain the same. He employs stochastic order relations to establish the monotonicity of value functions of MDPs with respect to changes in transition probabilities. This allows one to generate bounds on the value function of the MDP; however, he does not discuss using such bounds as an ingredient in a computational procedure. White and Schlussel (1981) present bounds and approximation procedures for so-called multimodule MDPs, which possess vector state spaces for which each element of the vector evolves independently, and each module is interconnected only through the cost structure. The problem we consider in this paper does not fall into this category. Lovejoy (1986) considers a slightly different problem, and proposes several approaches to generate bounds for the optimal policy.

1.2. Overview of Results and Outline

In this paper, we consider the seat-inventory control of multiple parallel single-leg flights in the presence of dynamic customer-choice behavior. The objective is to maximize the expected total revenue over all the flights. We formulate the problem as an MDP.

In our MDP formulation, we assume that the demand in each period on each flight depends on the numbers of open seats on *all* the flights. We then derive separable upper and lower bounds using stochastic comparison results and dynamic programming principles. We adopt a customer-choice model that assumes that the choice outcome of each customer is determined by a preference mapping together with the inventory availability. Within a fixed period, the choice model we employ is equivalent to the one developed by Mahajan and van Ryzin (2001a, b). However, we emphasize a preference ordering as the starting point, whereas their starting point is a utility-maximization assumption. One difference is that we consider a multi-period model; whereas Mahajan and van Ryzin study a single-period model (in a different problem context).

The bounds we develop, however, are not limited to the particular choice model, and are applicable to other settings, as long as certain assumptions are satisfied. In addition, we develop another upper bound using inventory-pooling ideas. Because the MDP is computationally intractable, we propose several solution approaches for our model and test them with numerical examples. Some of these approaches are based upon value-function approximations derived from the upper and lower bounds. Others involve a modification of a well-known linear programming formulation for network revenue management problems. Broadly speaking, the solution techniques involve simulation-based methods for approximately solving high-dimensional MDPs.

In summary, the primary contributions of this paper are (1) formulation of the revenue management problem for parallel flights in the presence of customer-choice behavior, (2) development of upper and lower bounds for the value function of the MDP, (3) description of linear-programming-based heuristics, (4) proposal of simulation-based solution techniques for the problem, and (5) discussion of numerical work that shows the proposed approaches are promising.

The remainder of the paper is organized as follows: Section 2 provides the basic formulation. Section 3 develops upper and lower bounds for the value function. Section 4 describes static and dynamic booking-limit policies. Section 5 introduces another upper bound from inventory pooling. Section 6 describes a choice model and relates it to the developments in §§2–5. Section 7 describes Some approximate solution procedures. Section 8 provides the linear programming formulation. Section 9 includes numerical results. Section 10 contains closing remarks.

2. Markov Decision Process Formulation

Throughout, we use superscripts to denote components of a vector and subscripts to denote time; ϵ^i is a vector whose ith element is one and all other elements are zeros, and ϵ^0 is a vector with all zeros. The dimensions of vectors should be evident from the context.

We consider the booking control of multiple parallel single-leg flights between a common origin and destination. Let $N = \{1, \ldots, n\}$ be the set of n flights. The initial capacity of flight i is c^i; $i = 1, \ldots, n$. Let c denote the n-vector with entries c^i. The seats on each flight can be sold to m possible classes. Note that we assume that the number of fare classes is the same on all flights. The fare for class-j is f_j. The m classes arrive in distinct time periods, with class-j demand arriving in period j. If one wants to allow, say, class Y to arrive both before and after class Z, then we simply relabel the classes as Y, Z, and Y' to put the problem into this context. Therefore, if an airline has 10 actual booking classes, then m could be larger than 10. As we move closer to the time of departure, the time index decreases, so the first time period is $t = m$, and the last is $t = 1$.

Within a time period, we assume that customers choose among the flights (or decide not to purchase) dynamically according to some choice model; however, we delay the description of any particular choice model to §6, because such specific information is not needed at this point. We do not consider the situation where customers select what class of ticket to purchase—rather, our model takes class to be exogenous. Although one could certainly object to this assumption, it is important to point out that most existing models also make the same assumption. Those that do allow choice behavior for class typically do not allow choice among flights. Belobaba (1989) briefly alludes to such horizontal shifts of passengers to different flights, but the same fare class. However, he does not describe solution approaches for the problem.

We formulate the booking control problem as an MDP. At time t, a state s represents the number of seats sold prior to time t. Here, s is an n-vector, and the ith component s^i is the number of seats sold prior to time t on flight i. For each period, we must select an action. The actions we consider are the number of seats to be open for sale on each flight at the beginning of the period. Let x be the (integer) n-vector of the numbers of open seats in period t. We require that $0 \leqslant x^i \leqslant c^i - s^i$ for $i = 1, \ldots, n$. Formally, the action space for state s is given by $\{x \in \mathbb{Z}^n: 0 \leqslant x \leqslant c - s\}$, where the inequalities should be interpreted componentwise. (Throughout, inequalities involving vectors should be interpreted componentwise.) Sometimes we will refer to an action x as an initial availability vector for period t.

Let Q_t^i be the integer-valued (random) demand for flight i in period t—shortly, we shall say more precisely what this means. Assume that conditional upon action x being selected in period t, the random n-vector Q_t is independent of the "past history" of the process. However, the conditional distribution of Q_t, given an action x, *does* depend upon x. Let $F_x^t(\cdot)$ denote this n-dimensional conditional distribution function. In other words, if for $i = 1, \ldots, n$ we make x^i seats available for flight i in period t, then the vector of demand will have conditional distribution $F_x^t(\cdot)$. At this juncture, it is important to note that by allowing this distribution to depend upon x, we are allowing demand to depend upon seat availability in a rather general manner. As we pointed out earlier, few revenue management models have previously allowed any such interactions.

Let X_t denote the action selected for period t. To simplify developments below, let $Q_t(x)$ denote a random vector with distribution $F_x^t(\cdot)$; i.e., $\mathrm{P}(Q_t \leqslant q \mid X_t = x) = \mathrm{P}(Q_t(x) \leqslant q) = F_x^t(q)$. Note that $Q_t^i(x)$, the demand for flight i when the action is x, depends not only on the number of seats open for flight i, but also on the availabilities for the other flights as well. In addition, we do allow individual components, say $Q_t^i(x)$ and $Q_t^j(x)$, to be dependent—this allows us to model the interactions of availability and demand across flights. For a generic function g, we shall use the notation $\mathrm{E}[g(Q_t(x))]$ for $\int_q g(q)\, dF_x^t(q) = \sum_q g(q)\mathrm{P}(Q_t(x) = q)$.

When action x is selected in period t and $Q_t = q$, the period-t revenue on flight i is

$$r_t^i(x^i, q^i) = f_t \min\{x^i, q^i\}, \tag{1}$$

where $\min\{x^i, q^i\}$ is the sales for flight i in period t. The total one-period revenue is

$$r_t(x, q) = \sum_{i=1}^{n} r_t^i(x^i, q^i). \tag{2}$$

Likewise, the expected revenue in period t, conditional upon the use of action x in period t, is given by $\mathrm{E}[r_t(x, Q_t(x))]$.

For each i, $Q_t^i(x)$ can be interpreted as the number of sales that would result for flight i if there were infinitely many seats available for flight i, and x^j seats available for each flight $j \neq i$. This notion is somewhat related to so-called *unconstrained demand*, which can be roughly defined as sales that would accrue if there were enough capacity on flight i to satisfy all incoming customers. *Within* a period the dynamics whereby a choice of action x combines with "randomness" to yield a realization of $Q_t(x)$ can be quite complicated—this is the subject of §6.

One might argue that it is more desirable to use the sales in period t as the basic random quantity, rather than to go through $Q_t(x)$, because $Q_t(x)$ may be hard to identify in practice. That is, instead of having a quantity $Q_t(x)$, one could start by defining the single-period revenue to be $\mathrm{E}[\sum_{i=1}^n f_t Y_t^i(x)]$, where the basic random quantity is a sales vector $Y_t(x)$ that should, of course, satisfy $\mathrm{P}(0 \leqslant Y_t(x) \leqslant x) = 1$. However, any such problem can be transformed into our framework, because $Y_t(x) = \min\{x, Y_t(x)\}$ componentwise under the assumption $\mathrm{P}(0 \leqslant Y_t(x) \leqslant x) = 1$. In other words, if one wants to define $Q_t(x)$ to be sales rather than demand, it will be consistent with our mathematical setup. We take $Q_t(x)$ as described above as a starting point for our model, because it gives us a simple correspondence between the model with and without choice behavior. Note that without choice behavior, $Q_t^i(x) = Q_t^i$ is indeed what is typically called the demand for flight i in period t.

A *policy* π prescribes the number of seats to be open on each flight in each period as a function of the booking history up to the time in question. We shall, without loss of optimality, restrict our attention to Markovian policies; i.e., those policies that base decisions upon only the current state (see Puterman 1994). Let $u_t^\pi(s)$ denote the expected total revenue from periods t to 1 under policy π given the state at the beginning of period t is s. Then,

$$u_t^\pi(s) = \mathrm{E}_s^\pi\left[\sum_{k=1}^{t} r_k(X_k, Q_k)\right]. \tag{3}$$

Here E_s^π denotes expectation with respect to the distribution induced by policy π when the state at time t is s. The objective of the MDP is to maximize the expected total revenue; i.e., to maximize $u_m^\pi(0)$. Let $v_t(s)$ be the maximum expected revenue obtainable from periods t to 1 given that the state at the beginning of period t is s; that is,

$$v_t(s) = \max_\pi u_t^\pi(s). \tag{4}$$

Well-known results from MDP theory (see, for instance, Puterman 1994) show that the problem can be reduced to that of iteratively solving the optimality equations

$$v_t(s) = \max_{0 \leqslant x \leqslant c-s} \mathrm{E}[r_t(x, Q_t(x)) + v_{t-1}(\min\{x, Q_t(x)\} + s)], \quad t = 1, \ldots, m, \tag{5}$$

and $v_0(s) = 0$; $s \in \{0, 1, \ldots, c^1\} \times \cdots \times \{0, 1, \ldots, c^n\}$. Before we proceed, observe that if we take $n = 1$ and assume that the distribution of $Q_t(x)$ does not depend upon x, then we get a model like those reviewed in Lautenbacher and Stidham (1999) and Li and Oum (2002).

In principle, one could implement an optimal policy by storing, for each s and t, a maximizing action from (5). Upon entry into a state s at time t, one would then need only look up the appropriate action. For moderately large n, however, the formulation above is rendered intractable by the well-known curse of dimensionality. Consider a case with $n = 10$ flights of 100 seats each and $m = 10$ time periods. For each given time t and state s, we need to solve a potentially difficult integer program. However, even if we were able to compute an optimal policy, storing it in a "look-up table" would require keeping track of $101^{10} \times 10 \approx 10^{21}$ actions. On the other hand, if we have 10 separate flights each with 100 seats and no choice behavior (so there are 10 decoupled problems), then we need only store $10 \times 101 \times 10 \approx 10^4$ actions (without exploiting any of the savings from the existence of structured optimal policies in the decoupled case). So, the inclusion of the customer-choice aspect moves the problem from computationally "easy" to intractable.

A significant portion of the remainder of the paper will be devoted to coming up with good (but suboptimal) policies that can be both computed and stored. Sections 3–6 focus on theoretical aspects of the problem. For instance, in §3, we derive separable upper and lower bounds for the value function. To apply the bounds, we need to find certain bounding sequences for the demands. We show how this can be done for a general choice model in §6. Sections 7–9 describe numerical examples and heuristic solution approaches (some based on ideas from §§3–6). Ultimately, we will focus on two general classes of approaches. In one, we provide methods for determining so-called static booking-limit policies, which require storage of just a few predetermined parameters (the booking limits). In the other, we describe techniques for obtaining easily computable approximations to the value function, which subsequently allow us to compute actions "on the fly"—so, there is no look-up table, but rather the action for a particular state is computed upon entry into the state in question.

3. Upper and Lower Bounds

In this section, we derive upper and lower bounds for the value function (4). To this end, recall that random variable X is stochastically smaller than random variable Y (written $X \leqslant_{st} Y$) means $\mathrm{E}g(X) \leqslant \mathrm{E}g(Y)$ for all increasing functions g for which the expectations exist. Equivalently, $\mathrm{P}(X \leqslant x) \geqslant \mathrm{P}(Y \leqslant x)$ for all x (see, e.g., Müller and Stoyan 2002). Let integer-valued random vectors $\{\bar{Q}_t\colon t = 1, \ldots, m\}$ (respectively, $\{\underline{Q}_t\colon t = 1, \ldots, m\}$) be such that $(\bar{Q}_1, \ldots, \bar{Q}_m)$ (respectively, $(\underline{Q}_1, \ldots, \underline{Q}_m)$) are independent. Moreover, suppose that

$$\underline{Q}_t^i \leqslant_{st} Q_t^i(x) \leqslant_{st} \bar{Q}_t^i \quad \text{for all } x \in \{0, 1, \ldots, c^1\} \times \cdots \times \{0, 1, \ldots, c^n\},\ \ i \in N,\ t = 1, \ldots, m.$$

Note that the distributions of $\bar{Q}_t$ and $\underline{Q}_t$ do not depend upon the action x. We will consider two separate MDPs for each flight using demand sequences $\{\bar{Q}_t\colon t = 1, \ldots, m\}$ and $\{\underline{Q}_t\colon t = 1, \ldots, m\}$. We follow the time convention and fare structure as in §2. In subsequent sections, we will describe methods to construct $\{\bar{Q}_t\colon t = 1, \ldots, m\}$ and $\{\underline{Q}_t\colon t = 1, \ldots, m\}$ for particular choice models.

Let $\bar{v}_t^i(s^i)$ be the maximum expected revenue obtainable from periods t to 1 on flight i where the demand in period k is $\bar{Q}_k^i$ for $k = 1, \ldots, t$ and the number of seats sold before time t is s^i. The MDP optimality equation for flight i is

$$\bar{v}_t^i(s^i) = \max_{0 \leqslant x^i \leqslant c^i - s^i} \mathrm{E}\big[r_t^i(x^i, \bar{Q}_t^i) + \bar{v}_{t-1}^i(\min\{x^i, \bar{Q}_t^i\} + s^i)\big], \quad t = 1, \ldots, m, \quad (6)$$

and $\bar{v}_0^i(s^i) = 0$; $s^i \in \{0, 1, \ldots, c^i\}$. Similarly, let $\underline{v}_t^i(s^i)$ be the maximum expected revenue obtainable from periods t to 1 on flight i where the demand in period k is $\underline{Q}_k^i$ for $k = 1, \ldots, t$ and the number of seats sold before time t is s^i. The MDP optimality equation for flight i is

$$\underline{v}_t^i(s^i) = \max_{0 \leqslant x^i \leqslant c^i - s^i} \mathrm{E}\big[r_t^i\big(x^i, \underline{Q}_t^i\big) + \underline{v}_{t-1}^i\big(\min\big\{x^i, \underline{Q}_t^i\big\} + s^i\big)\big], \quad t = 1, \ldots, m, \quad (7)$$

and $\underline{v}_0^i(s^i) = 0$; $s^i \in \{0, 1, \ldots, c^i\}$.

We will make use of the following result, which is a special case of Proposition 4 in Cooper and Gupta (2005), several times throughout the remainder of the paper.

Proposition 1. *Consider a single-flight MDP with no choice behavior, (one-dimensional) demand sequence $\{D_t\colon t = 1, \ldots, m\}$, and seat capacity κ. Let $w_t(\cdot)$ be the associated one-dimensional value function. That is, $w_t(\cdot)$ satisfies $w_t(s) = \max_{0 \leqslant x \leqslant \kappa - s} \mathrm{E}[f_t \min\{x, D_t\} + w_{t-1}(\min\{x, D_t\} + s)]$. Similarly, consider a single-flight MDP with no choice behavior, demand sequence $\{D_t'\colon t = 1, \ldots, m\}$, and capacity κ, with associated value function $w_t'(\cdot)$. Assume that demand distributions do not depend upon the actions selected. If $D_t \leqslant_{st} D_t'$ for all t, then $w_t(s) \leqslant w_t'(s)$ for all t and s.*

The following proposition is the main result of this section. As described in the previous section, the upper and lower bounds given below are simple to compute.

Proposition 2. *$\sum_{i=1}^n \underline{v}_t^i(s^i) \leqslant v_t(s) \leqslant \sum_{i=1}^n \bar{v}_t^i(s^i)$ for $t = 1, \ldots, m$.*

Proof. The proof is by induction. For $t = 1$, we have

$$v_1(s) = \max_{0 \leqslant x \leqslant c - s} \mathrm{E} \sum_{i=1}^n f_1 \min\{x^i, Q_1^i(x)\},$$

and

$$\sum_{i=1}^n \bar{v}_1^i(s^i) = \sum_{i=1}^n \max_{0 \leqslant x^i \leqslant c^i - s^i} \mathrm{E} f_1 \min\{x^i, \bar{Q}_1^i\} = \max_{0 \leqslant x \leqslant c - s} \mathrm{E} \sum_{i=1}^n f_1 \min\{x^i, \bar{Q}_1^i\}, \quad (8)$$

$$\sum_{i=1}^{n} \underline{v}_1^i(s^i) = \sum_{i=1}^{n} \max_{0 \leqslant x^i \leqslant c^i - s^i} \mathrm{E} f_1 \min\{x^i, \underline{Q}_1^i\}$$
$$= \max_{0 \leqslant x \leqslant c-s} \mathrm{E} \sum_{i=1}^{n} f_1 \min\{x^i, \underline{Q}_1^i\}. \tag{9}$$

The final equalities in (8) and (9) hold because $\overline{Q}_1$ and $\underline{Q}_1$ do not depend on x. Because $\underline{Q}_1^i \leqslant_{st} Q_1^i(x) \leqslant_{st} \overline{Q}_1^i$ and $f_1 \min\{x^i, q_1^i\}$ is increasing in q_1^i for all x^i, it follows that $\sum_{i=1}^{n} \underline{v}_1^i(s^i) \leqslant v_1(s) \leqslant \sum_{i=1}^{n} \bar{v}_1^i(s^i)$, thereby completing the base case.

For the inductive step, we assume for $t \geqslant 2$ that

$$\sum_{i=1}^{n} \underline{v}_{t-1}^i(s^i) \leqslant v_{t-1}(s) \leqslant \sum_{i=1}^{n} \bar{v}_{t-1}^i(s^i)$$

for all s. Fix s and let

$$\tilde{x} \in \operatorname*{argmax}_{0 \leqslant x \leqslant c-s} \mathrm{E} \sum_{i=1}^{n} \left[r_t^i(x^i, Q_t^i(x)) + \bar{v}_{t-1}^i(\min\{x^i, Q_t^i(x)\} + s^i) \right].$$

Then, using the induction hypothesis, we have

$$\begin{aligned} v_t(s) &= \max_{0 \leqslant x \leqslant c-s} \mathrm{E}\left[\sum_{i=1}^{n} r_t^i(x^i, Q_t^i(x)) + v_{t-1}(\min\{x, Q_t(x)\} + s) \right] \\ &\leqslant \max_{0 \leqslant x \leqslant c-s} \mathrm{E}\left[\sum_{i=1}^{n} r_t^i(x^i, Q_t^i(x)) + \sum_{i=1}^{n} \bar{v}_{t-1}^i(\min\{x^i, Q_t^i(x)\} + s^i) \right] \\ &= \sum_{i=1}^{n} \mathrm{E}[r_t^i(\tilde{x}^i, Q_t^i(\tilde{x})) + \bar{v}_{t-1}^i(\min\{\tilde{x}^i, Q_t^i(\tilde{x})\} + s^i)] \\ &\leqslant \sum_{i=1}^{n} \max_{0 \leqslant y^i \leqslant c^i - s^i} \mathrm{E}\left[r_t^i(y^i, Q_t^i(\tilde{x})) + \bar{v}_{t-1}^i(\min\{y^i, Q_t^i(\tilde{x})\} + s^i) \right]. \end{aligned} \tag{10}$$

Now consider a t-period one-dimensional MDP without choice behavior (i.e., with demand independent of action) that has demand vector $(Q_t^i(\tilde{x}), \overline{Q}_{t-1}^i, \ldots, \overline{Q}_1^i)$. Let $\tilde{v}_t^i(s^i)$ denote the value function of this MDP. Observe that the ith term in the summation (10) is precisely $\tilde{v}_t^i(s^i)$. Proposition 1 now implies that the ith term is bounded above by $\bar{v}_t^i(s^i)$. Hence, $v_t(s) \leqslant \sum_{i=1}^{n} \bar{v}_t^i(s^i)$. This completes the proof of the upper bound.

Next, we prove the lower bound. Let

$$\underline{x}^i = \min\left\{0 \leqslant k \leqslant c^i - s^i \colon f_t < \underline{v}_{t-1}^i(s^i + k) - \underline{v}_{t-1}^i(s^i + k + 1)\right\}. \tag{11}$$

By Theorem 4 and Corollary 2 of Lautenbacher and Stidham (1999),

$$\underline{x}^i \in \operatorname*{argmax}_{0 \leqslant x^i \leqslant c^i - s^i} \mathrm{E}\left[r_t^i(x^i, \underline{Q}_t^i) + \underline{v}_{t-1}^i(\min\{x^i, \underline{Q}_t^i\} + s^i) \right], \quad i = 1, \ldots, n.$$

Note that from (11), the value of $\underline{x}^i$ is determined by $\underline{v}_{t-1}^i(\cdot)$, the value of which does not depend on $\underline{Q}_t^i$. Hence, for each i, $\underline{x}^i$ depends upon the distributions of $(\underline{Q}_{t-1}^i, \ldots, \underline{Q}_1^i)$, but does *not* depend upon the distribution of $\underline{Q}_t^i$. Consequently, $\underline{x}^i$ also maximizes $\mathrm{E}[r_t^i(y^i, Q_t^i(\underline{x})) + \underline{v}_{t-1}^i(\min\{y^i, Q_t^i(\underline{x})\} + s^i)]$ over $y^i \in \{0, 1, \ldots, c^i - s^i\}$. Let $\hat{v}_t^i(s^i)$ denote the value function of a t-period one-dimensional MDP with no choice behavior that has demand vector $(Q_t^i(\underline{x}), \underline{Q}_{t-1}^i, \ldots, \underline{Q}_1^i)$. We now have

$$\begin{aligned} v_t(s) &= \max_{0 \leqslant x \leqslant c-s} \mathrm{E}\left[\sum_{i=1}^{n} r_t^i(x^i, Q_t^i(x)) + v_{t-1}(\min\{x, Q_t(x)\} + s) \right] \\ &\geqslant \max_{0 \leqslant x \leqslant c-s} \mathrm{E}\left[\sum_{i=1}^{n} r_t^i(x^i, Q_t^i(x)) + \sum_{i=1}^{n} \underline{v}_{t-1}^i(\min\{x^i, Q_t^i(x)\} + s^i) \right] \\ &\geqslant \sum_{i=1}^{n} \mathrm{E}\left[r_t^i(\underline{x}^i, Q_t^i(\underline{x})) + \underline{v}_{t-1}^i(\min\{\underline{x}^i, Q_t^i(\underline{x})\} + s^i) \right] \\ &= \sum_{i=1}^{n} \hat{v}_t^i(s^i) \geqslant \sum_{i=1}^{n} \underline{v}_t^i(s^i). \end{aligned}$$

In the above, the first inequality follows from inductive hypothesis, and the third inequality follows from Proposition 1. This completes the proof. $\square$

It is also possible to obtain other bounds for v_t. In §5, we describe an upper bound based upon inventory pooling. For another lower bound, define

$$W_t^x(s) = \mathrm{E}\left[r_t(a_{x,s}, Q_t(a_{x,s})) + W_{t-1}^x(\min\{a_{x,s}, Q_t(a_{x,s})\} + s) \right], \quad t = 1, \ldots, m,$$

where $a_{x,s} = \min\{x, c-s\}$ and $W_0^x(s) = 0$. Let

$$W_t(s) = \max_{0 \leqslant x \leqslant c-s} \mathrm{E}\left[r_t(x, Q_t(x)) + W_{t-1}^x(\min\{x, Q_t(x)\} + s) \right]. \tag{12}$$

It is straightforward to verify that $W_t(s) \leqslant v_t(s)$. Intuitively, $W_t(s)$ represents the situation where seat availability remains unchanged (with the exception that we do not oversell capacity) in periods $t, \ldots, 1$. In our experience with small numerical examples, (12) can be either larger or smaller than the lower bound in Proposition 2. However, the bounds in Proposition 2 are much easier to compute. Indeed, the maximization in (12) becomes intractable for large s, rendering (12) useless in such cases. In fact, exact evaluation of $W_t(s)$ requires roughly the same amount of computational effort as solving the multidimensional MDP exactly.

4. Booking Limits: Static Versus Dynamic

As described by Lautenbacher and Stidham (1999) and others, there is an optimal policy for models without customer choice with an appealing simple form. For our purposes, these results imply that an optimal policy $\underline{\pi}^i$ for flight i of the lower-bound problem (given by the maximizers in (7)) can be determined by a sequence $(\underline{b}_1^i, \underline{b}_2^i, \ldots, \underline{b}_m^i)$ as follows:

$$\underline{\pi}_t^i(s^i) = (\underline{b}_t^i - s^i)^+. \tag{13}$$

That is, if the state is s^i at time t for the one-dimensional problem with demand process $\{\underline{Q}_t^i\}$, it is optimal to make $(\underline{b}_t^i - s^i)^+$ seats available. See, e.g., Lautenbacher and Stidham (1999) for more on how to compute $(\underline{b}_1^i, \underline{b}_2^i, \ldots, \underline{b}_m^i)$. The existence of optimal policies of this type for the standard no-choice single-leg model allows one to store an optimal policy in terms of just m parameters.

In our formulation (4)–(5), we consider *only* booking-limit policies; implementing action x in state s gives a vector of booking limits of $s+x$. We can divide the class of booking-limit policies into what we shall term *static booking-limit policies* and *dynamic booking-limit policies*. A static booking-limit policy π^b is a policy of the form

$$(\pi^b)_t^i(s) = (b_t^i - s^i)^+, \tag{14}$$

where b is a matrix whose (i,t)th element is the booking limit for flight i in period t. Expression (14) means that at time t and state s, the policy π^b makes $(b_t^i - s^i)^+$ seats available for flight i; $i = 1, \ldots, n$. The modifier "static" captures the fact that the matrix b is fixed (and determined ahead of time). Policies that are not of this type we shall term dynamic booking-limit policies. Using this terminology, we can restate the result above as follows: Single-leg problems have an optimal static booking-limit policy. Unfortunately, this is not the case when we have choice behavior, as the following example shows.

Example 1. Consider a deterministic problem with two flights, each of which has two seats. There are more than two periods. In Period 2, there is one customer, who is willing to purchase a seat only on Flight 2. In Period 1, there are two customers, each of whom strictly prefers seats on Flight 1, but is willing to accept seats on Flight 2. Using the terminology of §2, we have

$$Q_2(x) = (0,1),$$

$$Q_1(x) = \begin{cases} (2,0) & \text{if } x^1 \geqslant 2, \\ (2,1) & \text{if } x^1 = 1, \\ (2,2) & \text{if } x^1 = 0. \end{cases}$$

Let the fare in Period 2 be $f_2 = 50$, and the fare in Period 1 be $f_1 = 100$. At the start of Period 2, if there are, respectively, one and two remaining seats on Flights 1 and 2 (i.e., if the state is $(1,0)$), then the optimal action in Period 2 is $x = (0,1)$. That is, only Flight 2 will be open and we will accept at most one booking. If instead the number of remaining seats is $(0,2)$ (i.e., if the state is $(2,0)$), then the optimal action is $(0,0)$—i.e., both flights should be closed. Hence, the optimal policy is not a static booking-limit policy.

The following example shows that the value function is not componentwise concave (for one-dimensional problems without choice behavior, the value function is concave).

Example 2. Consider the last period of a deterministic problem with two flights each with seat capacity 3. Suppose that the fare is $f_1 = 100$. There are two customer arrivals in the period. The first customer prefers Flight 1, but is willing to accept seats on Flight 2. The second customer is only willing to accept seats on Flight 1. The order of arrival matters. If the number of remaining seats is $(0,1)$ (i.e., if the state is $(3,2)$), then the total number of sales is 1, so $v_1(3,2) = 100$. If the number of remaining seats is $(1,1)$ (i.e., if the state is $(2,2)$), then the total number of sales is 1 and $v_1(2,2) = 100$. If the number of remaining seats is $(2,1)$ (i.e., if the state is $(1,2)$), then the total number of sales is 2 and $v_1(1,2) = 200$. Therefore, the value function is not componentwise concave. Because $v_1(0,2) = 200$, it is also not componentwise convex.

Given the simplicity of static booking-limit policies, it is of general interest to ask whether or not there are good policies within this narrow class. The following proposition says that the expected revenue from the simple static booking-limit policy (13) described above is at least as good as the optimal expected revenue for the lower-bound problem. Let $u_t^{\underline{\pi}}(s)$ be the expected revenue from using the booking limits from the lower-bound problem as an operating policy for periods $t, t-1, \ldots, 1$ in the actual problem; i.e., $\underline{\pi}_t(s) = (\underline{\pi}_t^1(s^1), \ldots, \underline{\pi}_t^n(s^n))$—see (3), (4), (7), and (13). Let $v_t^b(s) = \sup_b u_t^{\pi^b}(s)$ be the maximum expected revenue possible from a static booking-limit policy.

Proposition 3. $\sum_{i=1}^n \underline{v}_t^i(s^i) \leqslant u_t^{\underline{\pi}}(s) \leqslant v_t^b(s) \leqslant v_t(s)$.

Proof. The third inequality follows immediately from (4). The second inequality follows immediately from the definition of $v_t^b(s)$. Therefore, it remains only to prove the first inequality. This is done by induction. For $t = 1$,

$$\begin{aligned} u_1^{\underline{\pi}}(s) &= \mathrm{E}\sum_{i=1}^n f_1 \min\{c^i - s^i, Q_1^i(c-s)\} \\ &\geqslant \mathrm{E}\sum_{i=1}^n f_1 \min\{c^i - s^i, \underline{Q}_1^i\} = \sum_{i=1}^n \underline{v}_1^i(s^i). \end{aligned}$$

For the inductive step, assume that

$$u_{t-1}^{\underline{\pi}}(s) \geqslant \sum_{i=1}^n \underline{v}_{t-1}^i(s^i).$$

Let $\hat{v}_t^i(s^i)$ be the value function of t-period one-dimensional MDP with no choice behavior that has demand vector $(Q_t^i((\underline{b}_t-s)^+), \underline{Q}_{t-1}^i, \dots, \underline{Q}_1^i)$. Then,

$$\begin{aligned}u_t^{\pi}(s) &= \mathrm{E}\bigg[\sum_{i=1}^n f_t \min\{(\underline{b}_t^i-s^i)^+, Q_t^i((\underline{b}_t-s)^+)\}\\ &\qquad + u_{t-1}^{\pi}(\min\{(\underline{b}_t-s)^+, Q_t((\underline{b}_t-s)^+)\}+s)\bigg]\\ &\geqslant \sum_{i=1}^n \mathrm{E}[f_t\min\{(\underline{b}_t^i-s^i)^+, Q_t^i((\underline{b}_t-s)^+)\}\\ &\qquad + \underline{v}_{t-1}^i(\min\{(\underline{b}_t-s)^+, Q_t((\underline{b}_t-s)^+)\}+s)]\\ &= \sum_{i=1}^n \hat{v}_t^i(s^i) \geqslant \sum_{i=1}^n \underline{v}_t^i(s^i).\end{aligned}$$

In the above, the first inequality follows from the induction hypothesis, and the second inequality follows from Proposition 1. $\square$

5. Inventory Pooling

In this section, we describe how to obtain a different upper bound based upon inventory pooling. Let $Y_t^i(x) \equiv \min\{x^i, Q_t^i(x)\}$ be the random sales made for the ith flight in period t in the original problem with seat allocation x, and define $Y_t^{\mathrm{p}}(x) \equiv \sum_{i=1}^n Y_t^i(x)$. Observe that $Y_t^{\mathrm{p}}(x) \leqslant \min\{\sum_{i=1}^n x^i, \sum_{i=1}^n Q^i(x)\} \leqslant \sum_{i=1}^n x^i$. Hence, for each x, we have

$$Y_t^{\mathrm{p}}(x) = \min\left\{\sum_{i=1}^n x^i, Y_t^{\mathrm{p}}(x)\right\}. \tag{15}$$

Let $c^{\mathrm{p}} = \sum_{i=1}^n c^i$ be the total capacity of the n flights, and suppose that there exists an integer-valued random variable D_t that satisfies $Y_t^{\mathrm{p}}(x) \leqslant_{st} D_t$ for all x. The random variable D_t can be interpreted as the number of customers that are interested in buying a ticket in period t. In the next section, we provide a construction consistent with this interpretation. Let $v_t^{\mathrm{p}}(\cdot)$ be the MDP value function of the *pooled problem*, in which we have one flight with capacity c^{p} and independent one-dimensional demand sequence $\{D_t\}$. That is, $v_t^{\mathrm{p}}(\cdot)$ satisfies $v_t^{\mathrm{p}}(k) = \max_{0\leqslant z\leqslant c^{\mathrm{p}}-k} \mathrm{E}[f_t\min\{z, D_t\} + v_{t-1}^{\mathrm{p}}(\min\{z, D_t\} + k)]$, where k and z are integer scalars.

Proposition 4. *For each t, we have $v_t^{\mathrm{p}}(\sum_{i=1}^n s^i) \geqslant v_t(s)$ for all s.*

Proof. The proof is by induction. The statement holds for $t=1$. For the inductive step, assume that the result is true for $t-1$; that is, $v_{t-1}^{\mathrm{p}}(\sum_{i=1}^n s^i) \geqslant v_{t-1}(s)$ for all $s = (s^1, \dots, s^n)$. Next, consider an arbitrary state vector $s = (s^1, \dots, s^n)$ and let $s^{\mathrm{p}} = \sum_{i=1}^n s^i$. By the induction hypothesis,

$$\begin{aligned}v_t(s) &= \max_{0\leqslant x\leqslant c-s} \mathrm{E}\left[\sum_{i=1}^n f_t Y_t^i(x) + v_{t-1}(Y_t(x)+s)\right]\\ &\leqslant \max_{0\leqslant x\leqslant c-s} \mathrm{E}[f_t Y_t^{\mathrm{p}}(x) + v_{t-1}^{\mathrm{p}}(Y_t^{\mathrm{p}}(x)+s^{\mathrm{p}})] \qquad (16)\\ &= \mathrm{E}[f_t Y_t^{\mathrm{p}}(\tilde{x}) + v_{t-1}^{\mathrm{p}}(Y_t^{\mathrm{p}}(\tilde{x})+s^{\mathrm{p}})],\end{aligned}$$

where $\tilde{x}$ is the maximizer in (16). Relation (15) now implies

$$\begin{aligned}v_t(s) \leqslant \mathrm{E}\bigg[&f_t\min\left\{\sum_{i=1}^n \tilde{x}^i, Y_t^{\mathrm{p}}(\tilde{x})\right\}\\ &+ v_{t-1}^{\mathrm{p}}\left(\min\left\{\sum_{i=1}^n \tilde{x}^i, Y_t^{\mathrm{p}}(\tilde{x})\right\} + s^{\mathrm{p}}\right)\bigg].\end{aligned}$$

Because $0 \leqslant \tilde{x}^i \leqslant c^i - s^i$ for $i = 1, \dots, n$, it follows that $0 \leqslant \sum_{i=1}^n \tilde{x}^i \leqslant c^{\mathrm{p}} - s^{\mathrm{p}}$. Hence, from the above we have

$$\begin{aligned}v_t(s) &\leqslant \max_{0\leqslant z\leqslant c^{\mathrm{p}}-s^{\mathrm{p}}} \mathrm{E}\big[f_t\min\{z, Y_t^{\mathrm{p}}(\tilde{x})\}\\ &\qquad + v_{t-1}^{\mathrm{p}}(\min\{z, Y_t^{\mathrm{p}}(\tilde{x})\} + s^{\mathrm{p}})\big] \qquad (17)\\ &\leqslant \max_{0\leqslant z\leqslant c^{\mathrm{p}}-s^{\mathrm{p}}} \mathrm{E}\big[f_t\min\{z, D_t\}\\ &\qquad + v_{t-1}^{\mathrm{p}}(\min\{z, D_t\} + s^{\mathrm{p}})\big] \qquad (18)\\ &= v_t^{\mathrm{p}}(s^{\mathrm{p}}),\end{aligned}$$

where Inequality (18) follows from Proposition 1. Observe that the maximizations in (17) and (18) are over one-dimensional sets. The final equality above is simply the optimality equation for the pooled problem. $\square$

Before we proceed, it is tempting to consider the control strategy whereby total availability for all flights combined is given by the maximizer in (18). However, this is not allowed in formulation (5), because (5) assumes that the distribution of $Q_t(x)$ is induced by an initial availability x. In fact, the formulation can be extended in some cases (e.g., when the choice model in §6 is in place) to allow strategies of this form. However, in several of our test problems, the method implied by the maximization in (18) does not perform as well as other heuristics (please refer to §9, where the "pooled inventory" policy is called PBL, for details).

6. A Choice Model

In this section, we describe a customer-choice model, and we show how to derive from it upper and lower bounds on the distribution of $Q_t(x)$ as needed to apply the results of §3. We also identify D_t as needed for Proposition 4. The developments also show us how to, in principle, derive the distributions $F_x^t(\cdot)$ of $Q_t(x)$ from other more basic quantities that reflect consumers' underlying preferences. In addition, we provide a formal description of "demand" when consumer choice is affected by availability.

6.1. The Model

We study the customer-choice process in a generic booking period, so we suppress the subscript t in this section. We will focus on how the choice of initial availability vector x (recall x is an action) affects customer behavior.

Assume that customers arrive randomly throughout the period. Let $N_0 = \{0, 1, \dots, n\}$ and $M = \{1, 2, \dots, n+1\}$. The elements of N_0 represent the different flights, and 0 denotes the no-purchase option. A one-to-one mapping θ: $N_0 \to M$ for which $\theta(0) \neq 1$ is called a preference mapping. Each arriving customer is assumed to have such a mapping, which he uses to decide what (if anything) to purchase. The preference mapping of a customer is determined *before* he observes the inventory availability. For $k, l \in N_0$, we say that flight k is preferred to l if $\theta(k) < \theta(l)$. The requirement $\theta(0) \neq 1$ means that no customer's most preferred option is not to purchase a ticket. The requirement that a preference mapping is one to one ensures that each customer has strict preference order among the flights (and the no-purchase option). This model is essentially equivalent to ones that use random utility maximization as a starting point (see, e.g., Mahajan and van Ryzin 2001b, §2.2.1).

Let D be an almost surely finite nonnegative integer-valued random variable that represents the total number of customer arrivals in period t. For each k, let Θ_k denote the random preference mapping of the kth arriving customer. The basic random quantity is $(D, \{\Theta_k: k \in \mathbb{Z}^+\})$. We assume that the distribution of $(D, \{\Theta_k: k \in \mathbb{Z}^+\})$ is not affected by the inventory policy or the initial condition at the beginning of period t. The distributions of these vectors in *different* time periods are independent (this is needed for us to have an MDP). However, for a fixed period t, we allow a completely arbitrary dependence structure in the joint distribution of $(D, \{\Theta_k: k \in \mathbb{Z}^+\})$.

For availability n-vector y, define $A(y) = \{i: y^i > 0\} \cup \{0\}$ to be the set of available flights and the no-purchase option when the availability vector is y. A customer who has preference mapping θ and who arrives to find availability vector y will make the choice

$$\phi(y, \theta) = \operatorname*{argmin}_{i \in A(y)} \theta(i). \tag{19}$$

Note that $\phi(y, \theta)$ depends on y only through $A(y)$. The choice of customer k (i.e., the flight on which k buys a ticket) is $\Phi_{k,x} = \phi(X_x(k), \Theta_k)$, where $X_x(k)$ is the availability vector faced by customer k when the initial inventory allocation is x. The sequence $\{X_x(k)\}$ satisfies the recursion

$$X_x(k+1) = X_x(k) - \epsilon^{\Phi_{k,x}}, \quad k = 1, \ldots, D, \tag{20}$$

with the boundary condition $X_x(1) = x$. The total sales on flight i can be expressed as

$$Y^i(x) = \sum_{k=1}^{D} \mathbb{I}\{\Phi_{k,x} = i\}. \tag{21}$$

The choice model complies with the basic *consumer sovereignty property* (see McFadden 2000) in economics, which states that preferences are predetermined in any choice situation, and do not depend upon the alternatives available for selection. As mentioned in the introduction, the setup above is equivalent to that described by Mahajan and van Ryzin (2001a, b). As shown by Mahajan and van Ryzin, it subsumes a variety of models in the literature.

However, there are reasonable and realistic choice models that are not covered by the above framework. For example, psychologists have shown that when facing too many options, people may defer their decision or search for new options (see, e.g., Iyengar and Lepper 2000). In one of their studies (involving jars of jam), nearly 30% of customers who are provided with a limited number of choices subsequently make a purchase, whereas only 3% of customers who are given an extensive number of options do so. In such a case, a customer's preference may not be established before he observes the available choices. It should be pointed out that the results and the solution methodology presented in this paper do not depend upon a particular choice model; the setup described in this section is just one case where our results apply.

6.2. What Is Demand?

Revenue managers often use the term "unconstrained demand" to, in some sense, mean pure demand. Owing to the interaction among multiple flights, it is not altogether clear what these terms mean in situations when choice behavior is present. In this section, we provide one definition that is consistent with our earlier developments. In the following, the demand distribution on flight i does not depend on the seat availability on flight i, and depends only on the seat availability on all the other flights. So, what is demand when there is choice behavior?

Define

$$Q^i(x) = \sum_{k=1}^{D} \mathbb{I}\{\phi(X_{x+\infty\epsilon^i}(k), \Theta_k) = i\}. \tag{22}$$

In the above, $x + \infty\epsilon^i$ is the vector x with the ith element x^i replaced by ∞. We call $Q^i(x)$ the demand for flight i. By (21), $Q^i(x) = Y^i(x + \infty\epsilon^i)$. In other words, $Q^i(x)$ is the sales on flight i if the number of seats allocated on flight i is infinite and the seat allocation is x^j on each flight $j \neq i$. Observe that $Q^i(x)$ does not depend upon x^i, but it does, however, depend upon the other entries of x via $X^j_{x+\infty\epsilon^i}(1) = x^j$; $j \neq i$. Expression (22) gives us a pathwise definition, where the random variables $Q^i(x)$ are defined on the same probability space as the process $(D, \{\Theta_k: k \in \mathbb{Z}^+\})$. This allows computation of the distributions of $Q(x)$ through the formula

$$F^t_x(q) = \mathrm{P}\left(\sum_{k=1}^{D} \mathbb{I}\{\phi(X_{x+\infty\epsilon^i}(k), \Theta_k) = i\} \leqslant q^i;\, i = 1, \ldots, n\right). \tag{23}$$

The extent to which this computation is simple or not depends upon the probabilistic assumptions we place upon the underlying space. For more on this, see Mahajan and van Ryzin (2001a, b), who describe how many tractable models can indeed be put into this framework.

We next demonstrate that $Q^i(x)$ as defined in (22) is consistent with our earlier use of the symbol in §§2–5.

Specifically, we show that sales can be expressed as the minimum of demand and the number of open seats as in expressions (1)–(2).

LEMMA 1. $Y^i(x)=\min\{x^i, Q^i(x)\}$, *where* $Y^i(x)$ *is defined by* (21) *and* $Q^i(x)$ *is defined by* (22).

PROOF. By (22), we have

$$\begin{aligned} Q^i(x) &= \sum_{k=1}^{D} \mathbb{I}\{\phi(X_{x+\infty\epsilon^i}(k),\Theta_k)=i\} \\ &= \sum_{k=1}^{D} \big[\mathbb{I}\{\phi(X_{x+\infty\epsilon^i}(k),\Theta_k)=i\}\mathbb{I}\{X_x^i(k)>0\} \\ &\qquad + \mathbb{I}\{\phi(X_{x+\infty\epsilon^i}(k),\Theta_k)=i\}\mathbb{I}\{X_x^i(k)=0\}\big] \\ &= Y^i(x) + \sum_{k=1}^{D} \mathbb{I}\big\{\phi(X_{x+\infty\epsilon^i}(k),\Theta_k)=i\big\}\mathbb{I}\big\{X_x^i(k)=0\big\}. \end{aligned} \tag{24}$$

In the above, the last equality follows because

$$\phi(X_{x+\infty\epsilon^i}(k),\Theta_k)=\phi(X_x(k),\Theta_k)$$

when $X_x^i(k)>0$. Note also that

$$\begin{aligned} Y^i(x) &= \sum_{k=1}^{D} \mathbb{I}\{\phi(X_x(k),\Theta_k)=i\} \\ &= \sum_{k=1}^{D} \mathbb{I}\{\phi(X_x(k),\Theta_k)=i\}\mathbb{I}\{X_x^i(k)>0\}. \end{aligned}$$

In (24), observe that if $Y^i(x)<x^i$, then $X_x^i(k)>0$ for all k. Consequently, $Y^i(x)=Q^i(x)<x^i$. Similarly, if $Y^i(x)=x^i$, then $Q^i(x)\geqslant Y^i(x)=x^i$. This completes the proof. □

Before we proceed, it is worth pointing out there certainly could be other reasonable definitions of demand. From the standpoint of the models described in this paper, the key requirement is that (1)–(2) make sense in conjunction with the definition. This also points out the generality of the framework developed in §3, because the results there do not depend upon the particular details that give rise to $Q(x)$.

We are now ready to derive upper and lower bounds on demand as needed to apply the results of §3 to the choice model described in §6.1. Let

$$\underline{D}^i=\sum_{k=1}^{D}\mathbb{I}\{\Theta_k(i)=1\} \quad \text{and} \tag{25}$$

$$\overline{D}^i=Q^i(0). \tag{26}$$

Observe that $D=\sum_{i=1}^{n}\underline{D}^i$. We now have the following pathwise inequalities.

LEMMA 2. $\underline{D}^i\leqslant Q^i(x)\leqslant\overline{D}^i$ *for all* x.

PROOF. First note that $\overline{D}^i=Q^i(0)=\sum_{k=1}^{D}\mathbb{I}\{\phi(0+\infty\epsilon^i,\Theta_k)=i\}=\sum_{k=1}^{D}\mathbb{I}\{\Theta_k(i)<\Theta_k(0)\}$. The result now follows immediately from the fact that $\mathbb{I}\{\Theta_k(i)=1\}\leqslant\mathbb{I}\{\phi(X_{x+\infty\epsilon^i},\Theta_k)=i\}\leqslant\mathbb{I}\{\Theta_k(i)<\Theta_k(0)\}$. □

In the above, note that both $\underline{D}^i$ and $\overline{D}^i$ are random variables whose distribution is determined by exogenous choice process characteristics. We can now apply the result of §3 to obtain bounds for the value of the MDP for the choice models of this section. To this end, let $\{\underline{D}_t\colon t=1,\ldots,m\}$ and $\{\overline{D}_t\colon t=1,\ldots,m\}$ be the lower and upper bounds for demand in periods $t=1,\ldots,m$. Clearly, these sequences satisfy the conditions needed for $\{\underline{Q}_t\}$ and $\{\overline{Q}_t\}$ in Proposition 2 (see, e.g., Theorem 1.2.4 of Müller and Stoyan 2002), thereby yielding the following result.

PROPOSITION 5. *For each* $i=1,\ldots,n$, *construct an MDP for flight i with demand sequence* $\{\underline{D}_t^i\colon t=1,\ldots,m\}$ (*respectively,* $\{\overline{D}_t^i\colon t=1,\ldots,m\}$), *and denote the value function* $\underline{v}_t^i(s^i)$ (*respectively,* $\bar{v}_t^i(s^i)$). *Then,* $\sum_{i=1}^{n}\underline{v}_t^i(s^i)\leqslant v_t(s)\leqslant\sum_{i=1}^{n}\bar{v}_t^i(s^i)$ *for* $t=1,\ldots,m$.

It is also evident that $Q^i(x)$, $Y^i(x)$; $i=1,\ldots,n$ and D described in §§6.1 and 6.2 satisfy the conditions needed to apply Proposition 4. Therefore, the pooled upper bound is applicable to the choice model of §6.

The relationship between v_t^{p} in Proposition 4 and the upper bound in Proposition 5 depends upon the particulars of a given problem. To see this, note that the bound from the pooled problem is tight when there is full substitution (i.e., when $\Theta_k(0)=n+1$ for all k)—provided the original formulation is extended to allow such pooling policies and the definition of the "set of available flights" in §6.1 is suitably modified. The upper bound in Proposition 5 is tight when there is no substitution ($\Theta_k(0)=2$ for all k).

7. Solution Approaches

Because the state and action spaces are very large, exact solution of the MDP in §2 is, for practical purposes, impossible for moderate-sized problems. Consequently, we seek to identify heuristic methods that perform reasonably well. One simple heuristic method for our problem is to use the policy derived from the lower-bound booking limits. We showed in Proposition 3 that the expected revenue from following this policy is at least as large as the lower bound for the optimal expected revenue $\sum_{i=1}^{n}\underline{v}_m^i(0)$. In numerical experiments, we have observed that the lower-bound booking-limit policy outperforms other simple booking-limit policies (upper-bound booking-limit policy, hybrid booking-limit policy from lower- and upper-bound booking limits). So, we will use the lower-bound booking-limit policy as a baseline in our numerical example section.

7.1. Approximation Using Upper and Lower Bounds

In this section, we discuss how to approximate the value function using a weighted average of the upper and lower bounds. The approximation scheme takes the following form:

$$\tilde{v}_t(s) = \sum_{i=1}^{n} [\beta_i(s,t)\bar{v}_t^i(s^i) + (1-\beta_i(s,t))\underline{v}_t^i(s^i)]. \tag{27}$$

Alternatively, we could use the upper bound from the pooled problem instead of the separable upper bound. We will only discuss in detail approximation (27).

For $j = 1, \ldots, n$, we have

$$\begin{aligned}\Delta_j \tilde{v}_t(s) &= \tilde{v}_t(s) - \tilde{v}_t(s+\epsilon^j)\\ &= \sum_{i\neq j}(\beta_i(s,t) - \beta_i(s+\epsilon^j,t))(\bar{v}_t^i(s^i) - \underline{v}_t^i(s^i))\\ &\quad + \beta_j(s,t)\bar{v}_t^j(s^j) - \beta_j(s+\epsilon^j,t)\bar{v}_t^j(s^j+1)\\ &\quad + (1-\beta_j(s,t))\underline{v}_t^j(s^j)\\ &\quad - (1-\beta_j(s+\epsilon^j,t))\underline{v}_t^j(s^j+1).\end{aligned}$$

If we assume that $\beta_i(s,t)$ and $\beta_i(s+\epsilon^j,t)$ are very close, then the above suggests the approximation

$$\Delta_j \tilde{v}_t(s) \approx \beta_j(s,t)\Delta\bar{v}_t^j(s^j) + (1-\beta_j(s,t))\Delta\underline{v}_t^j(s^j). \tag{28}$$

When we use approximation scheme (27), expression (28) represents the approximate marginal value of a seat on flight j as determined by a weighted sum of the upper-bound marginal value $\Delta\bar{v}_t^j(s^j)$ and lower-bound marginal value $\Delta\underline{v}_t^j(s^j)$.

The choice of $\beta_i(s,t)$ will greatly affect the performance of our approximation. In our numerical experiments, constant weights, independent of state and time, work reasonably well. We will return to this issue in §9.

7.2. Static Booking-Limit Policy from Approximate Marginal Value

In (28), if we choose constant weight, then $\Delta_j \tilde{v}_t(s)$ does not depend on the value of s^i for $i \neq j$. In this case, let

$$\Delta\tilde{v}_t^j(s^j) \equiv \Delta_j \tilde{v}_t(s) = \beta\Delta\bar{v}_t^j(s^j) + (1-\beta)\Delta\underline{v}_t^j(s^j). \tag{29}$$

A booking limit for flight j in period t, b_t^j, can be determined by

$$b_t^j = \min\{0 \leqslant s^j \leqslant c^j : \Delta\tilde{v}_{t-1}^j(s^j) > f_t\}. \tag{30}$$

The method for determining booking limits in (30) is motivated by the results from single-flight revenue management models with no choice behavior. Note that absent customer-choice behavior, the marginal value given by (29) is exact, and (30) determines optimal booking limits (see Lautenbacher and Stidham 1999). With customer choice, choosing a "good" weight β is crucial for determining good booking limits. Alternatively, we can employ a heuristic search procedure by varying the weight β, computing the corresponding booking limits according to (30), then evaluating (via simulation) the resulting policies, and keeping the best one. The method is viable because the simulation step usually can be done very quickly even for big problems. The algorithm can be summarized as follows:

1. Initialize: set $b^* = 0$, $v^* = 0$.
2. Fix $0 \leqslant \underline{\beta} \leqslant \bar{\beta} \leqslant 1$ and $\delta > 0$. For $\beta = \underline{\beta}$ to $\bar{\beta}$ with step size δ do the following:
   (a) For $j = 1, \ldots, n$ and $t = 0, \ldots, m-1$, calculate the approximate marginal values $\Delta\tilde{v}_t^j(\cdot)$ using (29).
   (b) For $j = 1, \ldots, n$ and $t = 1, \ldots, m$, determine the booking limit on flight j in period t by
   $$b_t^j = \min\{0 \leqslant s^j \leqslant c^j : \Delta\tilde{v}_{t-1}^j(s^j) > f_t\}.$$
   (c) Let π^b be the static booking-limit policy where $(\pi^b)_t^j(s) = (b_t^j - s^j)^+$. Simulate the policy π^b for l replications and record the average total revenue from all the flights as $\hat{v}$. Then, $\hat{v}$ is an estimator for $u_m^{\pi^b}(0)$.
   (d) If $\hat{v} > v^*$, then $v^* = \hat{v}$ and $b^* = b$.

The output of the algorithm is a booking-limit matrix b^*, which is used to implement a static booking-limit policy. Note also that the algorithm is run "offline" before the start of the booking horizon.

In the simplest form of the algorithm, we have used lower and upper limits $\underline{\beta} = 0$ and $\bar{\beta} = 1$, with a step size δ between 0.01 and 0.05. The range of the weight may be reduced by preprocessing. For instance, one preprocessing method is to first run the algorithm with large δ (say $\delta = 0.1$) and observe the best weight $\tilde{\beta}$. Subsequently, a small step size can then be used to rerun the algorithm around a small range that contains $\tilde{\beta}$.

7.3. Dynamic Booking-Limit Policy from Approximate Marginal Value

The algorithm presented in §7.2 can also be used to derive a dynamic booking-limit policy if the algorithm is applied "on the fly" to the remaining t-period problem upon entry into each time period t. In this case, the weight is adjusted upon entry into each time period. This version of the method takes into account the state and time, and therefore has the potential of generating a better control policy. At the beginning of each time t in the booking process, observe the state s_t, and run the algorithm in §7.2 on the remaining t-period problem. The output of the algorithm is used to control bookings in period t according to $x_t = (b_t^* - s_t)^+$.

Observe that for each t the algorithm chooses an action for period t based upon the idea that it will use a static booking-limit policy from that time forward. However, this is not the case, because in subsequent periods the algorithm is run again. This is consistent with the basic philosophy of many real-world revenue management systems in which actions are selected by the repeated resolving of a particular formulation throughout the booking process.

It is possible to combine value function approximation and optimization techniques to compute actions on the fly in other ways as well. At the beginning of each period t, if we are in state s, we consider the following optimization problem:

$$\max_{0\leqslant x\leqslant c-s} \mathrm{E}[r_t(x,Q_t(x))+\tilde{v}_{t-1}(\min\{x,Q_t(x)\}+s)], \tag{31}$$

where $\tilde{v}_{t-1}(\cdot)$ could be (27) or any other approximation. The optimization problem in (31) can be solved approximately via heuristic search or simulation-based optimization. The resulting action can then be used to control the booking process in period t.

The method, however, performed poorly in most of our test problems (we do not report the numerical results here). The solution to (31) only provides an action for the current period, and thus fails to take into account the impact of future actions on value function approximation. This is different from the algorithm in §7.2, which does take into account the impact of future actions through the simulation procedure that iterates through values of β. As a result, the policy suggested by the method is, in a sense, myopic; the approximation error is accumulated period by period, and eventually leads to the large performance gaps observed in our numerical study. Furthermore, the method cannot be applied to certain situations where our model assumptions are violated. For instance, if the assumption that different fare classes arrive in distinct time periods is violated, it is necessary to set a booking limit for *each* fare class instead of just setting a booking limit for the current booking class (see, e.g., Belobaba 1989, p. 184).

8. Deterministic Linear Programming Formulation

One widely used solution approach in revenue management is deterministic linear programming. For an m-period problem with n single-leg flights and no choice behavior, one would need to solve n independent linear programming problems to get a seat allocation on each flight. For comparison purposes, we now write the n optimization problems in one formulation. Let y_j^i be the number of seats allocated in period j on flight i. (Recall that in our block-demand setting, time period and ticket class are synonymous.) Let $\mu_j^i=\mathrm{E}\underline{D}_j^i$ (hereafter we will work with the choice model of §6). The linear program for state s in period t is

$$\max_{y} \sum_{i=1}^{n}\sum_{j=1}^{t} f_j y_j^i \tag{32}$$

$$\text{s.t.} \sum_{j=1}^{t} y_j^i \leqslant c^i - s^i, \quad i=1,\ldots,n, \tag{33}$$

$$y_j^i \leqslant \mu_j^i, \quad j=1,\ldots,t;\, i=1,\ldots,n, \tag{34}$$

where constraint (33) ensures that the total number of seats allocated on each flight does not exceed the available capacity, and constraint (34) stipulates that the number of seats allocated to each class on each flight does not exceed the expected demand. Note that this is a standard formulation that is described in, e.g., Cooper (2002).

Next, we incorporate the choice effects in the linear programming model by considering a family of constraints parameterized by $\beta\in[0,1]$. The idea is to approximate the expected demand by a weighted sum of the upper- and lower-bound expected demand. Let $\mu_j=\mathrm{E}D_j=\sum_{i=1}^n \mu_j^i$ and $\lambda_j^i=\mu_j \mathrm{P}(\Theta(i)<\Theta(0))$. Note that here λ_j^i is the expected value of the upper-bound demand in (26). For state s and time t, we replace the demand constraint (34) with the following two constraints:

$$y_j^i \leqslant \beta\mu_j^i+(1-\beta)\lambda_j^i, \quad j=1,\ldots,t;\, i=1,\ldots,n, \tag{35}$$

$$\sum_{i=1}^{n} y_j^i \leqslant \mu_j, \quad j=1,\ldots,t. \tag{36}$$

Constraint (35) requires that the number of seats allocated does not exceed the expected demand approximated by a weighted sum of upper- and lower-bound expected demand, and constraint (36) ensures that the total number of seats allocated to each class does not exceed the total expected demand in that class. Observe that if there is no choice behavior among the flights, then (35)–(36) is equivalent to (34), because in that case $\lambda_j^i=\mu_j^i$, and hence constraint (36) is simply the sum of the constraints in (35). Note that regardless of β, the optimal objective function value is, in general, not the revenue associated with implementing the policy specified by the optimal solution of the linear program. Hereafter, for ease of reference, we call the linear program defined by (32)–(33) and (35)–(36) LPC.

8.1. Booking-Limit Policies Derived from the Seat Allocation Policy

The deterministic linear programming method divides seats on each flight so that the seats allocated to one booking class will not be available to other booking classes. Hence, it is possible that a higher booking class is closed while a lower booking class is still open. In particular, in our setting, this means seats allocated for a certain period, if not purchased in the period in question, will be left empty in the remaining periods. Under the reasonable assumption that the total number of bookings is increasing in the available inventory, i.e., when $\sum_{i=1}^n Q_t^i(x)$ is increasing in x, it is straightforward to show that higher expected revenue can be achieved by allowing the remaining booking classes to book the empty seats left from the previous periods. This can be done by converting a seat allocation matrix to a booking-limit matrix as follows. Let $\hat{y}_j^i$ be the number of seats allocated on flight i in period j; that is, $\{\hat{y}_j^i\}$ is an optimal solution to LPC. Recall that time is counted backwards. Let b be an $n\times t$ booking-limit matrix whose (i,j)th element b_j^i is the booking limit in period j on flight i, defined as

$$b_j^i=\sum_{k=j}^{t}\hat{y}_k^i, \quad j=2,\ldots,t, \tag{37}$$

$$b_1^i=c^i-s^i. \tag{38}$$

In the above, (38) ensures that the booking limits in the last period (Period 1) are equal to the remaining capacity. In our numerical procedures, we iterate through several values of β (as in the algorithm of §7.2), and evaluate via simulation the booking-limit policies implied by (37)–(38) and the respective solutions of LPC. We then keep the best one. We call the policy that results from this procedure LP. More formally, we obtain the policy LP by replacing Steps 2(a) and 2(b) of the algorithm in §7.2 by

2(a′)–(b′) Determine a booking-limit matrix b
by solving LPC and using (37)–(38).

The remainder of the algorithm remains the same.

8.2. Bid Pricing

Bid pricing is a revenue management method where threshold values (bid prices) are set for capacity on each flight, and a seat is sold only if the offered price is higher than the bid price. Talluri and van Ryzin (1998) discuss several different methods for producing bid prices. In our model, a set of bid prices can be produced from LPC. In particular, the bid prices are the shadow prices associated with the capacity constraint (33). Let $p_t(s) = (p_t^1(s), \ldots, p_t^n(s))$ be an n-vector, where $p_t^i(s)$ is the shadow price of capacity on flight i given the state is s in period t. A bid-price control policy specifies whether bookings should be accepted or denied in each period for each state. Let $u_t(s) = (u_t^1(s), \ldots, u_t^n(s))$ denote this decision, where $u_t^i(s) = 1$ if bookings are accepted on flight i in period t, and $u_t^i(s) = 0$ if bookings are rejected, so

$$u_t^i(s) = \begin{cases} 1 & \text{if } f_t \geqslant p_t^i(s), \\ 0 & \text{otherwise.} \end{cases}$$

After iterating through β as in the procedure to obtain LP, we will call the resulting best bid-price policy BP. Observe that, strictly speaking, formulation (5) does not allow bid-price policies, because the distribution of $Q_t(x)$ is assumed to be induced by an availability vector x. However, for choice processes as in §6, such a bid-price policy does make sense.

8.3. Re-solving

The linear program LPC can be re-solved each time period to take into account the adjustments in capacity and expected future demand. That is, prior to each re-solve, the capacity is decremented by the space taken up by booked customers, and the expected demand is replaced by the expected future demand. In the numerical experiments in §9, we consider re-solving versions of LP and BP. Re-solving for bid prices involves replacing the shadow prices with those of the reduced problem at the re-solving time.

At each re-solve point, we iterate through values of β, again evaluating the implied policies by simulation, and keeping the best one. In addition, the (adjusted) solution before re-solving is also included in the set of alternative solutions. That is, the booking limits (respectively, bid prices) considered in re-solving include those corresponding to different β values as well as the booking limits (respectively, bid prices) from the previous period adjusted by the most recent bookings.

9. Numerical Experiments

In this section, we present numerical examples to illustrate the application of the model, and to compare the performance of the different solution approaches. As discussed before, in general, static booking-limit policies are not optimal for our problem with multiple flights and dynamic substitution. However, one conclusion of our experiments is that such policies do typically work well. In the two-flight three-period setting, examples are constructed to test the methods under different levels of demand variability and interflight correlation. Understanding the relative performance of various heuristic methods in those test cases will provide some insights into the methods proposed in this paper. Finally, we report numerical results from a realistic-size problem with 16 parallel flights of 100 seats each.

It is of interest to compare the simulation results from different heuristics to the exact value of the MDP formulation. However, direct evaluation of the MDP is difficult because the joint distribution function for the demand is hard to evaluate, even if we start with very simple assumptions on demand and choice models. We use Monte Carlo methods to evaluate the value of the MDP by simulating the demand arrival and choice process to generate empirical estimates of the objective function on the right-hand side of (5). We do this for each feasible x and choose (using exhaustive search) the x that maximizes the right-hand side. This is done iteratively using the standard backwards induction algorithm for MDPs. Needless to say, this is possible only for small problems. We call this method MDPSIM.

Table 1 summarizes all the methods involved in our numerical examples. For LP and BP, we observed that re-solving almost always gives higher average revenue in the simulation, hence only the results for the version with re-solving are reported. Note that LP uses booking limits as

Table 1. Heuristics used in the numerical examples.

Method	Description
MDPSIM	Policy from evaluating MDP by simulation
LBL	Lower-bound booking limits from (7)
PBL	Booking limits from pooled upper bound
ABL	Booking limits from approximate value function as described in §7.2
DBL	Dynamic booking limits as described in §7.3
LP	Booking limits converted (see §8.1) from solutions of LPC with re-solving each period
BP	Bid pricing with bid prices generated from LPC and updated each period via re-solving

Table 2. Effects of demand variability.

Method	Simulated average			Distance from LBL			Distance from MDPSIM		
Variability	0.5	0.25	0.125	0.5 (%)	0.25 (%)	0.125 (%)	0.5 (%)	0.25 (%)	0.125 (%)
MDPSIM	9,970.79	10,431.70	10,566.90	2.05	1.28	1.13	0.00	0.00	0.00
LBL	9,770.78	10,299.64	10,448.32	0.00	0.00	0.00	−2.01	−1.27	−1.12
PBL	8,956.22	9,138.79	9,195.09	−8.34	−11.27	−11.99	−10.18	−12.39	−12.98
ABL	9,859.46	10,358.45	10,547.14	0.91	0.57	0.95	−1.12	−0.70	−0.19
DBL	9,859.46	10,358.45	10,547.14	0.91	0.57	0.95	−1.12	−0.70	−0.19
LP	9,643.18	10,127.30	10,337.45	−1.31	−1.67	−1.06	−3.29	−2.92	−2.17
BP	9,842.96	10,297.67	10,435.71	0.74	−0.02	−0.12	−1.28	−1.28	−1.24
LB	9,142.90	9,560.99	9,779.85	−6.43	−7.17	−6.40	−8.30	−8.35	−7.45
UB	11,052.62	11,802.43	12,182.21	13.12	14.59	16.59	10.85	13.14	15.29
PUB	10,239.35	10,680.63	10,794.10	4.80	3.70	3.31	2.69	2.39	2.15

described in §8.1. In our numerical results, the numbers for LB, UB, and PUB in Tables 2–5 are exact; the other numbers reflect averages over 10,000 simulations in Tables 2–4 and averages over 1,000 simulations in Table 5. (LB and UB are the bounds in Proposition 5. PUB is the bound in Proposition 4.)

As pointed out before, the relationship between the pooled upper bound and the separable upper bound depends upon the particulars of the problem. In the examples we consider, however, the separable upper bound is larger than the pooled upper bound. Our experience with many other examples is that the pooled upper bound is typically tighter. Nevertheless, we use the separable upper bound in the value function approximation because it is easier to work with. Also, as discussed below, the policy borne out of the pooled upper bound performs relatively poorly in many situations even when the no-purchase probability is very low (note that with full substitution, i.e., $P(\Theta(0)=n+1)=1$, the problem collapses to one dimension, and the pooled problem gives the optimal policy—modulo the point raised in the final paragraph of §5).

9.1. Effect of Demand Variability

We investigate the effect of demand variability by considering three two-flight three-period examples with the same capacity, fare, mean demand, and choice probability, but different demand variability. The capacity for each flight is 20. The fares are \$200, \$300, and \$400 for Periods 3, 2, and 1, respectively. The demand in each period is (truncated and rounded) normal with means of 8 and 4, respectively, for Flights 1 and 2. Within a period, we suppose that arrivals occur homogeneously through time. The choice probability is 0.8 for the first two periods, and is 0 for the last period. The examples differ in demand variability as measured by coefficient of variation (COV), which is the standard deviation divided by the mean. The COVs for the three examples are 0.5, 0.25, and 0.125, respectively. To put the choice model into the framework considered in §6, fix a period and let Z_i be the normally distributed number of customers with $\Theta(i)=1$. Then, the number of total arrivals is Z_1+Z_2. Given $Z_i=z_i$ for $i=1,2$, we sample $\sigma=(\sigma(1),\ldots,\sigma(z_1+z_2))$ from the set of $(z_1+z_2)!$ permutations of $\{1,2,\ldots,z_1+z_2\}$ uniformly at random. Then, let $M_1=\{i\colon \sigma(i)\leqslant z_1\}$ and $M_2=\{i\colon \sigma(i)>z_1\}$. For $i=1,2$, the set M_i contains the indices of customers with $\Theta(i)=1$. Given that a particular customer has $\Theta(i)=1$, we sample the remainder of his preference mapping according to $P(\Theta(i')<\Theta(0)\mid\Theta(i)=1)=0.8$ for $i'\in\{1,2\}\backslash\{i\}$ in Periods 3 and 2, and $P(\Theta(i')<\Theta(0)\mid\Theta(i)=1)=0$ in Period 1.

Table 2 reports data from the numerical examples. Observe that, with a few exceptions, almost all the values reported are decreasing in demand variability. For LB, UB, and PUB values, this result comes as no surprise and is consistent with the analysis of Cooper and Gupta (2005), who use stochastic comparisons to show that the value function of a single-leg block-demand MDP is decreasing in the demand variability. Their results do not apply to the multi-flight case with customer choice. However, our simulation results do show that the negative effects of variability typically do persist in this setting. It should be pointed out that the performance of PBL is quite poor, with revenue gap around 10%. It is also observed that the performance of LP is close to that of LBL.

The simulated averages from LBL and DBL are the same. This is, in general, the case for three-period examples with no customer choice in the last period (Period 1), because in this particular situation, the marginal revenue for the upper-bound problem and the lower-bound problem is the same (and both are exact) for $t=1$. Hence, the approximate marginal value in (29) does not depend upon the particular weight β, and therefore the booking limits from (30) are not updated at the beginning of Period 2. It follows that in this class of examples, the DBL policy coincides with the ABL policy.

9.2. Effects of Interflight Demand Correlation

We investigate the effect of interflight demand correlation by considering three two-flight three-period examples with the same capacity, fare, demand, and choice probability, but different interflight demand correlation. Note that demands across periods are still independent. The examples

Table 3. Effect of interflight demand correlation.

Method	Simulated average			Distance from LBL			Distance from MDPSIM		
Correlation	0.9	0	−0.9	0.9 (%)	0 (%)	−0.9 (%)	0.9 (%)	0 (%)	−0.9 (%)
MDPSIM	9,833.59	9,970.79	10,146.92	1.59	2.05	2.94	0.00	0.00	0.00
LBL	9,679.41	9,770.78	9,857.06	0.00	0.00	0.00	−1.57	−2.01	−2.86
PBL	8,913.56	8,956.22	9,001.72	−7.91	−8.34	−8.68	−9.36	−10.18	−11.29
ABL	9,766.57	9,859.46	9,984.26	0.90	0.91	1.29	−0.68	−1.12	−1.60
DBL	9,766.57	9,859.46	9,984.26	0.90	0.91	1.29	−0.68	−1.12	−1.60
LP	9,458.97	9,643.18	9,931.62	−2.28	−1.31	0.76	−3.81	−3.29	−2.12
BP	9,690.29	9,842.96	9,943.94	0.11	0.74	0.88	−1.46	−1.28	−2.00
LB	9,142.90	9,142.90	9,142.90	−5.54	−6.43	−7.25	−7.02	−8.30	−9.89
UB	10,807.74	11,052.62	11,413.56	11.66	13.12	15.79	9.91	10.85	12.48
PUB	9,955.55	10,239.35	10,659.79	2.85	4.80	8.14	1.24	2.69	5.05

have the same specifications as the example with COV 0.5 in §9.1. The interflight correlations are 0.9 (high positive correlation), 0 (no correlation), and −0.9 (high negative correlation), respectively. Table 3 reports data from the numerical examples.

It seems that the effect of demand correlation on revenue is moderate, with the case with high negative correlation giving slightly higher revenue. With high negative correlation, the demand of one flight tends to be low when the demand of another flight is high. The customer choice among flights enables the low-demand flight to accommodate some high-demand flight customers.

9.3. An Alternative Fare Structure

In this section, we consider a two-flight three-period example with the same capacity, demand specification, and choice probabilities as the example with COV 0.5 in §9.1. The fares are \$50, \$100, and \$400. Table 4 reports data from the numerical example. Consistent with the numerical results in §§9.1 and 9.2, PBL is inferior to LBL. To see why this is the case, note that with the pooled booking-limit policy, control is exerted on the aggregate capacity of all the flights, and booking requests are accepted as long as the total number of bookings on all the flights is less than the booking limit. Consequently, the popular flights are occupied by early-arriving low-fare customers, and late-coming inflexible high-fare customers are displaced.

Table 4. An alternative fare structure.

Method	Simulated average	Distance from LBL (%)	Distance from MDPSIM (%)
MDPSIM	6,224.63	0.09	0.00
LBL	6,218.82	0.00	−0.09
PBL	5,170.50	−16.86	−16.93
ABL	6,218.82	0.00	−0.09
DBL	6,218.82	0.00	−0.09
LP	6,012.75	−3.31	−3.40
BP	5,695.17	−8.42	−8.51
LB	5,984.74	−3.76	−3.85
UB	6,503.17	4.57	4.47
PUB	6,300.90	1.32	1.23

Somewhat different from the examples in §§9.1 and 9.2, here BP performs badly. In this particular example, BP opens both flights to all the classes in the three periods. This turns out to be a bad policy because the fare of the highest class (Class 1) is much higher than the fares of the two lower classes. This poor performance is a manifestation of the crudeness of bid-pricing policies mentioned at the end of §9.1. Also LBL, ABL, and DBL are the same in this case, and give average revenue that is only slightly lower than the MDPSIM value. This suggests that booking-limit type policies tend to be robust under different parameters.

9.4. A Large Example

For the example below, we assume a Markovian choice process, which is a special case of the choice model in §6. In a Markovian choice process, there is a random number of arriving customers, and each arriving customer has a random initial preference. If this initial preference cannot be satisfied, the customer chooses another flight with certain probabilities; if that flight is unavailable he checks another, and so on. This process repeats until a purchase is made or the customer goes away empty-handed. This situation appears to be related to Assumption 5* in Smith and Agrawal (2000). Let γ_i be the probability that a customer initially prefers flight i, and let α_{ij} be the probability a customer diverts from flight i to flight j in the choice process if flight i is unavailable. A customer continues looking for a flight until either he buys a ticket or does not purchase. We assume γ_i and α_{ij} satisfy

$$\sum_{i=1}^{n} \gamma_i = 1,$$

$$\sum_{j=0}^{n} \alpha_{ij} = 1 \quad \text{for } i = 0, 1, \ldots, n.$$

We also assume that (α_{ij}) is the transition matrix of an irreducible Markov chain with state space N_0, so that the choice process of each customer can be related to a

Markov chain $\{Z_t\}$ with $P(Z_0 = i) = \gamma_i$ and one-step transition matrix (α_{ij}). To connect this to the preference mappings of §6, define $\tau_i = \inf\{t: Z_t = i\}$, and for any permutation $i_0, i_1, \ldots, i_n$ of N_0, let

$$P\big(\Theta(i_0) = 1, \Theta(i_1) = 2, \ldots, \Theta(i_n) = n+1\big) = P\big(\tau_{i_0} < \tau_{i_1} < \cdots < \tau_{i_n}\big).$$

We also assume that different customers' preference mappings are independent.

Consider an 8-period problem with 16 flights. Each flight has 100 seats. The fares are \$300, \$400, \$500, \$1,200, \$1,400, \$1,600, \$1,800, and \$2,000 for Class 8 through Class 1. In each period $8, \ldots, 1$ the total number of arrivals D is Poisson, with respective means 320, 320, 320, 200, 200, 200, 200, and 200. Therefore, the total demand for all the flights during the booking horizon is Poisson with mean 1,960. From Period 8 to Period 6, each customer initially prefers flight i with probability $\gamma_i = 1/16 = 0.0625$ for $i = 1, \ldots, 16$. From Period 5 onward, each customer initially prefers each flight 1 through 8 with probability 0.025, and initially prefers each flight 9 through 16 with probability 0.1; i.e., $\gamma_i = 0.025$ for $i = 1, \ldots, 8$, and $\gamma_i = 0.1$ for $i = 9, \ldots, 16$. The customer-choice process is Markovian, as described earlier. The diversion probabilities are $\alpha_{ij} = 0.05$ for all $i = 0, 1, \ldots, 16$ and $j = 1, \ldots, 16$, and $\alpha_{ij} = 0.20$ for $i = 0, 1, \ldots, 16$ and $j = 0$.

Table 5 summarizes the results. The table shows that in this example static booking-limit policies do perform well, and this has been borne out in our other numerical experiments. The ABL policy, for example, gives average revenue that is only 2.13% from the upper bound. The DBL policy shows additional improvement over the ABL policy. The PBL policy, however, performs badly as in the smaller examples. The deterministic linear programming methods perform much better than in the small examples we tested. In particular, BP shows 0.89% improvement over LBL heuristics. In fact, simple (and easy-to-compute) heuristics often perform solidly, as evidenced by LBL. Mahajan and van Ryzin (2001b) have uncovered a similar phenomenon in stocking retail assortments under customer choice; that is, relatively simple policies (in their case, ones involving newsvendor-type heuristics) provide a strong baseline, which can be slightly improved upon using more sophisticated search techniques. Nevertheless, as described by Belobaba and Weatherford (1996) and Belobaba (2001), improvements in average revenue of 0.1% per flight can constitute a significant increase to the bottom line for an airline.

Table 5. 16-flight example.

Method	Simulated average	Distance from LBL average (%)	Distance from PUB value (%)
LBL	1,791,283.50	0.00	−3.32
PBL	1,651,388.00	−7.81	−10.87
ABL	1,813,504.40	1.24	−2.13
DBL	1,814,250.80	1.28	−2.08
LP	1,780,724.50	−0.59	−3.89
BP	1,807,225.70	0.89	−2.46
LB	1,729,126.01	−3.47	−6.68
UB	2,901,777.40	61.99	56.61
PUB	1,852,880.95	3.44	0.00

10. Conclusion and Future Directions

Airline revenue management has been an active research area over the past several years. However, most work assumes that the demand for each booking class follows some exogenously determined distribution. Among the papers that do consider some form of customer-choice behavior, most are concerned with how the customers decide which type of ticket to buy for a single specified flight. We have considered the problem of selecting seat-inventory control policies when the airline has many flights between a particular origin and destination within a short timespan. After formulating the problem as an MDP, we derived computable bounds for the value function, and developed some simulation-based approximate MDP procedures to obtain good policies for problems that are not amenable to an exact solution.

We have assumed that the fares for each flight during a certain period are the same. Nevertheless, examination of the proofs shows that the separable bounds we develop in Proposition 2 are valid even if the fares for different flights are different. The generality of the choice model in this paper has made it difficult to identify more specific structural characteristics of the MDP. Consequently, further study of the model with narrower assumptions on consumer behavior may yield additional results and insights. This is a direction for future work. Also, we have not considered overbooking and cancellations. Including these aspects in the model is another avenue for future research. In addition, we plan to examine situations in which customers choose among classes as well as flights.

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant DMI-0115385. The authors are grateful to the associate editor and referees for their helpful comments.

References

Belobaba, P. P. 1989. Application of a probabilistic decision model to airline seat inventory control. *Oper. Res.* **37**(2) 183–197.

Belobaba, P. P. 2001. Revenue and competitive impacts of O-D control: Summary of PODS results. Presentation at INFORMS Revenue Management Conference, New York.

Belobaba, P. P., L. R. Weatherford. 1996. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sci.* **27**(2) 343–363.

Bertsekas, D. P., J. N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.

Bodily, S. E., L. R. Weatherford. 1995. Perishable-asset revenue management: Generic and multiple-price yield management with diversion. *Omega* **23**(2) 173–185.

Boyd, E. A., I. C. Bilegan. 2003. Revenue management and e-commerce. *Management Sci.* **49**(10) 1363–1386.

Brumelle, S. L., J. I. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Oper. Res.* **41**(1) 127–137.

Brumelle, S. L., J. I. McGill, T. H. Oum, K. Sawaki, M. W. Tretheway. 1990. Allocation of airline seats between stochastically dependent demands. *Transportation Sci.* **24**(3) 183–192.

Chatwin, R. E. 1998. Multiperiod airline overbooking with a single fare class. *Oper. Res.* **46**(6) 805–819.

Cooper, W. L. 2002. Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.* **50**(4) 720–727.

Cooper, W. L., D. Gupta. 2005. Stochastic comparisons in airline revenue management. Working paper, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

Curry, R. E. 1990. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Sci.* **24**(3) 193–204.

Iyengar, S. S., M. R. Lepper. 2000. When choice is demotivating: Can one desire too much of a good thing? *J. Personality Soc. Psych.* **79**(6) 995–1006.

Karaesman, I., G. van Ryzin. 2004. Overbooking with substitutable inventory classes. *Oper. Res.* **52**(1) 83–104.

Lautenbacher, C. J., S. Stidham. 1999. The underlying Markov decision process in the single–leg airline yield management problem. *Transportation Sci.* **33**(2) 136–146.

Li, M. Z. F., T. H. Oum. 2002. A note on the single leg, multifare seat allocation problem. *Transportation Sci.* **36**(3) 349–353.

Littlewood, K. 1972. Forecasting and control of passenger bookings. *AGIFORS Sympos. Proc.* **12** 95–117.

Lovejoy, W. S. 1986. Policy bounds for Markov decision processes. *Oper. Res.* **34**(4) 630–637.

Mahajan, S., G. van Ryzin. 2001a. Inventory competition under dynamic consumer choice. *Oper. Res.* **49**(5) 646–657.

Mahajan, S., G. van Ryzin. 2001b. Stocking retail assortments under dynamic consumer substitution. *Oper. Res.* **49**(3) 334–351.

McFadden, D. 2000. Disaggregate behavioral travel demand's RUM side: A 30-year retrospective. Technical report, Department of Economics, University of California, Berkeley, CA.

McGill, J. I., G. J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.

Müller, A. 1997. How does the value function of a Markov decision process depend on the transition probabilities? *Math. Oper. Res.* **22**(4) 872–885.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. John Wiley and Sons, Chichester, UK.

Netessine, S., R. A. Shumsky. 2004. Revenue management games: Horizontal and vertical competition. *Management Sci.* Forthcoming.

Pfeifer, P. E. 1989. The airline discount fare allocation problem. *Decision Sci.* **20**(1) 149–157.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York.

Robinson, L. W. 1995. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Oper. Res.* **43**(2) 252–263.

Shumsky, R. A., F. Zhang. 2004. Dynamic capacity management with substitution. Working paper, Simon School of Business, University of Rochester, Rochester, NY.

Smith, S. A., N. Agrawal. 2000. Management of multi-item retail inventory systems with demand substitution. *Oper. Res.* **48**(1) 50–64.

Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Sci.* **44**(11) 1577–1593.

Talluri, K., G. van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* **50**(1) 15–33.

White, C. C., K. Schlussel. 1981. Suboptimal design for large scale, multimodule systems. *Oper. Res.* **29**(5) 865–875.

Wollmer, R. D. 1992. An airline seat management model for a single leg route when lower fare classes book first. *Oper. Res.* **40**(1) 26–37.

Zhao, W., Y.-S. Zheng. 2001. A dynamic model for airline seat allocation with passenger diversion and no-shows. *Transportation Sci.* **35**(1) 80–98.