

Revenue Management in a Dynamic Network Environment

Dimitris Bertsimas • Ioana Popescu

Sloan School of Management, MIT, Cambridge, Massachusetts 02139

INSEAD, Fontainebleau Cedex 77305, France

dbertsim@mit.edu • ioana.popescu@insead.edu

We investigate dynamic policies for allocating scarce inventory to stochastic demand for multiple fare classes, in a network environment so as to maximize total expected revenues. Typical applications include sequential reservations for an airline network, hotel, or car rental service. We propose and analyze a new algorithm based on approximate dynamic programming, both theoretically and computationally. This algorithm uses adaptive, nonadditive bid prices from a linear programming relaxation. We provide computational results that give insight into the performance of the new algorithm and the widely used bid-price control, for several networks and demand scenarios. We extend the proposed algorithm to handle cancellations and no-shows by incorporating oversales decisions in the underlying linear programming formulation. We report encouraging computational results that show that the new algorithm leads to higher revenues and more robust performance than bid-price control.

Introduction

Capacity constrained service industries, such as transportation, tourism, entertainment, media, and internet providers are constantly faced with the problem of intelligently allocating their limited, perishable inventories to demand from different market segments, with the objective of maximizing total revenues. Revenue management is concerned with the theory and practice underlying this type of problem. Following airline deregulation, revenue management techniques have had an important impact on the development of the industry, providing up to 4%–10% increases in company revenues (Fuchs 1987). For example, in 1997, American Airlines collected one billion dollars by implementing revenue management, representing most of the company's profit (Cook 1998).

Optimization techniques have been essential in the development of revenue management tools, particularly for seat allocation models. In this research, we are interested in investigating the design of dynamic

policies for allocating inventory to correlated, stochastic demand for multiple classes, in a network environment. Specifically, we design a decision support tool, based on stochastic and dynamic optimization techniques, that at each point in time accepts or rejects a reservation request, based on the currently available inventory, past sales and future potential demand, so as to maximize total expected revenues.

Problem Definition

The main problem we address in this paper is as follows. We are given an airline (hotel, car rental) network composed of l legs (pairs of consecutive days for hotels, car rentals), which are used to serve a total of m demand classes. The initial inventory is given by a vector $\mathbf{N} = (N_1, \dots, N_l)$ of leg capacities. The network is described by a $l \times m$ matrix \mathbf{A} and a m -vector $\mathbf{R} = (R_1, \dots, R_m)$: R_j is the fare category of class j , which utilizes a_{ij} units of resource (leg) i . In this way,

a demand class j is defined by its itinerary \mathbf{A}^j (a column of matrix \mathbf{A}) and its fare category R_j . In an airline network without group discounts, \mathbf{A} is a 0–1 matrix which may contain repeated columns for each fare class on a given itinerary. To account for special group fares, integer multiples of the itinerary-incidence vector are allowed.

For example, consider a very simple network corresponding to a weekend in a hotel: There are three nodes *Fri*, *Sat*, *Sun* and $l = 2$ legs (1) *Fri–Sat* and (2) *Sat–Sun* with total capacities $\mathbf{N} = (N_1, N_2)$. Suppose there is demand for all types of stays, i.e., “itineraries:” (1) *Fri–Sat*, (2) *Sat–Sun*, and (3) *Fri–Sat–Sun*, with two (*high* and *low*) fare classes for each type. Moreover, suppose there are discounts for groups of size $k_1 = 10$ for (1) *Fri–Sat* night stays, at a rate of $R_1^{(10)}$ per group, that is a total of $m = 7$ classes. The leg-class incidence matrix, together with the corresponding fare structure \mathbf{R} , is given by:

$$\begin{pmatrix} \mathbf{R} \\ \mathbf{A} \end{pmatrix} = \begin{pmatrix} R_1^h & R_1^l & R_2^h & R_2^l & R_3^h & R_3^l & R_1^{(10)} \\ 1 & 1 & 0 & 0 & 1 & 1 & 10 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}.$$

We assume a finite booking horizon of length T , with the time line sufficiently discretized so as to allow at most one request (reservation or cancellation) per time period almost surely (a.s.). Time is counted backwards: Time $t = T$ is the beginning of the booking horizon and time $t = 0$ is the end of the reservation period, where the no-shows are being counted. Customers who do not consume their reservations get full, partial, or no refund, depending on their fare class. If at the end of the horizon, the inventory is oversold, at time $t = -1$ redistribution decisions are being made. These can be class upgrades or customer bumping, in which case companies pay overbooking penalties. For the case of hotels and car rentals, an infinite time horizon could be more appropriate; however, one can decompose the problem in fixed length time periods (one month, one year, etc.).

The demand (to come) process at time t is denoted by \mathcal{D}^t , and $\overline{\mathcal{D}}^t$ represents the corresponding random vector of cumulative demands. That is, $\overline{\mathcal{D}}_j^t$ is a random variable representing the number of class j requests to come from time t until departure (order

does not matter). Usually, we have partial information about the demand process, which might consist of the expected demand to come $\mathbf{D}^t = E[\overline{\mathcal{D}}^t]$, and possibly other type of information available from forecasting tools, such as cancellation or no-show probabilities.

The state of the system \mathbf{S} is given by the time t (t periods to departure) and the sales-to-date record $\mathbf{s} = (s_1, \dots, s_m)$ for each demand class. If cancellations and no-shows are not allowed, then it is sufficient to define the state based on the remaining inventory $\mathbf{n} = (n_1, \dots, n_i)$. A common practice is to use the latter model and account for cancellations and no shows by incorporating a virtual increase, called overbooking pads, in the initial capacity definition.

The general stochastic, dynamic inventory control problem for network revenue management (NRM) can thus be stated as follows: At time t , and given that the state of the network is \mathbf{S} , should we accept or reject a new class j request? The overall objective is to maximize total expected revenues.

The decision to accept or reject determines an admission control policy, which is in general a function of the current network configuration (\mathbf{s} or \mathbf{n}), the time-to-go t , the currently requested fare R_j , as well as partial information from demand forecasts. If we accept the request, the new state becomes $(\mathbf{s} + k_j \cdot \mathbf{e}_j, t - 1)$, where k_j is the size of the class j group request, and \mathbf{e}_j is the j th unit vector; when cancellations are not allowed, the state of the network (\mathbf{n}, t) becomes simply $(\mathbf{n} - \mathbf{A}^j, t)$. If the request is rejected, only the time component of the state vector is changed to $t - 1$.

In reality, the control policy has an impact on the demand process, because customer choice may depend on the opportunity set. Capturing and quantifying this type of feedback phenomenon is a subtle task that goes beyond the scope of this work. We will thus make the simplifying assumption that the demand process is independent of the control policy.

Notation

We will use the following notation throughout the paper. Vectors will be denoted in bold, and random variables and processes in calygraphic style.

Time:

T = length of time horizon (number of time periods);

t = time periods left until departure (count-down).

Network:

l = number of legs in the network;

m = number of classes (itineraries with fare categories);

\mathbf{N} = total initial network capacity (l -vector);

\mathbf{A} = leg-class incidence ($l \times m$)-matrix;

a_{ij} = quantity of resource i utilized by bundle j .

Sales and inventory:

\mathbf{n} = remaining inventory (l -vector);

\mathbf{s} = sales to date vector (m -vector);

s_j^t = the number of itineraries sold to class j until time t ;

s_j^o = the number of class j itineraries overbooked at departure.

Fares, refunds, and penalties:

R_j = revenue collected for one class j sold;

R_j^c = the refund per class j cancellation;

R_j^{ns} = the refund per class j no-show;

C_j = the overbooking penalty per demand class j .

Demand:

p_j^t = probability of request for class j at time t ;

p_j^{ct} = the probability of a class j cancellation occurring at time t ;

p_0^t = the probability of no request (reservation or cancellation) at time t ;

p_j^{ns} = the probability that a class j reservation will not show up;

\mathcal{D}^t = demand (to come) process (m -dimensional);

$\bar{\mathcal{D}}^t$ = aggregate demand (to come) distribution (m -dimensional);

$\mathbf{D}^t = E[\bar{\mathcal{D}}^t]$ = expected aggregate demand to come (m -dimensional);

\mathcal{C} = the cancellation process, adapted to the sales history (not a decision);

\mathcal{NS} = the no-show distribution (adapted to final sales, adjusted for cancellations).

We will use the operator $(x)^+ = \max(x, 0)$, for $x \in \mathbb{R}$, which naturally extends for vectors: $(\mathbf{x})^+ = ((x_1)^+, \dots, (x_n)^+)$ for $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$.

Contributions

In this paper, we evaluate from different perspectives several policies for solving the dynamic and stochastic

NRM problem. To compare these different policies, we provide structural and computational results.

The most popular technique developed in the current literature is an additive bid-pricing approach. These are mechanisms whereby the opportunity cost of each itinerary is estimated as the sum of the shadow prices of the incident legs, obtained from a linear programming formulation of the problem (see Formulation (2) in §2.2). However, there are two obvious drawbacks to additive bid prices:

(a) They are not well defined if there are multiple dual solutions.

(b) They are restrictive in their way of taking into account bundles by their predefined additive structure. In particular, they do not account for changes of a dual basis in response to accepting large-group and multileg itinerary requests.

The contributions of this paper are as follows:

(1) We propose an efficient control policy, that is well defined if there are multiple dual solutions, and does not have an additive structure. The proposed control policy, which we call certainty equivalent control (CEC), belongs in the class of approximate dynamic programming mechanisms (see Bertsekas and Tsitsiklis 1998), in which the cost-to-go function is approximated by the value of a linear programming (LP) relaxation. We remark that this LP is equivalent to a network flow problem in the case of origin-destination (OD) demands or linear networks (where nodes can be ordered so that the arcs are pairs of consecutive nodes $(i, i+1)$, $i = 1, \dots, l$).

(2) We provide structural properties that compare the behavior of the proposed CEC policy with the additive bid-price approach. These results offer insight into the behavior of both methods.

(3) We propose several algorithmic improvements of the CEC policy based on approximate dynamic programming.

(4) We provide computational results that give insight into the performance of these algorithms and several variations, for different networks and demand scenarios. We observe that the CEC algorithm performs very well in practice, giving results that are very close to optimum. For high load factors, we observe an average 5%–10% improvement over existing policies (additive bid pricing). We describe and

simulate extensions of this algorithm that result in significantly higher improvements (up to 20%). Interestingly, the CEC policy appears to be significantly more robust to noise and bias in the demand forecast.

(5) We extend these algorithms to handle cancellations and no-shows by incorporating overbooking control in the underlying mathematical programming formulation. This extension preserves several structural properties. We report computational results that show that the proposed algorithm improves upon the performance of the bid-price control policy.

Structure

The remainder of the paper is organized as follows: In the next section, we present an overview of the literature. Section 2 describes several (dynamic, stochastic, linear, and network flow) models and formulations for the NRM problem and evaluates the relationships between them. Section 3 presents efficient algorithms for the NRM problem based on ideas from approximate dynamic programming. In §4, we present structural properties for the proposed CEC policy and contrast them with additive bid pricing. In §5, we extend our model and algorithms to handle cancellations, no-shows, and overbooking. Finally, in §6, we present computational results. The last section summarizes our conclusions.

1. Literature Review and Positioning

The NRM problem can be viewed as a particular instance of the general class of perishable asset revenue management problems (PARM). Initial developments in static single leg revenue management are due to Littlewood (1972), followed by Simpson (1989) and Belobaba (1987), who proposed a suboptimal policy for computing protection levels based on expected marginal seat revenues (EMSR). Curry (1989), Wollmer (1992), and Brumelle and McGill (1993), derive the optimal solution for the single leg static model. Robinson (1995) proposes an extension that handles nonmonotonic fare classes.

A characterization of the optimal dynamic policy based on a *threshold time property* is due to Diamond

and Stone (1991). An analogous discrete time solution is provided by Lee and Hersh (1993), who also provide a method for estimating arrival rates. Subramanian et al. (1999) propose a dynamic programming model that handles cancellations and overbooking, by analogy to a problem in the optimal control of admission in a queueing system.

A natural, but much harder question is to determine which of the single leg results can be extended to network (multiproduct) settings, and how. A major conceptual advance in the study and practice of network revenue management was introduced by *bid-price* control. These are additive, leg-based shadow prices used to approximate the opportunity cost of itinerary capacity. The concept was proposed by Simpson (1989), and further analyzed by Williamson (1992) in extensive simulation studies. A deficiency of such mathematical programming based models is that they do not account for nesting of the fare classes. To overcome this problem, Curry (1992) proposes a virtual nesting method. In all cases, however, the allocation of capacity to itinerary demand is decided by a one-time, static rule (one fixed set of bid prices).

The most realistic and relevant, yet least investigated model for NRM is the dynamic network model. Talluri and van Ryzin (1998) study a dynamic network model using bid-price control mechanisms and argue why bid-price policies are not optimal in general. They provide an asymptotic regime when certain bid-price controls, based on a probabilistic programming formulation of the problem, are asymptotically optimal. In the context of hotels, Bitran and Mondschein (1995) propose a dynamic policy that extends to multiple-night stays, but do not give any further analysis.

Chen et al. (2000) formulate the problem as a Markov decision problem, and use linear programming and regression splines to approximate the value function. Gunther et al. (2000) introduce a new method to compute bid price for single hub airline networks. Both studies report encouraging simulation results. None of these take into account cancellations.

For further pointers to related literature, we refer the interested reader the survey of McGill and

van Ryzin (1999), and for a schematic summary the Ph.D. thesis of Popescu (1999).

Relative Positioning. Our approach can be viewed as extending the single-leg models investigated by Lee and Hersh (1993) and Bitran and Mondschein (1995), who actually propose a similar LP-based heuristic for the multiple-night booking problem in a hotel, but without any further analysis. Our overbooking model is similar to the early single-leg DP formulation of Rothstein (1971, 1974). As far as static network models are concerned, our LP formulation is similar to the one proposed by Williamson (1992), but we handle multiple classes and group bookings. The network flow formulation we provide for linear networks and origin-destination fare (ODF) demands, is the same as those proposed by Glover et al. (1982) for airlines and Chen (1998) for hotels.

2. General Models and Relationships

The problem of dynamic inventory control for NRM belongs in the class of finite horizon decision problems under uncertainty. In this section, we present several optimization models for addressing the NRM problem, and discuss various relationships between them. Again, we assume that the control policy does not feed back into the evolution of the demand process. We restrict to the case when cancellations and overbooking are not permitted; this case is discussed in detail in §5.

2.1. The Dynamic Programming Model

The stochastic dynamic programming model provides the *optimal policy* for the NRM problem, by evaluating the whole tree of possibilities and making at each point in time the decision (to sell or not to sell) that would imply higher future expected revenues.

The states $\mathbf{S} = (\mathbf{n}, t)$ are defined by the current available capacity vector \mathbf{n} when the remaining time is t . The stochasticity is given by the demand process to come \mathcal{D}^t for the remaining t periods. We define $DP(\mathbf{n}, t)$ to be the maximum expected revenue to be collected from state (\mathbf{n}, t) . Assuming independent

demands, this can be computed via the Bellman equation as follows:

$$\begin{aligned} DP(\mathbf{n}, t) &= \sum_{j=1}^m p_j^t \max(DP(\mathbf{n}, t-1), R_j + DP(\mathbf{n} - \mathbf{A}^j, t-1)) \\ &\quad + p_0^t DP(\mathbf{n}, t-1) \\ &= DP(\mathbf{n}, t-1) + \sum_{j=1}^m p_j^t (R_j - DP(\mathbf{n}, t-1) \\ &\quad + DP(\mathbf{n} - \mathbf{A}^j, t-1))^+ \end{aligned} \quad (1)$$

for all $\mathbf{n} \leq N$, $t \leq T$, with the boundary conditions:

$$\begin{aligned} DP(\mathbf{n}, t) &= -\infty \quad \text{if } n_i < 0 \text{ for some } i, \quad \text{and} \\ DP(\mathbf{n}, 0) &= 0, \quad \text{for } \mathbf{n} \geq 0. \end{aligned}$$

We define the opportunity cost $OC_j(\mathbf{n}, t)$ to be

$$OC_j(\mathbf{n}, t) = DP(\mathbf{n}, t-1) - DP(\mathbf{n} - \mathbf{A}^j, t-1).$$

Then, the optimal policy accepts a request if and only if the corresponding fare R_j exceeds its current opportunity cost. The value function can thus be expressed as follows:

$$DP(\mathbf{n}, t) = DP(\mathbf{n}, t-1) + \sum_{j=1}^m p_j^t (R_j - OC_j(\mathbf{n}, t))^+.$$

One can easily show that the value function is non-decreasing in \mathbf{n} and t . Moreover, for the single leg case without batch arrivals, the value function is concave and opportunity costs decrease with t and \mathbf{n} (see Diamond and Stone 1991). Based on these properties, one can show that the optimal policy is characterized by threshold times. Threshold times are points in time during the booking horizon before which requests are rejected, and after which requests are accepted. In the general network case, however, this is not true, as we will show in an example in §4.1.

2.2. The Integer and Linear Programming Model (IP, LP)

The most frequently utilized formulations for the NRM problem are static models. These are deterministic analogues of the stochastic dynamic problem,

that use only expected demand information, and are usually much simpler to solve.

Suppose that all the available demand information consists of (unbiased) forecasts of the expected aggregate demand to come $\mathbf{D}^t = E[\overline{\mathcal{D}}^t]$ over the remaining t periods. The integer programming model computes the optimal allocation \mathbf{y}^* of available inventory \mathbf{n} to the expected itinerary demand \mathbf{D}^t , by maximizing total revenues subject to capacity and itinerary demand constraints. For all $\mathbf{n} \leq \mathbf{N}$ and $t \leq T$, we have

$$\begin{aligned} \text{IP}(\mathbf{n}, \mathbf{D}^t) = \max \quad & \mathbf{R}' \cdot \mathbf{y} \\ \text{s.t.} \quad & \mathbf{A} \cdot \mathbf{y} \leq \mathbf{n} \\ & \mathbf{0} \leq \mathbf{y} \leq \mathbf{D}^t \\ & \mathbf{y} \text{ integer.} \end{aligned}$$

The linear programming relaxation of this problem provides an efficient way to compute the best possible “fractional” allocation of inventory, and is defined by simply relaxing the integrality constraint:

$$\begin{aligned} \text{LP}(\mathbf{n}, \mathbf{D}^t) = \max \quad & \mathbf{R}' \cdot \mathbf{y} \\ \text{s.t.} \quad & \mathbf{A} \cdot \mathbf{y} \leq \mathbf{n} \\ & \mathbf{0} \leq \mathbf{y} \leq \mathbf{D}^t. \end{aligned} \quad (2)$$

Clearly, we have that $\text{IP}(\mathbf{n}, \mathbf{D}^t) \leq \text{LP}(\mathbf{n}, \mathbf{D}^t)$. This inequality can be strict, but there are particular instances when equality holds, one of which (linear networks) we describe in the next section. One can obtain further insight into the optimal LP-allocation by looking at the dual problem:

$$\begin{aligned} \text{LP}(\mathbf{n}, \mathbf{D}^t) = \min \quad & \mathbf{v}' \cdot \mathbf{n} + \mathbf{u}' \cdot \mathbf{D}^t \\ \text{s.t.} \quad & \mathbf{v}' \cdot \mathbf{A} + \mathbf{u}' \geq \mathbf{R}' \\ & \mathbf{u}, \mathbf{v} \geq \mathbf{0}. \end{aligned}$$

This formulation can be equivalently written as:

$$\begin{aligned} \text{LP}(\mathbf{n}, \mathbf{D}^t) &= \min_{\mathbf{v} \geq \mathbf{0}} \mathbf{v}' \cdot \mathbf{n} + (\mathbf{R}' - \mathbf{v}' \cdot \mathbf{A})^+ \cdot \mathbf{D}^t \\ &= \min_{k \in K} \mathbf{v}'_k \cdot \mathbf{n} + \mathbf{u}'_k \cdot \mathbf{D}^t, \end{aligned}$$

where K denotes the index set of extreme points $(\mathbf{v}_k, \mathbf{u}_k)$ of the dual polyhedron. Thus, the objective value of Model (2) is a *piecewise linear, concave, and nondecreasing* function of the expected demand to come \mathbf{D}^t and available capacity \mathbf{n} . We next relate the objective values of the dynamic and static formulations at any state (\mathbf{n}, t) .

PROPOSITION 1. $\text{DP}(\mathbf{n}, t) \leq \text{LP}(\mathbf{n}, \mathbf{D}^t)$.

For a full proof see Popescu (1999). The idea is that for each pathwise realization of demand, the revenue collected by the DP-policy along that path is upper bounded by the value of a “perfect hindsight” stochastic program. The perfect hindsight model, denoted PI, determines the optimal allocation of inventory for each particular realization of demand, and then computes the expected reward over all scenarios:

$$\begin{aligned} \text{DP}(\mathbf{n}, t) &\leq \text{PI}(\mathbf{n}, \mathcal{D}^t) = E_\pi[\text{LP}(\mathbf{n}, \mathcal{D}_\pi^t)] \\ &\leq \text{LP}(\mathbf{n}, E_\pi[\mathcal{D}_\pi^t]) = \text{LP}(\mathbf{n}, \mathbf{D}^t). \end{aligned}$$

Here, \mathcal{D}_π^t is the cumulative demand corresponding to scenario π . Since the LP value is concave in the demand vector, the second inequality follows from Jensen’s inequality. Furthermore, the two values converge as demand becomes very large (see Popescu 1999).

2.3. Origin-Destination RM and Linear Networks

Consider a particular case of the NRM problem, where demand is by ODF, as opposed to itinerary specific. This is by default the case for linear networks, where there is no question of routing, such as in hotel or car rental revenue management (where the nodes are days). In the case of airline revenue management this situation occurs when there are perfectly substitutable routes (same price, same travel time, and so on). In these cases, the Model (2) can be represented as a network flow problem. The advantage of this formulation is that it is very easy to solve in practice and reoptimization is very fast. For airlines, this model provides (for the company) *optimal routing of path-indifferent customers*.

The network representation is as follows: The nodes are the origins and destinations (days for hotels, {airport, time}-pairs for airlines). There are multiple forward arcs $(o, d)_f$ representing the flow from origin o to destination d from fare-class f . The capacity of each forward arc is the corresponding aggregate demand D_{od}^f . The revenue collected along arc $(o, d)_f$ is R_{od}^f . For each “leg” (i, j) (flight leg, respectively, pair of consecutive days) there is a backward arc of the type (j, i) ,

capacitated by n_{ij} , the available inventory of leg (i, j) . The revenue collected along backward arcs is zero. This model has been proposed by Glover et al. (1982) in the context of airlines, and for hotel revenue management by Chen (1998), who proves that it reduces to solving a network flow problem.

PROPOSITION 2 (INTEGRALITY OF SOLUTION). *For ODF models with integer data, the optimal LP-solution is integral, that is $IP(\mathbf{n}, \mathbf{D}^t) = LP(\mathbf{n}, \mathbf{D}^t)$.*

3. Approximate Dynamic Programming Algorithms

In this section, we present several algorithms for the NRM problem. Most of these algorithms belong in the generic class of approximate dynamic programming methods (see Bertsekas and Tsitsiklis 1998), in which an approximate value to the exact value function is used in the Bellman equation.

Given a certain efficient mathematical programming formulation (MP) of the NRM problem, a generic approximate DP algorithm for the NRM problem has the following structure:

Generic MP-Policy. Identify an efficient formulation MP of the NRM problem.

At any current state $\mathbf{S} = (\mathbf{n}, t)$,

(1) For a class j request, compute an MP-based estimate of the opportunity cost $OC_j^{\text{MP}}(\mathbf{S})$.

(2) Sell to class j if and only if its fare R_j exceeds its opportunity cost estimate, i.e.,

$$R_j \geq OC_j^{\text{MP}}(\mathbf{S}).$$

(3) Go to Step 1 and ITERATE.

We will denote the expected value of this policy at an initial state \mathbf{S} as $\Pi_{\text{MP}}(\mathbf{S})$. The difference between various algorithms comes from the approximating MP and the MP-based opportunity cost measure, that is from Step 1.

3.1. Bid-Price Control

Bid-price control is a popular method in NRM, whereby the opportunity cost (shadow price, or *bid price*) of an itinerary is approximated by the sum of opportunity costs of the legs along that itinerary. First, opportunity cost estimates (*bid prices*) are determined for each leg in the network, usually as the leg-shadow

prices \mathbf{v} from the linear programming formulation (2). Then itinerary bid prices are computed additively, at each state $\mathbf{S} = (\mathbf{n}, t)$, and for each class j request as

$$BP_j(\mathbf{S}) = (\mathbf{v}^{\mathbf{S}})' \cdot \mathbf{A}^j. \quad (3)$$

Notice that bid prices depend on the choice of optimal dual variables $\mathbf{v}^{\mathbf{S}}$. This technique was initially proposed by Simpson (1989), then studied by Williamson (1992) in her Ph.D. thesis. Several probabilistic models have been investigated by Glover et al. (1982), Williamson (1992), and Talluri and van Ryzin (1998). It has been observed (Williamson 1992) that the LP model achieves a better performance, and is more efficient. For recent work on bid-price control see Talluri and van Ryzin (1998) and Gunther et al. (2000).

3.2. Certainty Equivalent Control

The main disadvantages that are apparent from the definition of leg-based additive bid prices is that (a) they are not uniquely defined (several sets of shadow prices may be optimal), and (b) they provide an *additive* approximation of the opportunity costs, which are not necessarily additive due to “bundle effects” (group or multileg itinerary requests may determine basis changes in the dual LP).

We provide a different approximation scheme for opportunity costs, based on certainty equivalent adaptive control. The idea is to approximate the value function of the dynamic program $DP(\mathbf{n}, t)$ defined in Equation (1) by the value of the linear programming problem $LP(\mathbf{n}, \mathbf{D}^t)$ defined in §2.2. Thus, in Step 1 of the generic algorithm, a request for a given class j will be accepted if and only if its price R_j exceeds its current opportunity cost estimate given by:

$$OC_j^{\text{LP}}(\mathbf{n}, t) = LP(\mathbf{n}, \mathbf{D}^{t-1}) - LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1}).$$

Because the opportunity cost estimate is calculated in terms of LP objectives, it is uniquely defined, in that its value will not depend on the choice of (dual) solution. Thereby, the first drawback of bid prices is resolved.

This is the certainty equivalent control (CEC) policy, and we denote its expected value for an initial state \mathbf{S} as $\Pi_{\text{LP}}(\mathbf{S}) = \text{CEC}(\mathbf{S})$. For more on certainty

equivalence, see Bertsekas (1995). A similar policy is proposed by Bitran and Mondschein (1995) in the context of hotel revenue management, but no analysis is provided. Notice that in the case of linear networks, it is actually desirable to use the equivalent network flow formulation as described in §2.3 instead of the usual LP model because reoptimization at each stage becomes a much simpler task.

More generally, one can use this technique with virtually any mathematical programming (MP) model that provides an approximation of the value function. The corresponding OC estimate for an itinerary j request at the current state \mathbf{S} is defined as $\text{OC}_j^{\text{MP}}(\mathbf{S}) = \text{MP}(\mathbf{S}_{\text{rej}(j)}) - \text{MP}(\mathbf{S}_{\text{acc}(j)})$, where $\mathbf{S}_{\text{acc}(j)}$ and $\mathbf{S}_{\text{rej}(j)}$ are the states corresponding to the accept, and respectively reject decision for the current request j . For example, for the static LP model without cancellations and overbooking, if $\mathbf{S} = (\mathbf{n}, t)$, then $\mathbf{S}_{\text{acc}(j)} = (\mathbf{n} - \mathbf{A}^j, t - 1)$ and $\mathbf{S}_{\text{rej}(j)} = (\mathbf{n}, t - 1)$.

3.3. Extensions

In this section, we describe several extensions that provide improvements in the efficiency or accuracy of the adaptive policies described above.

Rollout Policy. The expected value $\Pi_H(\mathbf{S})$ of any heuristic H started in state \mathbf{S} provides an approximation of the $\text{DP}(\mathbf{S})$ value. Therefore, the expected heuristic values can in turn be used to provide estimates of opportunity costs, determined as follows:

$$\text{OC}_j^{R(H)}(\mathbf{n}, t) = \Pi_H(\mathbf{n}, t - 1) - \Pi_H(\mathbf{n} - \mathbf{A}^j, t - 1).$$

This opportunity cost estimation mechanism leads to a new approximate dynamic programming (ADP) heuristic, called the rollout of H , and denoted $R(H)$. This method is simply a form of policy iteration and is described in detail in Bertsekas and Tsitsiklis (1998). It has been observed in the dynamic programming literature that this procedure systematically improves heuristic performance (see also Bertsimas et al. 1999, Bertsekas et al. 1997, Bertsekas and Castanon 1998).

For practical purposes, we suggest using Monte Carlo simulation for evaluating the policy value Π_H for a subset of states, and then interpolating these in an online fashion. An interesting research idea is to investigate what types of preprocessing simulations would provide an insightful information database.

Simulations Using Monte Carlo Demand Estimation. One problem with the certainty equivalent policy is that it only considers expected demand information, and uses a deterministic approach to a highly stochastic problem. We propose a variation on the certainty equivalent policy that uses Monte Carlo demand estimation to capture demand variability.

Suppose we have information that the cumulative demand to come $\overline{\mathcal{D}}^{t-1}$ follows a certain distribution. We generate r samples from this distribution: $\widehat{\mathbf{D}}_1^{t-1}, \dots, \widehat{\mathbf{D}}_r^{t-1}$. In Step 1 of the generic algorithm, we calculate the opportunity cost estimate of itinerary j as a weighted average,

$$\text{OC}_j^{\text{MC}}(\mathbf{n}, t) = \sum_{i=1}^r \alpha_i \cdot \text{OC}_j^{\text{LP}}(\mathbf{n}, \widehat{\mathbf{D}}_i^{t-1}),$$

where $\alpha_i = P(\overline{\mathcal{D}}^{t-1} = \widehat{\mathbf{D}}_i^{t-1} \mid \overline{\mathcal{D}}^{t-1} \in \{\widehat{\mathbf{D}}_1^{t-1}, \dots, \widehat{\mathbf{D}}_r^{t-1}\})$ and $\text{OC}_j^{\text{LP}}(\mathbf{n}, \widehat{\mathbf{D}}_i^{t-1}) = \text{LP}(\mathbf{n}, \widehat{\mathbf{D}}_i^{t-1}) - \text{LP}(\mathbf{n} - \mathbf{A}^j, \widehat{\mathbf{D}}_i^{t-1})$.

One difficulty with implementing this procedure is that we might not have enough information about the aggregate demand and/or it may be too expensive to compute the actual value of the conditional probabilities α_i . For this reason, we run a simplified version of this policy, that assigns the same weights to all the trials: $\text{OC}_j^{\text{MC}}(\mathbf{n}, t) = (1/r) \cdot \sum_{i=1}^r \text{OC}_j^{\text{LP}}(\mathbf{n}, \widehat{\mathbf{D}}_i^{t-1})$.

4. Structural Properties

In this section, we derive several structural properties of the new approximate dynamic programming algorithm (CEC) and compare it with additive bid-price control algorithms (BPC) developed in the literature.

Given that the NRM problem requires a real-time response, it is desirable to use computationally inexpensive models to construct approximations of the opportunity cost. This motivates the choice for using the LP formulation described in §2.2. We are interested in a comparative structural assessment of the two LP-based policies described previously, BPC and CEC, and their corresponding opportunity cost approximations.

From an asymptotic point of view, it should be noted that the CEC policy is asymptotically optimal in the fluid scaling regime proposed by Talluri and van Ryzin (1998), whereby demand and capacities are simultaneously increased in a way that preserves

their relative values constant. They prove that in this regime, the additive bid-pricing policy converges to the optimum, as bid prices are being held fixed. By imitating their proof, one can show that the same property holds for the CEC policy, with deterministic prices, as OC estimates are being held fixed (see Popescu 1999).

The next result compares the opportunity cost approximations of a class j request at any given state.

PROPOSITION 3. *In any state (\mathbf{n}, t) , for any bid prices $BP_j(\mathbf{n}, t)$, $BP_j(\mathbf{n} - \mathbf{A}^j, t)$, the following inequalities hold:*

$$BP_j(\mathbf{n}, t) \leq OC_j^{LP}(\mathbf{n}, t) \leq BP_j(\mathbf{n} - \mathbf{A}^j, t). \quad (4)$$

Inequalities are strict if accepting class j must incur a change of basis in the LP dual.

PROOF. Recall that we have defined:

$$\begin{aligned} OC_j^{LP}(\mathbf{n}, t) &= LP(\mathbf{n}, \mathbf{D}^{t-1}) - LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1}) \\ &= (\mathbf{v}^{n,t})' \cdot \mathbf{n} + (\mathbf{u}^{n,t})' \cdot \mathbf{D}^{t-1} \\ &\quad - (\mathbf{v}^{n-A^j,t})' \cdot (\mathbf{n} - \mathbf{A}^j) - (\mathbf{u}^{n-A^j,t})' \cdot \mathbf{D}^{t-1}, \end{aligned}$$

where $(\mathbf{v}^{n,t}, \mathbf{u}^{n,t})$ and $(\mathbf{v}^{n-A^j,t}, \mathbf{u}^{n-A^j,t})$ are optimal dual solutions of $LP(\mathbf{n}, \mathbf{D}^{t-1})$ and $LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1})$, corresponding to the given bid prices: $BP_j(\mathbf{n}, t) = (\mathbf{v}^{n,t})' \cdot \mathbf{A}^j$ and $BP_j(\mathbf{n} - \mathbf{A}^j, t) = (\mathbf{v}^{n-A^j,t})' \cdot \mathbf{A}^j$, respectively. Because both solutions are feasible for both programs, we obtain the following upper bounds by evaluating each LP at the optimal solution of the other:

$$\begin{aligned} LP(\mathbf{n}, \mathbf{D}^{t-1}) &\leq (\mathbf{v}^{n-A^j,t})' \cdot \mathbf{n} + (\mathbf{u}^{n-A^j,t})' \cdot \mathbf{D}^{t-1}, \\ LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1}) &\leq (\mathbf{v}^{n,t})' \cdot (\mathbf{n} - \mathbf{A}^j) + (\mathbf{u}^{n,t})' \cdot \mathbf{D}^{t-1}. \end{aligned}$$

The upper bound in Equation (4) follows by applying the first of these inequalities in the OC formula, and the lower bound by the second inequality,

$$\begin{aligned} BP_j(\mathbf{n}, t) &= (\mathbf{v}^{n,t})' \cdot \mathbf{A}^j \leq OC_j^{LP}(\mathbf{n}, t) \leq (\mathbf{v}^{n-A^j,t})' \cdot \mathbf{A}^j \\ &= BP_j(\mathbf{n} - \mathbf{A}^j, t). \end{aligned} \quad (5)$$

In case the two dual optimal solutions coincide, we obtain equality throughout. \square

In general, if at a given state the CEC policy accepts a class j request, then at the same state, the bid-pricing

Table 1 The Behavior of the BPC and CEC Policies as a Function of an Optimal Primal Solution \mathbf{y}^*

y_j^*	$y_j^* \geq \min(D_j, 1)$ in some \mathbf{y}^*	$y_j^* < \min(D_j, 1)$ in all \mathbf{y}^* , but $\neq 0$ in some	$y_j^* = 0$ in all \mathbf{y}^*
CEC	accept	reject	reject
BPC	accept	accept	reject

policy will also accept, but not vice versa. This is because the following situation may occur: $BP_j(\mathbf{n}, t) \leq R_j < OC_j^{LP}(\mathbf{n}, t)$, (see §4.1).

Table 1 provides a comparative characterization of the bid-pricing policy (BPC) versus the CEC policy, in terms of the structure of the primal optimal solutions \mathbf{y}^* of the LP model (2). We assume the BPC policy is well defined, in that the dual optimal solution is unique. We also assume that the dual basis is not the same for $LP(\mathbf{n}, \mathbf{D}^{t-1})$ and $LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1})$; if the dual basis does not change, then the policies are identical. The following two propositions state and prove these results formally.

PROPOSITION 4 (STRUCTURAL PROPERTIES OF THE BPC POLICY). *At any state (\mathbf{n}, t) , if $LP(\mathbf{n}, \mathbf{D}^{t-1})$ has a unique dual optimal solution, then the corresponding bid-price policy accepts only classes j for which $y_j^* > 0$ in some primal optimal solution.*

PROOF. By analyzing the primal and dual LP, one can distinguish the following situations:

- In all optimal LP-solutions $y_j^* = 0$, then $u_j^{n,t} = 0$ from complementary slackness. By strict complementary slackness (see Bertsimas and Tsitsiklis 1997, p. 192), we have $(\mathbf{v}^{n,t})' \cdot \mathbf{A}^j + u_j^{n,t} > R_j$, i.e., $BP_j(\mathbf{n}, t) = (\mathbf{v}^{n,t})' \cdot \mathbf{A}^j > R_j$, in which case the bid-price policy rejects class j .

- In all optimal LP-solutions $y_j^* = D_j^{t-1}$, in which case $u_j^{n,t} = (R_j - (\mathbf{v}^{n,t})' \cdot \mathbf{A}^j)^+ > 0$, so the bid-price policy accepts class j .

- There is some optimal LP-solution such that $0 < y_j^* < D_j^{t-1}$, which implies that the dual constraint $(\mathbf{v}^{n,t})' \cdot \mathbf{A}^j - u_j^{n,t} \geq R_j$ is binding and $u_j^{n,t} = 0$, so $(\mathbf{v}^{n,t})' \cdot \mathbf{A}^j = R_j$, and thus the bid-price policy accepts class j . \square

PROPOSITION 5 (STRUCTURAL PROPERTIES OF THE CEC POLICY). *Suppose that $LP(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1})$ and $LP(\mathbf{n}, \mathbf{D}^{t-1})$ have different optimal dual bases. Then:*

(a) CEC accepts class j if $y_j^* \geq \min(D_j^{t-1}, 1)$ in some optimal solution of $\text{LP}(\mathbf{n}, \mathbf{D}^{t-1})$.

(b) CEC rejects class j if $y_j^* < \min(D_j^{t-1}, 1)$ in all optimal solutions of $\text{LP}(\mathbf{n}, \mathbf{D}^{t-1})$.

PROOF. If in some optimal solution \mathbf{y}^* of $\text{LP}(\mathbf{n}, \mathbf{D}^{t-1})$ we have that $y_j^* \geq \min(D_j^{t-1}, 1)$, then $\mathbf{y}^* - \min(D_j^{t-1}, 1) \cdot \mathbf{e}_j \geq \mathbf{0}$ is a feasible solution of $\text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}_{(j)}^{t-1})$, where $\mathbf{D}_{(j)}^{t-1} = \mathbf{D}^{t-1} - \min(D_j^{t-1}, 1) \cdot \mathbf{e}_j$, and hence,

$$\begin{aligned} \text{LP}(\mathbf{n}, \mathbf{D}^{t-1}) &= \mathbf{R}' \cdot \mathbf{y}^* \leq \min(D_j^{t-1}, 1)R_j + \text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}_{(j)}^{t-1}) \\ &\leq R_j + \text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1}), \end{aligned}$$

where the last inequality holds because any optimal primal solution of $\text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}_{(j)}^{t-1})$ is feasible to $\text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1})$. So, $\text{OC}_j^{\text{LP}}(\mathbf{n}, \mathbf{D}^t) = \text{LP}(\mathbf{n}, \mathbf{D}^{t-1}) - \text{LP}(\mathbf{n} - \mathbf{A}^j, \mathbf{D}^{t-1}) \leq R_j$. This proves Part (a).

For Part (b), it follows that $0 \leq y_j^* < D_j$ in all primal optimal solutions. By complementary slackness, we have that $u_j^{n,t} = 0$ and $(\mathbf{v}^{n,t})' \cdot \mathbf{A}^j = R_j$. Under the assumption that the optimal dual basis changes, we obtain by Proposition 3 that $\text{OC}_j^{\text{LP}}(\mathbf{n}, t) > \text{BP}_j(\mathbf{n}, t) = (\mathbf{v}^{n,t})' \cdot \mathbf{A}^j = R_j$, that is CEC rejects class j , which concludes the proof. \square

4.1. An Example

For single-leg instances of the NRM problem, the bid pricing and the CEC algorithms are the same. This is because under the CEC algorithm, the opportunity cost estimates are the same (no change of basis occurs in Equation 4). However, this is not true for the general network case. In this section, we provide an example that highlights the differences between the BPC and the CEC policies, and shows instances where each one is suboptimal. Moreover, we explain why the cross-concavity properties that insure the threshold time structure for the optimal single-leg policy cannot be extended to the network case. We will use the same example to exhibit the following situations:

- An instance when BPC accepts, but CEC rejects a given request in the same state;
- An instance when BPC is suboptimal;
- An instance when CEC is suboptimal;
- A counterexample of a cross-concavity property (“decreasing differences”) of the LP and DP-value functions for the NRM problem.

In general, if at a given state the CEC policy accepts a request from itinerary j , then at the same state the bid-pricing control policy will also accept, but not vice versa (see Proposition 3). The following situation may occur: $\text{BP}(j) \leq R_j < \text{OC}^{\text{LP}}(j)$, and so we will accept under the BPC policy but not under the CEC policy. The following is an example of such behavior, that provides insight into the structural properties of the two policies.

Consider a network with four nodes: a hub h , two origin nodes o_1, o_2 , and a destination node d . The legs of the network are (1) (o_1, h) , (2) (o_2, h) , (3) (h, d) . One can think of this as part of a bigger network where the other (connecting) flights have been sold out. Suppose there is demand from the origin nodes o_1, o_2 to the hub node h and to the destination node d , on the itineraries: (1) o_1h , (2) o_2h , (13) o_1hd , (23) o_2hd . Suppose that there is only one fare class per itinerary, and there is no demand from h to d . We assume the pricing structure is such that subitineraries cost less: $R_1 < R_{13}, R_2 < R_{23}$. Furthermore, assume with almost no loss of generality that $R_{13} + R_2 < R_{23} + R_1$ (the other case is symmetric, unless equality holds).

Suppose that the available capacity in the current state t is $\mathbf{n} = (n_1, n_2, n_3)$ and the expected demand for each fare class in the remaining $t - 1$ periods is positive. The static LP and its dual can be formulated as follows:

$$\begin{aligned} \text{LP}(\mathbf{n}, \mathbf{D}) &= \max \quad R_1 y_1 + R_2 y_2 + R_{13} y_{13} + R_{23} y_{23} \\ \text{s.t.} \quad &y_1 + y_{13} \leq n_1 \\ &y_2 + y_{23} \leq n_2 \\ &y_{13} + y_{23} \leq n_3 \\ &\mathbf{0} \leq \mathbf{y} \leq \mathbf{D} \\ &= \min \quad \mathbf{n}'\mathbf{v} + \mathbf{u}'\mathbf{D} \\ \text{s.t.} \quad &v_1 + u_1 \geq R_1 \\ &v_2 + u_2 \geq R_2 \\ &v_1 + v_3 + u_{13} \geq R_{13} \\ &v_2 + v_3 + u_{23} \geq R_{23} \\ &\mathbf{u}, \mathbf{v} \geq \mathbf{0}. \end{aligned}$$

Suppose that there is one seat left on each leg, so $\mathbf{n} = (1, 1, 1)$, and $D_1 > 1$ and $D_{23} > 1$, so that the demand for the high-paying mix is large enough for the corresponding constraints to be nonbinding in an optimal solution. Then the optimal LP solution is $y_1^* = y_{23}^* = 1$, $y_2^* = y_{13}^* = 0$, and the value of the LP is $R_1 + R_{23}$.

Since the demand constraints are nonbinding, we must have that $\mathbf{u} = \mathbf{0}$, and hence $v_1 \geq R_1, v_2 \geq R_2, v_3 \geq \max(0, R_{13} - v_1, R_{23} - v_2)$. Therefore, in an optimal solution, the shadow prices are equal to $v_1^* = R_1, v_2^* = R_2, v_3^* = R_{23} - R_2$. We can compute the bid prices and opportunity costs as follows:

$$\begin{aligned} \text{OC}_1^{\text{LP}} &= \text{LP}(1, 1, 1, \mathbf{D}) - \text{LP}(0, 1, 1, \mathbf{D}) \\ &= R_1 = \text{BP}(1), \\ \text{OC}_2^{\text{LP}} &= \text{LP}(1, 1, 1, \mathbf{D}) - \text{LP}(1, 0, 1, \mathbf{D}) \\ &= R_1 + R_{23} - R_{13} > \text{BP}(2) = R_2, \\ \text{OC}_{13}^{\text{LP}} &= \text{LP}(1, 1, 1, \mathbf{D}) - \text{LP}(0, 1, 0, \mathbf{D}) \\ &= R_1 + R_{23} - R_2 = \text{BP}(13) > R_{13}, \\ \text{OC}_{23}^{\text{LP}} &= \text{LP}(1, 1, 1, \mathbf{D}) - \text{LP}(1, 0, 0, \mathbf{D}) \\ &= R_{23} = \text{BP}(23). \end{aligned}$$

Therefore, the two policies disagree on the acceptance of class (2): Under the BPC policy, we will accept, whereas under the CEC policy, we will reject a class (2) request at state $\mathbf{n} = (1, 1, 1)$ as long as there is sufficient forthcoming demand for classes (1) and (23).

The question is which one of the policies is better? Clearly, in the case when demand is deterministic, the BPC policy is suboptimal, by giving away at time t one unit of capacity that would bring higher revenues in the future. The CEC policy, however, is by definition optimal in the deterministic case because it is equivalent to the DP (certainty equivalence). In the stochastic case, the bid-pricing policy is suboptimal when there is sufficient demand to come from the high-fare mix. Otherwise, we may be better off accepting class (2) right away, in which case our policy is suboptimal.

We can also observe on this example that the LP, and thus the DP value do *not* exhibit a certain type of cross-concavity property called decreasing differences (see Karaesman and van Ryzin 1998):

DEFINITION 1. A function $f: S \subset R^n \rightarrow R$ satisfies decreasing differences on S if for any $\mathbf{s} \in S$, and $i \neq j \in \{1, \dots, n\}$, and for all $\delta_i, \delta_j > 0$ with $\mathbf{s} + \delta_i \mathbf{e}_i, \mathbf{s} + \delta_j \mathbf{e}_j$ and $\mathbf{s} + \delta_i \mathbf{e}_i + \delta_j \mathbf{e}_j \in S$, the following relation holds: $f(\mathbf{s} + \delta_i \mathbf{e}_i + \delta_j \mathbf{e}_j) - f(\mathbf{s} + \delta_i \mathbf{e}_i) \leq f(\mathbf{s} + \delta_j \mathbf{e}_j) - f(\mathbf{s})$.

This cross-concavity property reduces in the univariate case to concavity. This is the key observation underlying the proof of the threshold times property for the optimal single-leg policy (see Diamond and

Stone 1991), and would provide a sufficient condition for the property to extend to the network case.

However, the decreasing differences property is violated in our example:

$$\begin{aligned} \text{LP}(1, 1, 1) - \text{LP}(0, 1, 1) \\ &= \max(R_1 + R_{23}, R_2 + R_{13}) - R_{23} < R_{13} \\ &= \text{LP}(1, 0, 1) - \text{LP}(0, 0, 1). \end{aligned}$$

Notice that the assumption here is that “subitinerary” fares are cheaper ($R_1 < R_{13}, R_2 < R_{23}$). Network effects imply that the opportunity cost of Itinerary (13) under the CEC policy decreases with capacity. Furthermore, if the demand is large enough, the above relation transfers to the corresponding DP values because the two are asymptotically equal. Surprisingly, this says that incremental revenues (opportunity costs) may decrease by decreasing capacity along a certain direction (see also Feng and Lin 2000).

5. Cancellations and Overbooking

The control policies we have discussed so far do not take into account the fact that a significant fraction of customers cancel their reservations during the booking period (cancellations) or simply do not utilize their reservation (no-shows). In these cases customers get full, partial, or no refund, depending on the fare category. In either case, extra capacity becomes available and could be used to accommodate other potential customers. To counterbalance this phenomenon, a common revenue management practice is overbooking. Airlines, hotels, and so on, oversell their inventories to account for reservations that will not materialize. If at the end of the horizon, more demand has materialized than the inventory can accommodate, companies practice class upgrades, pay overbooking penalties to unsatisfied reservations or even perform aircraft changes.

The typical overbooking method practiced by airlines, hotels, and so on, is to decide an initial allocation of overbooking pads, which are virtual increases in leg-capacity. This is usually performed, in practice as well as in the literature, as a static, one-time

decision made at the beginning of the booking horizon. In the following, we propose a new method for dynamic overbooking control, where oversales decisions are dynamic and implicit in the admission control mechanism.

5.1. General Models and Relationships

First, we extend the various models described in §2 to incorporate cancellations, no shows and overbooking.

A Dynamic Programming Model. In the case of perfect state information, this model provides the *optimal control policy*. By allowing cancellations, the state space of the DP-model becomes much larger because it is necessary to keep track of the past sales record \mathbf{s} . The random quantities involved are the demand, cancellations, and no-show processes. Given the initial network inventory \mathbf{N} , the maximum expected revenue (less refunds and penalties) to be collected from state (\mathbf{s}, t) onward (cost-to-go), is given by

$$\begin{aligned}
 DP_{\mathbf{N}}^o(\mathbf{s}, t) &= \sum_j p_j^t \cdot \max(DP_{\mathbf{N}}^o(\mathbf{s}, t-1), \\
 &\quad R_j + DP_{\mathbf{N}}^o(\mathbf{s} + \mathbf{e}_j, t-1)) \\
 &\quad + \sum_{j|s_j \geq 1} p_j^{ct} \cdot (DP_{\mathbf{N}}^o(\mathbf{s} - \mathbf{e}_j, t-1) - R_j^{ct}) \\
 &\quad + p_0^t \cdot DP_{\mathbf{N}}^o(\mathbf{s}, t-1) \\
 DP_{\mathbf{N}}^o(\mathbf{s}, 0) &= \sum_{\tilde{\mathbf{s}}} P(\tilde{\mathbf{s}} \text{ bookings out of } \mathbf{s} \text{ materialize}) \\
 &\quad \cdot DP_{\mathbf{N}}^o(\tilde{\mathbf{s}}, -1) \\
 DP_{\mathbf{N}}^o(\mathbf{s}, -1) &= - \min_{\mathbf{s}^o} C' \cdot \mathbf{s}^o \\
 &\quad \text{s.t. } \mathbf{A} \cdot (\mathbf{s} - \mathbf{s}^o) \leq \mathbf{N} \\
 &\quad 0 \leq \mathbf{s}^o \leq \mathbf{s}. \tag{6}
 \end{aligned}$$

The difference from the basic model is that cancellation events are incorporated in the Bellman Equation (6), and the boundary conditions are changed to account for no-shows (time $t = 0$) and final bumping decisions (time $t = -1$). The final bumping decision is made so as to minimize total penalties, while keeping the actual capacity restrictions satisfied.

If customer-walking penalties are payed per leg $\mathbf{c} = (c_1, \dots, c_l)$, rather than per itinerary, the boundary condition at $t = -1$ is simply $\mathbf{c}'(\mathbf{A}\mathbf{s} - \mathbf{N})^+$. When bumping penalties are itinerary specific (not leg-additive), then an optimization problem needs to be

solved to decide which passengers should be refused boarding so as to incur least penalties. For example, in a two-leg network which is oversold by one seat on each leg, it is better to bump a connecting passenger rather than two different passengers on each leg whenever overbooking penalties are leg-subadditive.

This model does not incorporate secondary effects associated with overbooking (e.g., image damage and loss of goodwill, customer value, demand shifting, and so on).

The Integer and Linear Programming Models. In the case of the NRM problem with cancellations and no shows, we observe that both the final revenue gained from, and capacity occupied by a class j reservation are not deterministic quantities because they depend on cancellations and no-shows which are random processes. To define a model that is consistent with these observations, we define \tilde{R}_j and $\tilde{\mathbf{A}}^j$ to be the expected revenue gained from, and expected capacity occupied by a class j reservation, before the overbooking period. Let p_j^c and p_j^{ns} denote the probability that a given class j reservation is cancelled at some point in the booking period, and respectively does not show up for the flight. We assume that these quantities are independent on the time the reservation was made, and so is the cancellation penalty. We denote the revenue expected (or "adjusted") from booking a class j customer as $\tilde{R}_j = R_j - (1 - p_j^{ns}) \cdot p_j^c \cdot R_j^c - (1 - p_j^c) \cdot p_j^{ns} \cdot R_j^{ns}$, which accounts for potential cancellation and no-shows events and respective refunds R_j^c, R_j^{ns} (but not overbooking penalties). Let $\tilde{\mathbf{A}}^j = (1 - p_j^c) \cdot (1 - p_j^{ns}) \cdot \mathbf{A}^j$ denote the expected capacity occupied by a class j reservation at the end of the horizon. Finally $\tilde{C}_j = (1 - p_j^c) \cdot (1 - p_j^{ns}) \cdot C_j$ is the average overbooking cost of one class j reservation.

With these notations, we can formulate the following integer programming approximation model, that maximizes expected revenues subject to expected capacity constraints:

$$\begin{aligned}
 IP_{\mathbf{N}}^o(\mathbf{s}, t) &= \max_{\mathbf{y}, \mathbf{z}^o} \tilde{\mathbf{R}}' \cdot \mathbf{y} - \mathbf{C}' \cdot \mathbf{z}^o \\
 \text{s.t. } &0 \leq \tilde{\mathbf{A}} \cdot (\mathbf{y} + \mathbf{s}) - \mathbf{A} \cdot \mathbf{z}^o \leq \mathbf{N} \\
 &0 \leq \mathbf{y} \leq \mathbf{D}^t \\
 &\mathbf{y}, \mathbf{z}^o \text{ integer.}
 \end{aligned}$$

The vector \mathbf{y} decides how many requests to be accepted in the future, whereas \mathbf{z}^o determines which

customers (itinerary requests) should be bumped at the end of the horizon, if necessary. The “virtual” overbooking pad for each leg is $\mathbf{A} \cdot \mathbf{z}^o$. The capacity constraint requires that cancellation-adjusted past and future sales, less oversales, should not exceed the initial network capacity.

With the change of variable $s_j^o = z_j^o / (1 - p_j^c) \cdot (1 - p_j^{ns})$, $j = 1, \dots, n$, the corresponding linear programming relaxation, and its dual, can be written as follows:

$$\begin{aligned} \text{LP}_N^o(\mathbf{s}, t) &= \max \quad \tilde{\mathbf{R}}' \cdot \mathbf{y} - \tilde{\mathbf{C}}' \cdot \mathbf{s}^o \\ &\text{s.t.} \quad \tilde{\mathbf{A}} \cdot (\mathbf{s} + \mathbf{y} - \mathbf{s}^o) \leq \mathbf{N} \\ &\quad \mathbf{0} \leq \mathbf{y} \leq \mathbf{D}^t \\ &\quad \mathbf{0} \leq \mathbf{s}^o \leq \mathbf{y} + \mathbf{s} \\ &= \min \quad \mathbf{v}' \cdot \mathbf{N} - \mathbf{u}' \cdot \mathbf{s} + (\tilde{\mathbf{R}}' - \mathbf{u}') \cdot \mathbf{D}^t \\ &\text{s.t.} \quad \mathbf{u}' = \min(\tilde{\mathbf{C}}', \mathbf{v}' \cdot \tilde{\mathbf{A}}) \\ &\quad \mathbf{v} \geq \mathbf{0}. \end{aligned}$$

5.2. Adjusted Policies

Following the spirit of the basic adaptive bid pricing and CEC policies, we adjust these to incorporate cancellations and overbooking.

Adjusted BPC Policy. We modify the DP model so that at any given state \mathbf{S} a class j request is accepted if and only if its “adjusted” fare \tilde{R}_j is higher than either its adjusted bid price, or the adjusted overbooking penalty \tilde{C}_j , i.e.,

$$\tilde{R}_j \geq \min(\tilde{C}_j, \text{BP}_j^o(\mathbf{S})).$$

Here, bid prices are computed as $\text{BP}_j^o(\mathbf{S}) = (\mathbf{v}^{\mathbf{S}})' \cdot \tilde{\mathbf{A}}^j$, where $\mathbf{v}^{\mathbf{S}}$ represent shadow prices for the leg-capacity constraints in $\text{LP}^o(\mathbf{S})$.

Adjusted CEC Policy. Similarly, we can adapt the CEC policy to accept a class j request if and only if its “adjusted” revenue exceeds its “adjusted” opportunity cost. $\tilde{R}_j \geq \text{OC}_j^o(\mathbf{S}) = \text{LP}_N^o(\mathbf{s}, t - 1) - \text{LP}_N^o(\mathbf{s} + \mathbf{e}_j, t - 1)$. Again, adjustment accounts for the fact that capacity and revenue might not be realized or may be overbooked.

The same type of arguments can be used to extend Proposition 1 for the case of cancellations, no shows and overbooking. Moreover, the structural properties proved in §4 are preserved in the overbooking-adjusted policies.

6. Computational Results

In this section, we present computational results that illustrate the relative practical performance of the previously described admission control policies for the NRM problem, under different demand regimes. Our objective is to evaluate the different models and policies proposed, in terms of the following criteria: (1) running time, (2) quality of approximation, and (3) robustness. In addition, we would like to further investigate the effectiveness of several extensions and improvements. In §6.1, we report computational results without cancellations and overbooking, while in §6.2, we allow cancellations and overbooking.

6.1. Computational Performance Without Cancellations and Overbooking

We performed expected value calculations for two-leg instances and simulation runs for larger networks. To have a consistent and fairly accurate base for comparing various policies, these were simulated simultaneously on the same realizations of the demand process. The computations were performed in MATLAB 4.0 on an Intel Pentium II Celeron 450 MHz (128 MB RAM, WinNT 4.0). We restricted our attention to smaller instances with the explicit objective to obtain insights into the behavior of both BPC and CEC.

6.1.1. Models and Assumptions. We considered the following types of networks:

N2. A two-leg network with nodes o, h, d , legs (1) oh (2) hd and capacities $\mathbf{N} = (N_1, N_2)$. The available itineraries are (1) oh , (2) hd , and (3) ohd , with one class per itinerary. The leg-class incidence matrix, together with the fare structure \mathbf{R} , is:

$$\begin{pmatrix} \mathbf{R} \\ \mathbf{A} \end{pmatrix} = \begin{pmatrix} R_1 & R_2 & R_3 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

N3. The three-leg network described in the example of §4.1, with four nodes: a hub h , two origin nodes o_1, o_2 , and a destination node d , and the legs (1) o_1h , (2) o_2h , (3) hd . There is demand from the origin nodes

o_1, o_2 to the hub node h and to the destination node d , on the itineraries: (1) o_1h , (2) o_2h , (13) o_1hd , (23) o_2hd . Suppose that there is only one fare class per itinerary, and there is no demand from h to d . The leg-class incidence matrix, together with the fare structure \mathbf{R} , is:

$$\begin{pmatrix} \mathbf{R} \\ \mathbf{A} \end{pmatrix} = \begin{pmatrix} R_1 & R_2 & R_{13} & R_{23} \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}.$$

N4. A four-leg network, with two origins o_1, o_2 , two destinations d_1, d_2 and a hub h . The legs in the network are (1) o_1h , (2) o_2h , (3) hd_1 , (4) hd_2 , with capacities $\mathbf{N} = (N_1, N_2, N_3, N_4)$. There is demand for all the eight itineraries, with one fare class per itinerary. The leg-class incidence matrix, together with the fare structure \mathbf{R} , is:

$$\begin{pmatrix} \mathbf{R} \\ \mathbf{A} \end{pmatrix} = \begin{pmatrix} R_1 & R_2 & R_3 & R_4 & R_{13} & R_{14} & R_{23} & R_{24} \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}.$$

N4.2. The same as (N4), except there are two fares per itinerary, so 16 demand classes in all. The leg-class incidence matrix, together with a high-low fare structure $\mathbf{R} = (\mathbf{R}^h, \mathbf{R}^l)$, is:

$$\begin{pmatrix} \mathbf{R} \\ \mathbf{A} \end{pmatrix} = \begin{pmatrix} R_1^l & R_1^h & R_2^l & R_2^h & R_3^l & R_3^h & R_4^l & R_4^h & R_{13}^l & R_{13}^h & R_{14}^l & R_{14}^h & R_{23}^l & R_{23}^h & R_{24}^l & R_{24}^h \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

We used the following alternative scenarios to model the arrival process:

HP. Homogeneous Poisson arrivals with arrival rates given by a constant vector \mathbf{p} (i.i.d. Bernoulli trials).

NHL. Nonhomogeneous Poisson with high-low demand. We assume arrival rates increase for high-fare classes ($\mathbf{p}^t = \mathbf{p}/\log_a(a + t \cdot \rho)$) and decrease ($\mathbf{p}^t = \mathbf{p}/\log_a(a + (T - t) \cdot \rho)$) for low-fare classes, as we approach departure.

NLH. Nonhomogeneous Poisson with low-high demand. We assume arrival rates decrease for high-fare classes ($\mathbf{p}^t = \mathbf{p}/\log_a(a + (T - t) \cdot \rho)$) and increase ($\mathbf{p}^t = \mathbf{p}/\log_a(a + t \cdot \rho)$) for low-fare classes, as we approach departure.

The log-factors are given by ρ , and a is a constant, equal to the base of the logarithm (we take $a = 2$).

6.1.2. Running Time. Exact calculations of the optimal expected revenue (DP), and expected values of the proposed policies (CEC, BPC) are practically impossible. Already for two-leg networks (N2) with one class per itinerary, 50 seats per leg initial capacity and $T = 300$ time periods the computation takes about two hours.

A tractable approach for measuring performance of the proposed policies, however, is provided by simulation. A simulation run of the CEC or BPC policies for a two-leg network takes less than a minute (less than a second per iteration). The largest network we simulated was a four-leg network (N4.2) with two demand classes per itinerary. When the initial capacities are in the order of $N = 50-100$ and the time horizon has $T = 300$ time periods, a full simulation run for such an instance takes a few minutes (i.e., a few seconds per iteration).

Table 2 Expected Value Calculation for ($N2$) with $\mathbf{N} = (50, 50)$; $\mathbf{R} = (25, 20, 35)$; $\mathbf{p} = (0.4, 0.3, 0.1)$; $T = 10-200$

T	(HP)				(NLH)				(NHL)			
	LP	DP	CEC	BPC	LP	DP	CEC	BPC	LP	DP	CEC	BPC
10	195	195	195	195	201.4	201.4	201.4	201.4	249.7	249.7	249.7	249.7
20	390	390	390	390	413.5	413.5	413.5	413.5	498.8	498.8	498.8	498.8
30	585	585	585	585	632.9	632.9	632.9	632.9	747.3	747.3	747.3	747.3
40	780	780	780	780	857.6	857.6	857.6	857.6	995.1	995.1	995.1	995.1
50	975	975	975	975	1,086.5	1,086.5	1,086.5	1,086.5	1,242.1	1,242.1	1,242.1	1,242.1
60	1,170	1,170	1,170	1,170	1,318.6	1,318.6	1,318.6	1,318.6	1,488.3	1,488.1	1,488.1	1,488
70	1,365	1,365	1,365	1,365	1,553.4	1,552.2	1,552.1	1,552.2	1,733.6	1,703	1,702.8	1,690.4
80	1,560	1,559.5	1,559.5	1,559.5	1,790.4	1,755.3	1,754.5	1,754.9	1,834.3	1,800.9	1,799.6	1,753.4
90	1,755	1,745.7	1,745.4	1,745.7	1,885.3	1,863.3	1,860.8	1,840.9	1,846.1	1,831.4	1,829.5	1,791.1
100	1,950	1,897.5	1,896.4	1,897.1	1,925.1	1,904.7	1,901.4	1,852.1	1,858.4	1,845.8	1,843.1	1,809.6
110	2,020	1,998.9	1,996.6	1,993.4	1,937.3	1,922	1,917	1,863.3	1,871.1	1,858.7	1,855	1,822.5
120	2,090	2,064.6	2,061.6	2,031.6	1,949.1	1,934.3	1,927.1	1,876.9	1,884.4	1,872.1	1,867.1	1,837
130	2,140	2,110.6	2,107.3	2,021.9	1,960.6	1,945.9	1,936.7	1,889.5	1,898.3	1,886.2	1,879.8	1,851.5
140	2,170	2,145.5	2,140.5	2,018.3	1,971.7	1,957.1	1,946.2	1,899.5	1,913	1,901	1,893.2	1,864
150	2,200	2,175.7	2,167.6	2,018.9	1,982.5	1,968.1	1,955.6	1,907.2	1,928.7	1,916.7	1,907.6	1,875.8
160	2,230	2,202.4	2,192.9	2,019.3	1,993.1	1,978.7	1,964.9	1,913.1	1,945.5	1,933.6	1,923.2	1,887.8
170	2,250	2,222.8	2,216.9	2,019.4	2,003.4	1,989.1	1,974	1,918.1	1,963.8	1,952	1,940.2	1,900.8
180	2,250	2,236.2	2,233	2,019.4	2,013.5	1,999.3	1,983	1,922.9	1,984.2	1,972.5	1,959.3	1,915.5
190	2,250	2,243.8	2,242.3	2,019.4	2,023.5	2,009.3	1,991.8	1,928.6	2,007.4	1,995.8	1,981.4	1,932.4
200	2,250	2,247.5	2,246.9	2,019.4	2,033.2	2,019	2,000.5	1,934.7	2,035.3	2,023.7	2,008.1	1,952.8

Note. The first section considers homogeneous arrivals, and the next two are nonhomogeneous arrivals with $\rho = (50, 50, 30)$ and $a = 2$.

For more realistic airline networks [20 legs (10 origins, 10 destinations, and one hub), 600 classes (100 itineraries, 6 fare classes)] the time per iteration is still in the order of a few minutes. This supports the idea that these dynamic policies can be used in an online fashion to provide an adaptive control mechanism where opportunity cost estimates are rapidly recalculated at each iteration.

6.1.3. Quality of Approximation

Expected Value Computations for Two-Leg Networks. The results in Table 2 show expected value calculations for the case of the two-leg network, with different initial capacities, fare structures, and demand regimes. In all cases, we observe that when the time horizon is large, the value of the CEC policy is near optimal, whereas the bid-pricing policy appears to be off by a constant. When the time horizon is small, both policies perform optimally. For intermediate stages when both some demand and capacity constraints become binding, it is not clear which policy is better. Note that the BPC value is non-monotonic over time. We believe that this is because

of the dependence of the policy on the choice of dual solutions (not enforced in our program).

Simulation Results. For two-leg networks we provide in Table 3 a comparison between simulation results and exact calculations for the DP value function and expected values of the CEC and BPC policies. The high dimensionality of the problem does not allow for exact expected value calculations for larger networks. In Table 4, we present results from

Table 3 Simulations and Expected Value Calculation for Two-Leg Network

T	LP	DP EXP	CEC		BPC	
			EXP	SIM (σ)	EXP	SIM (σ)
($N2$)						
50	975	975	975	973.3 (71)	975	973.2 (71)
100	1,950	1,897.5	1,896.4	1,904.05 (83)	1,897.4	1,903.75 (82)
150	2,200	2,175.5	2,167.6	2,174.9 (45)	2,018.8	2,021.2 (44)
200	2,250	2,247.5	2,246.9	2,246.4 (18)	2,019.4	2,023.3 (17)
300	2,250	2,250	2,250	2,250 (0)	2,200	2,200 (0)

Note. ($N2$): $\mathbf{N} = (50, 50)$; $\mathbf{R} = (25, 20, 35)$; $\mathbf{p} = (0.4, 0.3, 0.1)$.

Table 4 Simulations for Larger Networks (Top-Down)

T	LP	CEC (σ)	BPC (σ)
<i>(N3)</i>			
100	2,300	2,273 (140)	2,268 (140)
200	3,200	3,155 (184)	3,148 (181)
300	3,600	3,508 (178)	3,476 (177)
<i>(N4.1)</i>			
100	3,500	3,462 (150)	3,460 (154)
300	5,800	5,615 (78)	5,635 (74)
600	6,050	5,951 (85)	5,842 (94)
<i>(N4.2)</i>			
100	4,070	4,087 (213)	4,085 (216)
200	6,125	5,914 (241)	5,856 (233)
300	11,110	10,859 (302)	10,846 (309)

Note. (N3): $\mathbf{N} = (60, 60, 100)$; $\mathbf{R} = (10, 20, 50, 30)$ and $\mathbf{p} = (0.2, 0.3, 0.4)$; (N4.1): $N_i = 50$, $\mathbf{R} = (20, 20, 50, 30, 60, 40, 45, 55)$, $\mathbf{p} = (0.2, 0.25, 0.05, 0.05, 0.1, 0.15, 0.1, 0.1)$; (N4.2): $N_i = 50$, $\mathbf{R}' = (20, 20, 50, 30, 60, 40, 45, 55)$, with $\mathbf{R}^h = 2 \cdot \mathbf{R}'$ and $\mathbf{p} = (0.15, 0.02, 0.2, 0.01, 0.05, 0.015, 0.05, 0.01, 0.1, 0.01, 0.15, 0.015, 0.1, 0.01, 0.1, 0.01)$.

simulating various admission control policies for larger networks. For simulation runs, we perform $NT = 100$ trials and average the results to compute an expected value estimate. We also compute the standard deviation σ of the estimation. In all cases, we observe that our CEC policy performs better than the BPC policy, by up to 2%.

6.1.4. Robustness in Estimation. So far, we have assumed that the CEC and BPC policies use correct demand information, in the sense that the mean-demand forecast is unbiased. While this provides a reasonable standard for comparing policies, in reality, the (expected demand) measurements obtained from forecasting tools are seldom exact. We feel that it is of practical value to assess the robustness of the CEC and BPC policies to noise and bias in the demand forecast.

To obtain a qualitative measure of how bias and noise in the demand forecast affect these policies, we assess the impact on the underlying LP-value. Let δ be an m -variate random variable representing the noise in demand forecast for each class. To measure the robustness of the LP-value to bias and

noise in demand forecasts, we compare $LP(\mathbf{n}, \mathbf{D})$ and $E_\delta[LP(\mathbf{n}, \mathbf{D} + \delta)]$. From the dual formulation we have that

$$\begin{aligned} E_\delta[LP(\mathbf{n}, \mathbf{D} + \delta)] - LP(\mathbf{n}, \mathbf{D}) \\ = E_\delta[OC_\delta(\mathbf{n}, \mathbf{D})] \leq (\mathbf{R}' - \mathbf{v}' \cdot \mathbf{A})^+ \cdot E[\delta], \end{aligned}$$

where \mathbf{v} are the leg-shadow prices of $LP(\mathbf{n}, \mathbf{D})$. So there is a sublinear effect on the value function, associated with the forecasting bias of the demand for those classes whose fares exceed their bid prices.

Furthermore, computational experiments show that the CEC policy is surprisingly robust to certain forms of noise and bias in the demand data, much more so than BPC.

Correlated Random Noise. At each point in time t we generate a multivariate random variable $\delta \sim N(\mathbf{b}, \sigma)$, to introduce noise into the arrival forecast. We propose two alternatives for adding noise to the forecast:

- Constant rate noise: Generate δ and add it to the arrival rate at each time;
- Log-rate noise: Generate δ and add $\delta^t = \log(a + \rho \cdot (t - a)/T) \cdot \delta$ to the arrival rate.

From the results in Table 5, we observe that CEC is constantly more robust than BPC to noise and bias. The columns LP_δ compute the LP-values with noisy forecasts: $LP_\delta(\mathbf{n}, \mathbf{D}^t) = LP(\mathbf{n}, \mathbf{D}^t + E[\delta^t])$.

Uncorrelated Random Noise. The goal of this computational exercise is to illustrate and compare the robustness of the BPC and CEC policies, to increasing bias in the demand forecast. We start with an initial (uncorrelated) bias per arrival \mathbf{b} , and in each consecutive experiment subscripted $(\alpha\mathbf{b})$, we increase the bias by a factor of $\alpha = 1, 2, 4$. In the last experiment (δ) , we introduce correlated noise (δ same as above) in the arrival rate estimation. We observe from Table 6 that CEC is constantly more robust than BPC to noise and bias.

6.1.5. Extensions

Rollout Heuristics. We examine how previously computed values of the CEC and BPC policies can be used to provide estimates of opportunity costs, leading to the roll-out heuristics R(CEC) and R(BPC). We

Table 5 Computation with Noisy Forecasts

T	Constant Rate					Log Rate				
	LP	LP _δ	DP	CEC	BPC	LP	LP _δ	DP	CEC	BPC
10	195	213.7	195	195	195	195	221.1	195	195	195
20	390	416.7	390	390	390	390	408.7	390	390	390
30	585	611.9	585	585	585	585	623.3	585	585	585
40	780	784.6	780	780	780	780	765.5	780	780	780
50	975	975.1	975	975	975	975	957.2	975	975	975
60	1,170	1,175.5	1,170	1,170	1,170	1,170	1,123.5	1,170	1,170	1,170
70	1,365	1,366.4	1,365	1,365	1,365	1,365	1,315.3	1,365	1,365	1,365
80	1,560	1,605.6	1,559.5	1,559.5	1,559.5	1,560	1,569	1,559.5	1,559.5	1,559.5
90	1,755	1,770.3	1,745.7	1,745.2	1,745.3	1,755	1,685.5	1,745.7	1,745.2	1,745.1
100	1,950	1,977.8	1,897.5	1,895.8	1,895.4	1,950	1,878.5	1,897.5	1,895.8	1,894.6
110	2,020	2,086.1	1,998.9	1,996.1	1,961.4	2,020	2,008.5	1,998.9	1,996	1,974.7
120	2,090	2,115.2	2,064.6	2,061.7	1,971.7	2,090	2,069.2	2,064.6	2,061.6	1,967.2
130	2,140	2,139.4	2,110.6	2,107.1	1,980.1	2,140	2,105.5	2,110.6	2,106.3	1,981.2
140	2,170	2,172	2,145.5	2,139.6	1,987.7	2,170	2,147.1	2,145.5	2,137.4	1,987.9
150	2,200	2,197.8	2,175.7	2,166.9	1,994.5	2,200	2,168	2,175.7	2,162.5	1,993.5
160	2,230	2,227.8	2,202.4	2,192.8	1,999.5	2,230	2,178.8	2,202.4	2,185.5	2,030.2
170	2,250	2,250	2,222.8	2,216.8	2,000.2	2,250	2,203.2	2,222.8	2,208.4	2,151.1
180	2,250	2,250	2,236.2	2,233	2,000.2	2,250	2,240.5	2,236.2	2,227.3	2,208.9
190	2,250	2,250	2,243.8	2,242.3	2,130.6	2,250	2,235.2	2,243.8	2,240.3	2,223.4
200	2,250	2,250	2,247.5	2,246.9	2,227	2,250	2,248.5	2,247.5	2,246.8	2,230.6

Note. (N2): $\mathbf{N} = (50, 50)$; $\mathbf{R} = (25, 20, 35)$; $\mathbf{p} = (0.4, 0.3, 0.1)$; $\mathbf{b} = (0.07, -0.05, 0.03)$; $\sigma = \frac{1}{5} \begin{pmatrix} 0.1 & -0.02 & 0.04 \\ -0.02 & 0.08 & 0.01 \\ 0.04 & 0.01 & 0.06 \end{pmatrix}$.

observe that the rollout procedure produces a significant improvement in the quality of the value function approximation, especially for the BPC policy (10%). This can be observed in Table 7 where we performed calculations for various two-leg networks.

In practice, however, it is too expensive to compute and store all the H -values. In order to implement this idea effectively, we suggest storing H -values from several insightful a priori simulations, and interpolating these in an online fashion (as needed), for the “second time around” policy. It is an interesting idea to investigate what types of preprocessing simulations would provide an insightful information database.

Simulations Using Monte Carlo Demand Estimation. We run a simplified version of the Monte Carlo policy described in §3.3, that assigns the same weights to all the trials and defines:

$$OC_j^{MC}(\mathbf{n}, t) = \frac{1}{r} \cdot \sum_{i=1}^r OC_j(\mathbf{n}, \widehat{\mathbf{D}}_i^{t-1}).$$

When the number of trials r is large enough (so we have a reasonably good MC-demand estimate), this policy provides a visible improvement (1.6%) over the original version, as can be observed from the following computational examples. However, when the number of trials is not large enough (30 or less), we have observed that this policy delivers a poor performance. The tables below provide estimates for the value of different models (LP, PI^{LP}, MC, CEC, BPC), as well as standard deviations, lower and upper bounds obtained through simulation. We also calculate the estimated average ratio P between different models, as well as the standard deviation, minimum, and maximum ratio observed during simulation. One can observe for instance that the ratio CEC/BPC is always ≥ 1 , whereas the ratio MC/PI^{LP} is surprisingly high, ranging between 96%–99%. The computational results presented in Table 8 show an improvement of up to 20% of MC over CEC.

Table 6 Increasing Bias in Arrival Rates

T	DP	LP	CEC	BPC	LP_b	CEC_b	BPC_b	LP_{2b}	CEC_{2b}	BPC_{2b}
10	190	190	190	190	196.227	190	190	202.454	190	190
20	380	380	380	380	387.9062	380	380	395.8123	380	380
30	568.0307	570	567.9671	568.0053	578.9176	567.958	568.0149	587.8351	567.9583	568.0271
40	713.9429	740	712.8882	711.4275	751.7866	712.6471	711.3734	763.5732	712.654	710.6125
50	784.8309	815	782.7964	745.0979	818.9707	782.1006	743.3085	822.9414	781.6062	737.3545
60	819.8745	845	817.9585	749.6329	849.1515	817.5189	748.0799	853.303	816.9229	742.664
70	841.6218	855	839.8851	749.8441	855	839.6463	748.1471	855	839.3091	742.664
80	851.0686	855	850.2119	749.8441	855	850.1653	748.1471	855	850.1115	742.664
90	854.0811	855	853.8794	749.8441	855	853.8777	748.1471	855	853.8765	742.664
99	854.7912	855	854.7523	749.8441	855	854.7523	748.1471	855	854.7523	742.664

T	DP	LP_{4b}	CEC_{4b}	BPC_{4b}	LP_{δ}	CEC_{δ}	BPC_{δ}
10	190	214.9079	190	190	191.8341	190	190
20	380	411.6247	380	380	361.348	380	380
30	568.0307	605.6702	567.9585	568.0273	671.4161	567.9394	567.8188
40	713.9429	787.1465	712.5613	706.1965	707.4845	712.1564	708.9488
50	784.8309	830.8828	780.0541	723.385	801.1911	779.7808	752.618
60	819.8745	855	815.0074	728.399	831.9429	815.4683	757.3472
70	841.6218	855	838.5854	728.399	855	838.5076	757.5589
80	851.0686	855	850.0349	728.399	855	849.9272	757.804
90	854.0811	855	853.8752	728.399	855	853.8605	761.142
99	854.7912	855	854.7523	728.399	855	854.7513	763.8937

Note. (N2): Data: $\mathbf{N} = (19, 19)$; $\mathbf{R} = (25, 20, 35)$; $\mathbf{p} = (0.3, 0.4, 0.1)$, $\mathbf{b} = (0.07, -0.05, 0.03)$.

6.2. Computational Performance Under Cancellations and Overbooking

The objective in this section is to understand the relative performance of BPC and CEC in an environment with cancellations and overbooking. We considered a booking horizon of 15 periods for a hub and spoke network with five cities and two classes, of the type (N4.2). The arrival process for the highest fare class is nonhomogeneous Poisson with rate 0.5 for Periods 1–13 and 5 for Periods 14, 15. The arrival process for the lowest fare class is homogeneous Poisson with Rate 3. For simplicity, we kept the fare of the higher class in a single-leg itinerary equal to \$100 and of the lower class equal to \$80. We varied the fare of two-leg itineraries. After experimentation we identified the following parameters that affected the relative performance of CEC and BPC: We varied the following parameters: (a) The overbooking penalty (C_j), (b) the probability of cancellation (p_c), and (c) the fare of a two-leg itinerary compared to a single-leg itinerary.

The implementation was done in C and the algorithms were run on a Dell Pentium III 600MHz operating under Linux.

In Table 9, we report the behavior of CEC and BPC as a function of the overbooking penalty. We observe that CEC leads to consistently higher revenue by approximately 1%.

In Table 10, we report the behavior of CEC and BPC as a function of the cancellation probability. With the exception of very high cancellation rate (0.30), CEC outperforms BPC.

In Table 11, we report the behavior of CEC and BPC as a function of the ratio ρ of the fare of a two-leg itinerary versus the fare of a single-leg itinerary. CEC outperforms BPC, but the level of overperformance decreases as ρ ranges from 1 to 2.

7. Conclusions

In this paper, we have presented several models and algorithms for solving the stochastic and dynamic

Table 7 **Rollout Heuristics**

<i>T</i>	LP	DP	CEC	BPC	R(CEC)	R(BPC)
<i>(N2.1)</i>						
1	19.5	19.5	19.5	19.5	19.5	19.5
10	195	195	195	195	195	195
20	390	390	390	390	390	390
30	585	585	585	585	585	585
40	780	780	780	780	780	780
50	975	975	975	975	975	975
60	1,170	1,170	1,170	1,170	1,170	1,170
70	1,365	1,365	1,365	1,365	1,365	1,365
80	1,560	1,559.5	1,559.5	1,559.5	1,559.5	1,559.5
90	1,755	1,745.7	1,745.4	1,745.7	1,745.6	1,745.7
100	1,950	1,897.5	1,896.4	1,897.1	1,897.2	1,897.4
110	2,020	1,998.9	1,996.6	1,993.4	1,998	1,995.5
120	2,090	2,064.6	2,061.6	2,031.6	2,063.2	2,042.2
130	2,140	2,110.6	2,107.3	2,021.9	2,108.9	2,052.8
140	2,170	2,145.5	2,140.5	2,018.3	2,143	2,104.2
150	2,200	2,175.7	2,167.6	2,018.9	2,172.9	2,163.9
160	2,230	2,202.4	2,192.9	2,019.3	2,199.6	2,197.2
170	2,250	2,222.8	2,216.9	2,019.4	2,220.6	2,217.4
180	2,250	2,236.2	2,233	2,019.4	2,234.8	2,230.2
190	2,250	2,243.8	2,242.3	2,019.4	2,243.1	2,238.1
200	2,250	2,247.5	2,246.9	2,019.4	2,247.2	2,241.9
<i>(N2.2)</i>						
1	19	19	19	19	19	19
10	190	190	190	190	190	190
20	380	380	380	380	380	380
30	570	568.0307	567.9671	568.0053	568.0285	568.0258
40	740	713.9429	712.8882	711.4275	713.8018	713.2517
50	815	784.8309	782.7964	745.0979	784.3939	774.4301
60	845	819.8745	817.9585	749.6329	819.2368	815.8581
70	855	841.6218	839.8851	749.8441	840.6865	837.072
80	855	851.0686	850.2119	749.8441	850.6009	846.4442
90	855	854.0811	853.8794	749.8441	853.9772	848.7245
100	855	854.8245	854.7925	749.8441	854.8099	849.8084

Note. (*N2.1*): $\mathbf{N} = (50, 50)$, $\mathbf{R} = (25, 20, 35)$, $\mathbf{p} = (0.4, 0.3, 0.1)$; (*N2.2*): $\mathbf{N} = (19, 19)$, $\mathbf{R} = (25, 20, 35)$, $\mathbf{p} = (0.3, 0.4, 0.1)$.

NRM problem. We proposed a new efficient algorithm, based on a certainty equivalent approximation and compared it with the widely used bid-price control policy. This policy conceptually improves the current NRM-approach based on additive bid pricing, by using more insightful, piecewise linear approximations of opportunity cost. It is just as easy to compute, but as opposed to BPC, it is more “robust” in the following sense:

- The CEC opportunity cost estimates are uniquely defined at each state, whereas for additive bid prices, there may be several dual optimal solutions.
- The CEC policy is optimal in the deterministic regime, whereas BPC is not.
- Computationally we observe that the CEC policy outperforms the BPC policy when the load factors (the ratio of expected demand to available capacity) tend to be large. When the load factor is small, both policies perform optimally. There is a critical range

Table 8

(N2.1)	LP	PI ^{LP}	MC	CEC	BPC	P	AvgP	StdP	MinP	MaxP
EXP	2,200	2,191.2	2,170	2,162.8	2,161.6	CEC/BPC	1.0006	0.0035	0.9931	1.0074
Std	0	44.48	50.45	52.41	54.19	CEC/PI ^{LP}	0.987	0.0093	0.9569	1
LB	2,200	2,070	2,035	2,000	2,000	MC/PI ^{LP}	0.9903	0.0084	0.9645	1
UB	2,200	2,250	2,250	2,250	2,240	MC/CEC	1.003	0.0067	0.9909	1.0225
(N2.2)	LP	PI ^{LP}	MC	CEC	BPC	P	AvgP	StdP	MinP	MaxP
EXP	225	215	208	205	202.5	CEC/BPC	1.01	0.028	1	1.075
Std	0	10.55	16.53	18.7	19.76	CEC/PI ^{LP}	0.9521	0.0498	0.8537	1
LB	225	195	185	175	175	MC/PI ^{LP}	0.9664	0.038	0.9024	1
UB	225	225	225	225	225	MC/CEC	1.016	0.034	1	1.0857
(N2.3)	LP	PI ^{LP}	MC	CEC	BPC	P	AvgP	StdP	MinP	MaxP
EXP	8,000	7,922	7,785	7,771	7,766	CEC/BPC	1.0007	0.011	0.975	1.02
Std	0	224.3	238.7	285.2	283.8	CEC/PI ^{LP}	0.9808	0.0166	0.94	1
LB	8,000	7,500	7,200	7,050	7,050	MC/PI ^{LP}	0.9828	0.017	0.939	1
UB	8,000	8,600	8,400	8,500	8,400	MC/CEC	1.002	0.0184	0.9615	1.0425
(N4)	LP	PI ^{LP}	MC	CEC	BPC	P	AvgP	StdP	MinP	MaxP
EXP	6,800	n.a.	6,694.4	6,588	6,568.8	CEC/BPC	1.0029	0.0051	0.9935	1.0117
Std	0	n.a.	214.44	213.73	215	MC/CEC	1.01624	0.0127	0.9908	1.0421
LB	6,800	n.a.	6,260	6,140	6,160					
UB	6,800	n.a.	7,065	6,950	6,930					

Note. (N2): $\mathbf{N} = (50, 50)$, $\mathbf{R} = (25, 20, 35)$, $\mathbf{p} = (0.3, 0.4, 0.1)$, $T = 150$, $NT = 100$, $r = 50$; (N2.1): $\mathbf{N} = (5, 5)$, $\mathbf{R} = (25, 20, 35)$, $\mathbf{p} = (0.3, 0.4, 0.1)$, $T = 20$, $NT = 10$, $r = 50$; (N2.3): $\mathbf{N} = (20, 20)$, $\mathbf{R} = (25, 20, 35)$, $\mathbf{p} = (0.2, 0.3, 0.5)$, $T = 50$, $NT = 50$, $r = 50$; (N4): $\mathbf{N} = (50, 50, 100, 100)$, $\mathbf{R} = (20, 20, 50, 30, 60, 40, 45, 55)$, $\mathbf{p} = (0.15, 0.15, 0.15, 0.15, 0.1, 0.1, 0.1, 0.1)$, $T = 200$, $NT = 50$, $r = 60$.

when the load factor is close to one, where it is not clear which policy provides a better performance, but the difference between the two is very small. Moreover, computational exercises show that the value of the CEC policy is very close to the value function DP, and to the perfect information upper bound PI^{LP}.

- The CEC outperforms the BPC policy in an environment where cancellations and overbooking are present under various scenarios of cancellation probabilities, disparity in fares and overbooking penalties.

- The CEC policy is significantly more robust than BPC to bias and (correlated) noise in the demand forecast.

The following extensions provide insights towards further developments:

- The rollout procedure produces an important improvement in the value of both policies, virtually closing the optimality gap.

- The Monte Carlo simulation procedure incorporates demand variability, and produces a significant improvement in the value of the CEC policy when the number of trials is sufficiently large.

Table 9 The Expected Revenue in 200 Simulation Runs as a Function of the Overbooking Penalty

C_j	100	110	120	130
CEC	24,480	24,450	23,975	23,890
BPC	22,355	22,270	21,920	21,567

Note. The cancellation probability was 0.01 and the fare of a two-leg itinerary was equal to the fare of a single-leg itinerary.

Table 10 The Expected Revenue in 200 Simulation Runs as a Function of the Overbooking Penalty

p_c	0.05	0.10	0.20	0.30
CEC	22,382	21,095	18,760	11,280
BPC	19,680	20,160	17,262	12,960

Note. The overbooking penalty was \$130 and the fare of a two-leg itinerary was equal to the fare of a single-leg itinerary.

Table 11 The Expected Revenue in 200 Simulation Runs as a Function of the Ratio ρ of the Fare of a Two-Leg Itinerary Versus the Fare of a Single-Leg Itinerary

ρ	1	1.5	2
CEC	22,522	23,400	24,457
BPC	20,587	21,437	24,187

Note. The cancellation probability was 0.1 and the overbooking penalty \$130.

Acknowledgments

The authors thank Eduardo Messmacher for implementing the algorithms with cancellation and overbooking. They are also grateful to two anonymous referees for many useful suggestions.

References

- Belobaba, P. 1987. Air travel demand and airline seat inventory management. Ph.D. thesis, MIT, Cambridge, MA.
- Bertsekas, D. 1995. *Dynamic Programming and Optimal Control*, Vol. I. Athena Scientific, Belmont, MA.
- , D. Castañón. 1999. Rollout algorithms for stochastic scheduling problems. *J. Heuristics* 5 89–108.
- , J. Tsitsiklis. 1998. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.
- , ——, C. Wu. 1997. Rollout algorithms for combinatorial optimization. *J. Heuristics* 3 245–262.
- Bertsimas, D., J. Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA.
- , C. Teo, R. Vohra. 1999. Greedy, randomized and approximate dynamic programming algorithms for facility location problems. Working paper, <http://www.fba.nus.edu.sg/qm/wkpaper/Index98.html>.
- Bitran, G., S. Mondschein. 1995. An application of yield management to the hotel industry considering multiple day stays. *Oper. Res.* 43(3) 427–443.
- Brumelle, S., J. McGill. 1993. Airline seats allocation with multiple nested fare classes. *Oper. Res.* 41 127–137.
- Chen, D. 1998. Network flows in hotel yield management. Technical report TR1225, Cornell University, School of Operations Research and Industrial Engineering, <http://citeseer.nj.nec.com/363544.html>.
- Chen, V., D. Gunther, E. Johnson. 2000. An approach to airline yield management based on a Markov decision problem. Working paper, <http://citeseer.nj.nec.com/chen98markov.html>.
- Chi, Z. 1995. Airline yield management in a dynamic network environment. Ph.D. thesis, MIT, Operations Research Center, Cambridge, MA.
- Cook, T. 1998. Sabre soars. *ORMS Today* 6 27–31.
- Curry, R. 1989. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Sci.* 24(3) 193–204.
- . 1992. Real-time revenue management: Bid price strategies for origins/destinations and legs. *Scorecard* 2(Q).
- Diamond, M., R. Stone. 1991. Dynamic yield management on a single leg. Unpublished manuscript, Northwest Airlines.
- Feng, Y., P. Lin. 2000. Optimal booking control on a two-leg flight with multiple fares. Working paper, <http://www.math.nus.edu.sg/~matlinp/www/fenglinor.pdf>.
- Fuchs, S. 1987. Managing the seat auction. *Airline Bus.* 7 40–44.
- Glover, F., R. Glover, J. Lorenzo, C. McMillan. 1982. The passenger-mix problem in scheduled airlines. *Interfaces* 12 73–79.
- Gunther, D., V. Chen, E. Johnson. 2000. Airline yield management: Optimal bid prices for single-hub problems without cancellations. Working paper, <http://gunther.smeal.psu.edu/13888.html>.
- Karaesman, I., G. van Ryzin. 1998. Overbooking with substitutable inventory classes. Working paper, <http://www-1.gsb.columbia.edu/divisions/dro/working%20papers/DRO-2002-11.pdf>.
- Lee, T., M. Hersh. 1993. A model for airline seat inventory control with multiple seat bookings. *Transportation Sci.* 27 1252–1265.
- Littlewood, K. 1972. Forecasting and control of passengers. *Proc. AGIFORS 12th Annual Sympos.*, 95–128.
- McGill, J., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* 33 233–256.
- Popescu, I. 1999. Applications of optimization in probability, finance and revenue management. Ph.D. thesis, Applied Mathematics Department and Operations Research Center, MIT, Cambridge, MA.
- Robinson, L. 1995. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Oper. Res.* 43(2) 252–263.
- Rothstein, M. 1971. An airline overbooking model. *Transportation Sci.* 5 180–192.
- . 1974. Hotel overbooking as a Markovian sequential decision process. *Decision Sci.* 5 389–404.
- Simpson, R. 1989. Using network flow techniques to find shadow prices for market and seat inventory control. Memorandum M89-1, MIT Flight Transportation Laboratory, Cambridge, MA.
- Subramanian, S. S. J., C. Lautenbacher. 1999. Airline yield management with overbooking, cancellations and no-shows. *Transportation Sci.* 33 147–167.
- Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Sci.* 44 1577–1593.
- Williamson, E. 1992. Airline network seat control. Ph.D. thesis, MIT, Cambridge, MA.
- Wollmer, R. 1992. An airline seat management model for a single leg route when lower fare classes book first. *Oper. Res.* 40 26–37.

Received: October 2000; revision received: September 2001; accepted: January 2002.

Copyright 2003, by INFORMS, all rights reserved. Copyright of Transportation Science is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.