Revenue Management in a Dynamic Network Environment

Dimitris Bertsimas * and Ioana Popescu †

September, 2000

Abstract

We investigate dynamic policies for allocating inventory to correlated, stochastic demand for multiple classes, in a network environment so as to maximize total expected revenues. Typical applications include airline networks or sequential reservations for a hotel or car rental service. We propose and analyze both theoretically and computationally a new algorithm, based on approximate dynamic programming, which provides structural insights into the optimal policy by using adaptive, non-additive 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 overbooking 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.

*Boeing Professor of Operations Research, Sloan School of Management, E53-363, Massachusetts Institute of Technology, Cambridge, MA 02139. dbertsim@mit.edu. Research partially supported by NSF grant DMI-9610486 and the Singapore-MIT alliance.

†Assistant Professor of Decision Sciences, INSEAD, Fontainebleau Cedex 77305, FRANCE. ioana.popescu@insead.fr.

1 Introduction

Capacity constrained service industries, such as transportation, tourism, entertainment, media and internet providers are constantly faced with the problem of intelligently allocating their available 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 [21]. For example, in 1997, American Airlines collected one billion dollars by implementing revenue management, representing most of the company's profit [15].

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 $N = (N_1, \ldots, N_l)$ of leg capacities. The network is described by a $l \times m$ matrix A and a m-vector $R = (R_1, \ldots, R_m)$ defined as follows: $a_{ij} = k_j \neq 0$ if leg i is part of bundle (e.g. itinerary) j in a fare class for groups of size k_j; R_j is the fare category of class j. In this way, a demand class j is defined by its itinerary A^j (a column of matrix A) and its fare category R_j. Allowing the matrix A to contain repeated columns for each fare class, and to be integer multiples of the itinerary-incidence vector, we can account for special fares for group requests.

For example, consider a very simple network corresponding to a weekend in a hotel: there are 3 nodes Fri, Sat, Sun and $l = 2$ legs (1) $Fri - Sat$ and (2) $Sat - Sun$ with total capacities $N = (N_1, N_2)$. 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 R, is:

$$\left(\begin{array}{c} R \\ A \end{array} \right) = \left(\begin{array}{ccccccc} R_1^h & R_1^l & R_2^h & R_2^l & R_3^h & R_3^l & R_1^{(10)} \\ 1 & 1 & 0 & 0 & 1 & 1 & 10 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 \end{array} \right).$$

We assume a finite booking horizon of length T, with the time line sufficiently discretized so as to allow at most one request 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 booking period. We also let time $t = -1$ account for the time when final overbooking decisions are being made. 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).

At each point in time $t > 0$ a class j request can arrive with a certain probability p_j^t or a class j reservation may be cancelled with probability p_j^{ct}. At time 0, the no-shows are being counted: we assume that each reservation has a no-show probability p_j^{ns} that is independent of the time the reservation was made. Customers who do not consume their reservations get full, partial or no refund, depending on the fare category and time of cancellation. If at the end of the horizon more demand has materialized than the inventory can accommodate, overbooking decisions are being made. These can be class upgrades or customer bumping, in which case companies pay overbooking penalties.

The demand process at time t is denoted by $\mathcal{D}^t$, and $\mathcal{D}_c^t$ represents its corresponding aggregate distribution. That is, $\mathcal{D}_c^t(j)$ is a random variable representing the number of class j requests to come from time t until departure (order doesn't matter). Usually, we have partial information about the demand process, which might consist of the expected demand to come $D^t = E[\mathcal{D}_c^t]$, and possibly other type of information available from forecasting tools, such as cancellation or no-show probabilities.

The state of the system at any given moment in time t (t periods to departure), is given by the sales-to-date record $S = (s_1, \ldots, s_m)$ for each demand class. If cancellations and no-shows are not allowed, then it is sufficient to define the state of the network by the remaining inventory $n = (n_1, \ldots, n_l)$. Alternatively, 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 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 time-to-go t, the state of the network (S, or n), partial information from demand forecasts $\mathcal{F}^t$, as well as the currently requested fare R_j. If we accept the request, the new state at time $t-1$ becomes $S + k_j \cdot e_j$, where k_j is the size of the class j group request, and e_j is the jth unit vector (when cancellations are not allowed, the state of the network n becomes simply $n - A^j$). If the request is rejected the state of the network does not change.

In reality, the control policy has an impact on the demand process. 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.

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);

$N =$ total initial network capacity (l-vector);

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

$a_{ij} = k_j \neq 0$ iff leg i is part of the itinerary j, in a fare class for groups of size k_j.

Sales and inventory:

$n =$ remaining inventory (l-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);

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

$D^t = E[\mathcal{D}_c^t]$ = 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 also use the operator $(x)^+ = \max(x, 0)$, for $x \in \mathcal{R}$, which naturally extends for vectors: $(x)^+ = ((x_1)^+, \ldots, (x_n)^+)$ for $x = (x_1, \ldots, x_n) \in \mathcal{R}^n$.

Contributions

In this paper we evaluate from different perspectives several policies for solving the dynamic and stochastic NRM Problem. In order 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 Section 3.2). 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 pre-defined additive structure. In particular, they do not account for changes of dual basis due to large group and multi-leg 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 it 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 methods, in which the cost-to-go function is approximated by the solution value of a linear programming (LP) relaxation. In the case of linear networks[1], or origin-destination (OD) demands, we remark that this LP is equivalent to a network flow problem.

2. We provide structural properties that compare the behavior of the proposed CEC policy with the additive bid-price approach. These results offer insight in 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 decisions 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.

[1] A linear network is one whose nodes can be ordered so that the arcs are pairs of consecutive nodes $(i, i+1),\ i = 1, \ldots, l$.

Structure

The remainder of the paper is organized as follows: In the next section we present an overview of the literature. Section 3 describes several (dynamic, stochastic, linear and network flow) models and formulations for the NRM Problem and evaluates the relationships between them. Section 4 presents efficient algorithms for the NRM problem based on ideas from approximate dynamic programming. In Section 5, we present structural properties for the proposed CEC policy and contrast them with additive bid-pricing. In Section 6, we extend our model and algorithms to handle cancellations, no shows and overbookings. Finally, in Section 7, we present computational results. The last section summarizes our conclusions.

2 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). We propose here a structured classification of the relevant literature, by the following main coordinates:

- Single Leg vs. Network Models.
- Static vs. Dynamic Controls (or Decision Rules).

Static decision models are those that function at the strategic level, by deciding in advance the maximum number of reservations to accept on each market segment, and setting protection levels to save seats for late-coming high-paying customers. The same rule applied repetitively with updated forecasts is still considered static. Dynamic rules are decisions to be taken at the operational level: at each point in time, the controller decides whether to accept or reject a current request, given the current sales record, the congestion on the network and the future potential demand.

Static Single Leg. Most of the work in the literature so far addresses the single-leg model. The usual static approach consists in assuming customers booking in increasing order of fare classes, and computing protection levels for high-fare customers. Initial developments are due to Littlewood in 1972 [31], followed by Simpson [40] and Belobaba [2], who proposed a suboptimal policy for computing protection levels based on expected marginal seat revenues (EMSR). Curry [16], Wollmer [50] and Brumelle and McGill [9], derive the optimal solution for the single leg static model. Robinson [36] proposes an extension that handles non-monotonic fare classes.

Dynamic Single-Leg. The principle behind the dynamic version of the single-leg model consists in determining an optimal on-line control policy based on a dynamic programming model, as opposed to allocating a certain amount of seats per demand class. A characterization of the optimal policy, based on a *threshold time property* is due to Diamond and Stone in 1991 [18], and later Feng and Gallego in 1995 [20]. An analogous discrete time solution, that also handles group reservations, is provided by Lee and Hersh [28], who also provide a method for estimating arrival rates. Subramanian et al. [41] 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.

Static Network. The allocation problem is far more complex over the entire network than on individual routes. Several static models have been proposed. Buhr [10] investigates a two sector flight model. A static network flow model with deterministic demands was proposed by Glover et al. [22]. In the same context, a rolling horizon model is suggested by Dror et al. [19]. Their formulation is very complex, allowing for cancellations and overbooking, but no solution is suggested. The case of random demands has been addressed by Wang [45] and Wollmer [50]. Both use very complex formulations based on EMSR.

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 in 1989 [40], and further analyzed by Williamson in 1992 [47] 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 in 1992 [17] 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).

Dynamic Network. The most realistic and relevant, yet least investigated model for NRM is the dynamic network model. A segmentation (or decomposition) model for a single class is proposed by Wong et al. in 1993 [51]. These authors propose an EMSR-based flexible assignment approach, that works well for the 2-leg case, but becomes very cumbersome for higher dimension networks. Alstrup et. al. in 1986 [1] have a similar approach that handles multiple classes. Notably, none of these methods provides an efficient control policy. In the context of hotels, Bitran and Mondschein in 1995 [8] propose a dynamic policy that extends to multiple-night stays, but do not give any further analysis.

Talluri and van Ryzin [42] study a dynamic network model using bid-price control mechanisms and argue why bid-price policies are not optimal in general, and provide an asymptotic regime when certain bid-price controls, based on a probabilistic programming formulation of the problem, are asymptotically optimal. Their model does not consider cancellations and overbooking as part of the dynamic allocation decision.

Chen, Gunther and Johnson [13] formulate the problem as a Markov decision problem, and use linear programming and regression splines to approximate the value function. Gunther, Chen and Johnson [23] introduce a new method to compute bid price for single hub airline networks. Both studies report encouraging simulation results.

All previous studies do not study cancellations.

Related Problems. It is important to acknowledge that inventory control is only one part of the general NRM problem, which has direct interconnections with several other classes of problems. These include pricing, forecasting, network planning, aircraft and flight frequency assignment, crew scheduling etc. Pricing decisions have a definite impact on demand, and demand forecasting is an essential input to any inventory control policy. In turn, network congestion and capacity constraints influence the pricing structure. These inherent interactions suggest that solving each of these problems separately will only lead to suboptimal solutions. However, the complexity of each of these individual problems has hindered advances in the literature so far.

For further pointers to related literature, we refer the interested reader to the tables below, and the survey of McGill and van Ryzin [33].

Relative Positioning. Our approach can be viewed as extending the single-leg models investigated by Lee and Hersh [28], and Bitran and Mondschein [8], who actually propose a similar LP-based heuristic for the multiple-night case, but without any further analysis. Our overbooking model is similar to the early single-leg DP formulation of Rothstein [37], [38]. As far as static network models are concerned, our LP formulation is similar to the one proposed by Williamson [47], but we handle multiple classes and group bookings. The network flow formulation we provide for linear networks and ODF demands, is the same as those proposed by Glover et al [22] for airlines, and Chen [12] for hotels.

Schematically, the relevant literature is summarized as follows:

	Single Leg	Network
Static	**Two Class Model** Littlewood [31] Titze & Greisshaber [43]	**Two Sector Flights** Buhr [10]
	Monotone Multi-Class (Nested) **EMSR:** Simpson [40], Belobaba [2],[3] **Optimal:** Brumelle & McGill [9] Curry [16], Wollmer [50]	**Network Partition** Wang [45], Wollmer [49], [48] Glover et. al [22], D. Chen [12] Dror, Trudeau & Ladany [19]
	Non-monotone Multi-Class Robinson [36] **Adaptive P-Levels** van Ryzin & McGill [44]	**Bid-Pricing** **Virtual nesting:** Curry [17] Simpson [40] **LP vs. PNLP:** Williamson [47]
Dynamic	**Optimal: Threshold Times** Diamond & Stone [18] Feng & Gallego [20] Lee & Hersh [28], Chi [14]	**EMSR w/Segmentation** **Single Class:** Wong et al. [51] **Multiple:** Alstrup et al. [1]
	Other DP Models Chatwin [11], Janakiram et al. [24]	**Additive Bid-Pricing** Talluri & van Ryzin [42], Gunther et al. [23] **DP-approximations:** Chen et. al. [13], **OURS**
Survey	Weatherford & Bodily[46], van Ryzin & McGill [33]	

Overbooking		
	Single Leg	**Network**
Static	Schlifer & Vardi [39]	Dror [19] Karaesman & van Ryzin [25]
Dynamic	Chatwin [11], Rothstein [37], [38] Subramanian et al. [41]	**OURS**

Other Industries	
Hotels	Bitran & Mondschein [8] , Bitran & Gilbert [7] , Chen [12] Ladany [27] , Kimes [26] , Liberman & Yechiali [30]
Car Rental	Karaesman & van Ryzin [25]
Maritime	Maragos [32]
Internet	Paschalidis & Tsitsiklis [34], Leida [29]

3 General Models and Relationships

The problem of dynamic inventory control for network revenue management (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 or overbooking are not permitted; this case is discussed in detail in Section 6.

3.1 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 are defined by the current available capacity vector 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. In general this is not known exactly, thus the imperfect information DP. We define $DP(n,t)$ to be the maximum expected revenue to be collected from state (n,t). This can be computed via the Bellman equation as follows:

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

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

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

We define the opportunity cost $\mathrm{OC}_j(n,t)$ to be

$$
\mathrm{OC}_j(n,t) = DP(n,t-1) - DP(n-A^j,t-1).
$$

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:

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

One can easily show that the value function is nondecreasing in n and t. Moreover, for the single leg case, the value function is concave and opportunity costs decrease with t and n (see Chi [14] or Diamond and Stone [18]). Based on these properties one can show that the optimal policy is characterized by threshold times[2]. In the general network case, however, this is not true, as we will show in an example in Section 5.1.

The exact computation of the opportunity cost suffers from the curse of dimensionality, requires exponential time, and thus cannot be used for practical purposes.

3.2 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 $D^t = E[\mathcal{D}_c^t]$ over the remaining t periods (that is $\mathcal{F}^t$ is the class of processes with aggregate expected value given by D^t). Alternatively, one can think of this model in terms of demand being a deterministic quantity. The integer programming model computes the optimal allocation y^* of available inventory n to the expected itinerary demand D^t, by maximizing total revenues subject to capacity and itinerary demand constraints. For all $n \le N$ and $t \le T$ we have:

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

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

[2] Threshold times are points in time on the booking horizon before which requests are rejected, and after which requests are accepted.

integrality constraint:

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

Clearly, we have that $IP(n, D^t) \leq LP(n, D^t)$. 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:

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

This formulation can be equivalently written as:

$$
LP(n, D^t) = \min_{v \geq 0} \; v' \cdot n + (R' - v' \cdot A)^+ \cdot D^t \; = \min_{k \in K} \; v_k' \cdot n + u_k' \cdot D^t \,,
$$

where K denotes the index set of extreme points (v_k, u_k) of the dual polyhedron.

From the last formulation one can see that the objective value of Model (2) is a **piecewise linear, concave and nondecreasing** function of the expected demand to come D^t and available capacity n.

3.3 Origin-Destination RM and Linear Networks.

Consider a particular case of the NRM problem, where demand is by Origin-Destination-Fare (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, etc.). 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 re-optimization is very fast. For airlines, this model provides optimal (for the airline) 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^f_{od}. The revenue collected along arc $(o,d)_f$ is R^f_{od}. For each "leg" (i,j) (flight leg, resp. 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. [22] in the context of airlines, and for hotel revenue management by Chen [12], who proves that it reduces to solving a network flow problem.

Proposition 1 *(Integrality of solution) For ODF models with integer data, the optimal LP-solution is integral, that is $IP(n, D^t) = LP(n, D^t)$.*

3.4 The Perfect Information Model (PI)

The perfect information model ("perfect hindsight" or "re-optimization" model) determines the optimal allocation of inventory for each particular realization of demand, and then computes the expected reward over all scenarios:

$$PI(n,t) = PI(n,\mathcal{D}^t) = E_\pi[IP(n, D^t_\pi)].$$

Here π is a scenario corresponding to a sequence of realizations of demand D^t_π, so D^t_π is a deterministic quantity, representing the actual aggregate demand corresponding to scenario π. For each pathwise realization of the demand process, the best possible capacity allocation is computed, together with its corresponding reward, $IP(n, D^t_\pi)$. The PI-value is the average reward over all possible demand scenarios π. An upper bound can be obtained by solving the optimal LP-allocation instead of the integer program:

$$PI^{LP}(n,t) = E_\pi[LP(n, D^t_\pi)] \geq E_\pi[IP(n, D^t_\pi)] = PI(n,t).$$

These models are based on perfect information on the realization of the demand process and do not provide an actual policy, so in this respect they are of little practical value. The objective values, however, provide absolute upper bounds on the performance of any policy in each possible scenario.

3.5 The value of demand information.

So far we have presented various models for the dynamic and stochastic NRM Problem. In this section, we investigate the various relationships between these models and their objectives, and their dependence on the available demand information.

The following result captures the relationship between the objective values of these formulations at any state (n,t):

Theorem 1 $DP(n,t) \leq PI(n,t) \leq PI^{LP}(n,t) \leq LP(n,D^t)$.

Proof: For each pathwise realization of demand, the revenue collected by the DP-policy along that path is upper bounded by the "perfect hindsight" value PI. This is because the PI uses perfect information, thus being optimal for each particular scenario. Therefore, $DP(n,t) \leq PI(n,t)$.

Since $PI^{LP}(n,t) = E_\pi[LP(n,D^t_\pi)]$, and $LP(n,D^t)$ is concave in D^t, we obtain

$$PI(n,t) \leq PI^{LP}(n,t) = E_\pi[LP(n,D^t_\pi)] \leq LP(n,E[D^t_\pi]) = LP(n,D^t)\,,$$

where the last part follows by Jensen's inequality. ■

We examine in Section 7 the value of robustness in the estimation of demand parameters, and the cost of bias and noise in the forecasts.

4 Approximate Dynamic Programming Algorithms

In this section, we present several algorithms for the NRM problem. These algorithms belong in the class of approximate dynamic programming methods (see Bertsekas and Tsitsiklis [4] for a general introduction), in which an approximate value to the exact value function is used in the Bellman equation.

4.1 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 from a certain mathematical programming formulation. Then itinerary bid-prices are computed additively. This technique was initially proposed by Simpson [40], then studied by Williamson [47] in her PhD thesis. For recent work on bid-price control see Talluri and van Ryzin [42] and Gunther et. al. [23].

Most of the developments in the literature compute static bid-prices and use these to determine a one-time static allocation of inventory. Motivated by work by practitioners, Talluri and van Ryzin study adaptive bid-prices, that provide a dynamic policy for inventory control. At each point in time, they solve the Lagrangean relaxation of a probabilistic program, and define the bid-prices as the corresponding Lagrangean multipliers. Their underlying model is a probabilistic program that incorporates randomness in the fares; in the case of deterministic fares, their model coincides with the LP model from Section 3.2.

Given a certain mathematical programming formulation (MP) of the NRM problem, one can describe the corresponding static, and respectively dynamic, bid-price policies generically as follows:

Static Bid-Price Control:

At the initial state S^0 (e.g. $S^0 = (N, T)$),

- Solve $MP(S^0)$ and find the shadow prices for the leg-capacity constraints v^0.
- Compute the static itinerary bid-prices: $\mathrm{BP}_j^0 = (v^0)' \cdot A^j$.

At any given state S^t,

- Sell only to those classes j whose fares exceed their bid-price BP_j^0 : $R_j \geq \mathrm{BP}_j^0$.

ITERATE.

Notice that this policy decides at the very beginning what classes to sell to and in what proportions, ignoring the inherently dynamic and stochastic component of the problem.

Adaptive Bid-Price Control:

At any current state S^t (e.g. $S^t = (n, t)$),

- Solve $MP(S^t)$ and find a set of shadow prices for the leg-capacity constraints v^{S^t}.
- Compute adaptive bid prices for a class j request additively : $\mathrm{BP}_j(S^t) = (v^{S^t})' \cdot A^j$.
- Sell to class j if and only if its fare R_j exceeds its current bid-price $\mathrm{BP}_j(S^t)$, i.e.,

$$R_j \geq \mathrm{BP}_j(S^t).$$

ITERATE.

Notice that this approach is not well specified if the optimal dual variables are not unique. In general, the underlying MP model is a linear program, such as the one proposed in Section 3.2. Several probabilistic models have been investigated by Glover et al. [22], Williamson [47], etc. and Talluri and van Ryzin [42]. It has been observed (Williamson) that the LP model achieves a better performance, and is more efficient.

4.2 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 multi-leg 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 $DP(n,t)$ defined in Eq.(1) by the value of the linear programming problem $LP(n, D^t)$. Thus, a request for a given class j will be accepted if and only if its price R_j exceeds its current estimate for opportunity cost given by:

$$\mathrm{OC}_j^{LP}(n,t) = LP(n, D^{t-1}) - LP(n - A^j, D^{t-1}).$$

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 S is defined as $\mathrm{OC}_j^{MP}(S) = MP(S_{rej(j)}) - MP(S_{acc(j)})$, where $S_{acc(j)}$ and $S_{rej(j)}$ 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 $S = (n,t)$, then $S_{acc(j)} = (n - A^j, t-1)$ and $S_{rej(j)} = (n, t-1)$.

This approximation defines the following adaptive inventory control policy:

Approximate Dynamic Programming (ADP): At any current state $S = S^t$,

- If a class j request arrives, compute the opportunity cost estimate

$$\mathrm{OC}_j^{MP}(S) = MP(S_{rej(j)}) - MP(S_{acc(j)}).$$

- Accept if and only if its fare exceeds its opportunity cost: $R_j \geq \mathrm{OC}_j^{MP}(S)$.

ITERATE.

For example, in the case when the underlying MP model is the LP described in Section 3.2, this is the **Certainty Equivalent Control (CEC)** policy. A similar policy is proposed by Bitran and Mondschein [8] in the context of hotel revenue management, but no analysis is provided. Notice that in the case of linear networks (e.g. hospitality) it is actually desirable to use the equivalent network flow formulation as described in Section 3.3 instead of the usual LP model, since re-optimization at each stage becomes a much simpler task.

4.3 Extensions

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

Squared Heuristics. The expected value of any heuristic H provides an approximation of the value function (DP). Therefore, conceptually, the H-values provide estimates of opportunity costs that are determined as follows:

$$\mathrm{OC}_j^{H}(n,t) = H(n,t-1) - H(n - A^j, t-1).$$

This opportunity cost estimation mechanism leads to a new approximate dynamic programming **(ADP)** heuristic, called squaring H, or H^2. This method is a form of rollout algorithm, or policy iteration. It has been observed in the dynamic programming literature that this procedure tends to improve heuristic performance (see Bertsekas and Tsitsiklis [4], and Bertsimas, Teo and Vohra [5]).

For practical purposes, we suggest using Monte Carlo simulation for evaluating the policy H at a set of values, and interpolating these in an on-line 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.

MC Policy (CEC with MC demand):

At the current state (n, t), $t = T, \ldots, 0$, suppose we have information that the aggregate demand to come $\mathcal{D}_c^{t-1}$ follows a certain type of $\mathcal{F}^{t-1}$-distribution and suppose a class j arrival occurs.

- Generate r trials from the aggregate demand distribution: $\hat{D}_1^{t-1}, \ldots, \hat{D}_r^{t-1} \sim \mathcal{F}^{t-1}$. Compute the OC_j value for each of these potential realizations of demand:

$$\mathrm{OC}_j(n, \hat{D}_i^{t-1}) = LP(n, \hat{D}_i^{t-1}) - LP(n - A^j, \hat{D}_i^{t-1}).$$

- Estimate the opportunity cost of itinerary j as a weighted average of these values:

$$\mathrm{OC}_j^{MC}(n, t) = \sum_{i=1}^{r} \alpha_i \cdot \mathrm{OC}_j(n, \hat{D}_i^{t-1}),$$

  where $\alpha_i = P(\mathcal{D}_c^{t-1} = \hat{D}_i^{t-1} | \mathcal{D}_c^{t-1} \in \{\hat{D}_1^{t-1}, \ldots \hat{D}_r^{t-1}\})$.

- Accept the class j request if and only if $R_j \geq \mathrm{OC}_j^{MC}$.

ITERATE.

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. In order to escape this issue, we actually run a simplified version of this policy, that assigns the same weights to all the trials and defines: $\mathrm{OC}_j^{MC}(n, t) = \frac{1}{r} \cdot \sum_{i=1}^{r} \mathrm{OC}_j(n, \hat{D}_i^{t-1})$.

A Neighbor Search Policy. Even though each iteration of CEC (or BPC) only requires solving an LP, this may be too time-expensive for an on-line task, especially in the case of large networks with many demand classes. For practical purposes, it may be preferable to update the set of OC estimates only once (overnight) or perhaps a few times a day, when we are close to departure, and even more seldom at the beginning of the horizon. We propose the following improved policy that divides the time horizon in pre-established periods of decreasing length, and only solves an LP once at the beginning of each period k. However, instead of computing just one set of dual variables at each update, this new policy will keep track of the neighboring dual bases as well, thus providing more accurate OC estimates:

Neighbor Search Policy (NSP):

At the beginning of each period k, suppose the current state is $(n_k, t_k + 1)$. Compute and store the optimal dual variable (v_k, u_k) corresponding to $LP(n_k, D^{t_k})$, and its neighboring vertices in the dual polyhedron $(v, u) \in \mathcal{N}_k$, where $\mathcal{N}_k$ denotes the set of neighboring vertices of (v_k, u_k), or perhaps only a subset of these.

- During period k, at any current state $(n, t+1)$, suppose a class j arrival occurs.
- Let $LP_k(n, D^t) = \min_{(v,u)\in\mathcal{N}_k} (v'n + u'D^t)$, $\quad LP_k(n{-}A^j, D^t) = \min_{(v,u)\in\mathcal{N}_k} (v'(n - A^j) + u'D^t)$.
  Define $\qquad \mathrm{OC}_j^{NS}(n, t) = LP_k(n, t) - LP_k(n - A^j, t)$.
- Accept the class j request if and only if $R_j \geq \mathrm{OC}_j^{NS}$.

ITERATE

The advantage of this method over the current practice of recomputing bid-prices overnight is that it keeps track of the entire basis (as opposed to just the shadow prices v), and it accounts for possible changes of the basis during the day, by optimizing over neighboring bases. To improve this further, and incorporate variance effects, one can perform a Monte Carlo hybrid at each sub-iteration.

5 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 Section 3.2. 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 [42], 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 [35]).

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

Proposition 2 *In any state $S = (n,t)$ the following inequalities hold:*

$$BP_j(n,t) \leq OC_j^{LP}(n,t) \leq BP_j(n - A^j, t). \tag{3}$$

Both equalities hold if accepting class j does not incur a change of basis in the LP dual.

Proof: Recall that we have defined:

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

where $(v^{n,t}, u^{n,t})$ and $(v^{n-A^j,t}, u^{n-A^j,t})$ are optimal dual solutions of $LP(n, D^{t-1})$ and $LP(n - A^j, D^{t-1})$ respectively. Since both solutions are feasible for both programs, we obtain the following upper bounds by evaluating each LP at the optimal solution of the other:

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

The upper bound in Eq. (3) follows by applying the first of these inequalities in the OC formula, and the lower bound by the second inequality. In case the two dual optimal solutions coincide, we obtain equality throughout. ■

In general, if at a given state the CEC policy accepts a class j request, then at the same state the bid-pricing policy will also accept, but not vice versa. This is because the

y_j^*	$y_j^* \geq 1$ in some y^*	$y_j^* < 1$ in all y^*, but $\neq 0$ in some	$y_j^* = 0$ in all y^*
CEC	accept	reject	reject
BPC	accept	accept	reject

Table 1: The behavior of the BPC and CEC controls as a function of an optimal primal solution $y*$.

following situation may occur: $\mathrm{BP}_j(n,t) \leq R_j < \mathrm{OC}_j^{LP}(n,t)$, as we will see in an actual example described in Section 5.1.

Table 1 provides a comparative characterization of the bid-pricing policy (BPC) versus the certainty equivalent control (CEC) policy, in terms of the structure of the primal optimal solutions 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 $LP(n, D^{t-1})$ and $LP(n - A^j, D^{t-1})$ (if the dual basis does not change, then the policies are identical).

The following two propositions state and prove these results formally.

Proposition 3 *(Structural Properties of the BPC Policy) At any state (n,t), if $LP(n, D^{t-1})$ 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 [6], p. 192), we have $(v^{n,t})' \cdot A^j + u_j^{n,t} > R_j$, i.e., $\mathrm{BP}_j(n,t) = (v^{n,t})' \cdot A^j > R_j$, in which case the bid-price policy rejects class j.

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

■

Proposition 4 *(Structural properties of the CEC Policy) Suppose that* $LP(n - A^j, D^{t-1})$ *and* $LP(n, D^{t-1})$ *have different optimal dual bases. Then:*

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

(b) If $y_j^* < 1$ *in all primal optimal solutions of* $LP(n, D^{t-1})$*, then CEC rejects class* j*.*

Proof: If in some optimal solution y^* of $LP(n, D^{t-1})$, $y_j^* \geq 1$, then $y^* - e_j \geq 0$ is a feasible solution of $LP(n - A^j, D^{t-1} - e_j)$, and hence

$$LP(n, D^{t-1}) = R' \cdot y^* \leq R_j + LP(n - A^j, D^{t-1} - e_j) \leq R_j + LP(n - A^j, D^{t-1}),$$

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

For part (b), since $0 \leq y_j^* < 1$ in all primal optimal solutions we have that $u_j^{n,t} = 0$ and $v^{n,t} \cdot A^j = R_j$ by complementary slackness. Under the assumption that the optimal dual basis changes, we obtain that:

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

that is CEC rejects class j, which concludes the proof. ■

5.1 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 Eq.(3)). 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 2). The following situation may occur: $\mathrm{BP}(j) \leq R_j < \mathrm{OC}(j)$, 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 4 nodes: a hub H, two origin nodes O_1, O_2, and a destination node D. The legs of the network are (1) O_1H, (2) O_2H, (3) HD. 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. The pricing structure is such that: $R_1 < R_{13}$, $R_2 < R_{23}$, $R_{13} + R_2 < R_{23} + R_1$.

Suppose that the available capacity in the current state t is $n = (n_1, n_2, n_3)$ and the demand to come in the remaining $t-1$ periods is $D = D^{t-1}$. The static LP can be formulated as follows:

$$\begin{aligned} LP(n, D) = \max \quad & R_1y_1 + R_2y_2 + R_{13}y_{13} + R_{23}y_{23} \\ \text{s.t.} \quad & y_1 + y_{13} \leq n_1 \\ & y_2 + y_{23} \leq n_2 \\ & y_{13} + y_{23} \leq n_3 \\ & 0 \leq y \leq D. \end{aligned}$$

Let v be the vector of dual variables corresponding to the capacity constraints, and u the vector of dual variables corresponding to the demand constraints. The dual problem can be formulated as follows:

$$
\begin{array}{rlcl}
LP(n,D) = \min & n'v + u'D & & \\
\text{s.t.} & v_1 + u_1 & \geq & R_1 \\
& v_2 + u_2 & \geq & R_2 \\
& v_1 + v_3 + u_{13} & \geq & R_{13} \\
& v_2 + v_3 + u_{23} & \geq & R_{23} \\
& u,\ v & \geq & 0\,.
\end{array}
$$

Suppose that there is one seat left on each leg, so $n = (1,1,1)$, 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 non-binding 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 non-binding, we must have that $u = 0$, 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 $v_1^* = R_1$, $v_2^* = R_2$, $v_3^* = R_{23} - R_2$. And hence the bid-prices for each itinerary are:

$$
\begin{array}{rcccl}
\mathrm{BP}(1) & = & v_1^* & = & R_1, \\
\mathrm{BP}(2) & = & v_2^* & = & R_2, \\
\mathrm{BP}(13) & = & v_1^* + v_3^* & = & R_1 + R_{23} - R_2 > R_{13} \\
\mathrm{BP}(23) & = & v_2^* + v_3^* & = & R_{23}.
\end{array}
$$

So the bid-pricing policy accepts all classes but (13).

We can also compute the opportunities costs as follows:

$$
\begin{array}{rclcccccc}
\mathrm{OC}(1) & = & LP(1,1,1,D) - LP(0,1,1,D) & = & R_1 & = & \mathrm{BP}(1) & & \\
\mathrm{OC}(2) & = & LP(1,1,1,D) - LP(1,0,1,D) & = & R_1 + R_{23} - R_{13} & > & \mathrm{BP}(2) & = & R_2, \\
\mathrm{OC}(13) & = & LP(1,1,1,D) - LP(0,1,0,D) & = & R_1 + R_{23} - R_2 & = & \mathrm{BP}(13) & > & R_{13}, \\
\mathrm{OC}(23) & = & LP(1,1,1,D) - LP(1,0,0,D) & = & R_{23} & = & \mathrm{BP}(23). & &
\end{array}
$$

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 $n = (1,1,1,)$ 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, since 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 [25]):

Definition 1 *A function $f : S \subset R^n \to R$ is said to satisfy decreasing differences on S if for any $s \in S$, and $i \neq j \in \{1, \ldots, n\}$, and for all $\delta_i > 0$ and $\delta_j > 0$ with $s + \delta_i e_i$, $s + \delta_j e_j$ and $s + \delta_i e_i + \delta_j e_j \in S$, the following relation holds:*

$$f(s + \delta_i e_i + \delta_j e_j) - f(s + \delta_i e_i) \leq f(s + \delta_j e_j) - f(s),$$

where e_i is the $i-th$ unit vector.

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 [18] or [14]), and would provide a sufficient condition for the property to extend to the network case.

In the example we have just described here, we can see that the following relation violates the decreasing differences property:

$$LP((1,1,1),D) - LP((0,1,0),D) = R_1 + R_{23} - R_2 > R_{13} = LP((1,0,1),D) - LP((0,0,0),D).$$

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, since the two are asymptotically equal. Surprisingly, this says that incremental revenues (opportunity costs) may increase by decreasing capacity along a certain direction.

6 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, etc. 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 etc. 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 present 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, by incorporating the overbooking decision in the admission control mechanism.

6.1 General Models and Relationships

First, we extend the various models described in Section 3 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, since it is necessary to keep track of the past sales record s. The random quantities involved are the demand, cancellations and no-show processes. Given the initial network inventory N, the maximum expected revenue (less refunds and penalties) to be collected from state (s, t)

onward (cost-to-go), is given by:

$$\begin{aligned} DP^o_N(s,t) &= \sum_j p^t_j \cdot \max(DP^o_N(s,t-1), R_j + DP^o_N(s+e_j,t-1)) \\ &+ \sum_{j|s_j\geq 1} p^{ct}_j \cdot (DP^o_N(s-e_j,t-1) - R^{ct}_j) \\ &+ p^t_0 \cdot DP^o_N(s,t-1) \end{aligned} \tag{4}$$

with the boundary conditions:

$$\begin{aligned} DP^o_N(s,0) &= \sum_{\tilde{s}} P(\tilde{s} \text{ bookings out of } s \text{ materialize }) \cdot DP^o_N(\tilde{s},-1) \\ DP^o_N(s,-1) &= -\min\ C' \cdot s^o \\ & \quad \text{s.t.}\ \ A \cdot (s - s^o) \leq N \\ & \qquad\quad 0 \leq s^o \leq s\,. \end{aligned}$$

The difference from the basic model is that cancellation events are incorporated in the Bellman equation (4), and the boundary conditions are changed to account for no-shows (time $t = 0$) and final overbooking decisions (time $t = -1$). The overbooking decision is made so as to minimize total penalties, while keeping the actual capacity restrictions satisfied. The model does not incorporate secondary effects associated with overbooking (e.g. image damage and loss of goodwill, demand shifting etc.).

The Perfect Information Model

The perfect information, or perfect hindsight model uses exact knowledge of the realization of demand on each particular path, and computes the best inventory allocation for each path. In these conditions, there will be no overbooking situation at the end of the horizon, whereas cancellations and no-shows can be initially adjusted in the demand process. Therefore, if we define the adjusted demand process as $\tilde{\mathcal{D}} = \mathcal{D} - \mathcal{C} - \mathcal{NS}$, the perfect information model $PI^o_N(n,t)$ can be reformulated as $PI(n, \tilde{\mathcal{D}})$.

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, since they depend on cancellations and no-shows which are random processes. In order to define a model that is consistent with these observations, we define $\tilde{R}_j$ and $\tilde{A}^j$ to be the expected revenue gained from, and expected capacity occupied by a class j reservation, before the overbooking period. We denote the expected (or "adjusted") revenue from class j, after discounting potential cancellations and no shows (but not overbooking penalties) as $\tilde{R}_j = R_j - p^c_j \cdot R^c_j - (1-p^c_j) \cdot p^{ns}_j \cdot R^{ns}_j$. Also $\tilde{A}^j = (1-p^c_j) \cdot (1-p^{ns}_j) \cdot A^j$ is the expected capacity occupied by a class j reservation at the end of the horizon, after discounting potential cancellation and no show events. Finally $\tilde{C}_j = (1-p^c_j) \cdot (1-p^{ns}_j) \cdot C_j$ is the average cost of one overbooked class j reservation.

Let p^c_j and p^{ns}_j denote the probability that a given class j reservation is canceled 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.

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

$$
\begin{aligned}
IP^o_N(s,t) = \max \quad & \tilde{R}' \cdot y - C' \cdot z^o \\
\text{s.t.} \quad & 0 \le \tilde{A} \cdot (y+s) - A \cdot z^o \le N \\
& 0 \le y \le D^t \\
& y, z^o \text{ integer.}
\end{aligned}
$$

The vector y decides how many requests to be accepted in the future, whereas z^o determines which itinerary requests should be overbooked at the end of the horizon, if necessary. The final overbooking pad for each leg is $A \cdot z^o$. The capacity constraint requires that cancellation-adjusted past and future sales, less overbookings (a non-negative vector of leg-sales) should not exceed the initial network capacity.

With the change of variable $s^o_j = z^o_j / (1-p^c_j) \cdot (1-p^{ns}_j)$, $j = 1, \ldots, n$, the corresponding linear programming relaxation, and its dual, can be written as follows:

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

6.2 Adjusted Policies

Adjusted BPC Policy. Following the spirit of the basic adaptive bid-pricing policy, we adjust the BPC policy to incorporate cancellations and overbooking as follows:

1. In state $S = (s,t)$, $t > 0$,

   - Solve $LP^o(S)$ and find the shadow prices for the leg-capacity constraints, v^S
   - If a class j request arrives, compute adjusted bid prices : $\mathrm{BP}^o_j(S) = (v^S)' \cdot \tilde{A}^j$. Accept the request if and only if its "adjusted" price $\tilde{R}_j$ is higher than either its adjusted bid-price $\mathrm{BP}^o_j(S)$, or the adjusted overbooking penalty $\tilde{C}_j$, i.e.,

   $$\tilde{R}_j \geq \min(\tilde{C}_j, \mathrm{BP}^o_j(S)).$$

   - If a class j cancellation is requested, adjust the sales record $s = s - e^j$.

   ITERATE.

2. In state $S = (s,0)$, subtract no-shows (ns) from the current sales record: $s = s - ns$.

3. In state $S = (s,-1)$, solve $DP^o_N(s,-1)$ (a simple LP) and determine the optimal overbooking structure.

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. Again, adjustment accounts for the fact that capacity and/or revenue might not be realized or may be overbooked.

1. In state $S = (s,t)$, $t > 0$,

   - If a class j request arrives, compute the opportunity cost estimate

   $$\mathrm{OC}^o_j(s,t) = LP^o_N(s,t-1) - LP^o_N(s+e_j,t-1).$$

   Accept the request if and only if $\tilde{R}_j \geq \mathrm{OC}^o_j(S)$.

   - If a class j cancellation is requested, adjust sales record $s = s - e^j$.

ITERATE.

2. In state $S = (s, 0)$, subtract no shows (ns) from the current sales record $s = s - ns$.

3. In state $S = (s, -1)$, solve the linear program $DP_N^o(s, -1)$ and determine the optimal overbooking structure.

The same type of arguments can be used to extend Theorem 1 for the case of cancellations, no shows and overbooking, showing the following relations:

$$DP_N^o(n,t) \leq PI_N^o(n,t) \leq LP_N^o(n, D^t).$$

Moreover, the structural properties proved in Section 5 as well as the threshold time properties of Section 6.1 are preserved in the overbooking-adjusted policies.

7 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 Section 7.1 we report computational results without cancellations and overbooking, while in Section 7.2 we allow cancellations and overbooking.

7.1 Computational Performance without Cancellations and Overbooking

We performed expected value calculations for two leg instances and simulation runs for larger networks. In order 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.

7.1.1 Models and Assumptions

We considered the following types of networks:

(N2) A 2-leg network with nodes O, H, D, legs $(1)OH$, $(2)HD$ and capacities $N = (N_1, N_2)$. 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 R, is:

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

(N3) The 3-leg network described in the example of Section 5.1, with 4 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 R, is:

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

(N4) A 4-leg network, with 2 origins O_1, O_2, 2 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 $N = (N_1, N_2, N_3, N_4)$. There is demand for all the 8 itineraries, with one fare class per itinerary. The leg-class incidence matrix, together with the fare structure R, is:

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

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

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

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

(HP) **H**omogeneous **P**oisson arrivals with arrival rates given by a constant vector p (i.i.d. Bernoulli trials).

(NHL) **N**on-homogeneous Poisson with **h**i-**l**ow demand. We assume arrival rates increase for high fare classes ($p^t = p/log_a(a + t \cdot \rho)$) and decrease ($p^t = p/log_a(a + (T - t) \cdot \rho)$) for low-fare classes, as we approach departure.

(NLH) **N**on-homogeneous Poisson with **l**ow-**h**i demand. We assume arrival rates decrease for high fare classes ($p^t = p/log_a(a + (T - t) \cdot \rho)$) and increase ($p^t = p/log_a(a + t \cdot \rho)$) 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 $a = 2$).

7.1.2 Running Time

Exact calculations of the optimal expected revenue (DP), and expected values of the proposed policies (CEC, BPC) are practically impossible. Already for 2-leg networks (N2) with one class per itinerary, 50 seats per leg initial capacity and $T = 300$ time periods the computation takes about 2 hrs.

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 2-leg network takes less than a minute (less than a second per iteration). The largest network we simulated was a 4-leg network (N4.2) with 2 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).

	(HP)				(NLH)				(NHL)			
T	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	1086.5	1086.5	1086.5	1086.5	1242.1	1242.1	1242.1	1242.1
60	1170	1170	1170	1170	1318.6	1318.6	1318.6	1318.6	1488.3	1488.1	1488.1	1488
70	1365	1365	1365	1365	1553.4	1552.2	1552.1	1552.2	1733.6	1703	1702.8	1690.4
80	1560	1559.5	1559.5	1559.5	1790.4	1755.3	1754.5	1754.9	1834.3	1800.9	1799.6	1753.4
90	1755	1745.7	1745.4	1745.7	1885.3	1863.3	1860.8	1840.9	1846.1	1831.4	1829.5	1791.1
100	1950	1897.5	1896.4	1897.1	1925.1	1904.7	1901.4	1852.1	1858.4	1845.8	1843.1	1809.6
110	2020	1998.9	1996.6	1993.4	1937.3	1922	1917	1863.3	1871.1	1858.7	1855	1822.5
120	2090	2064.6	2061.6	2031.6	1949.1	1934.3	1927.1	1876.9	1884.4	1872.1	1867.1	1837
130	2140	2110.6	2107.3	2021.9	1960.6	1945.9	1936.7	1889.5	1898.3	1886.2	1879.8	1851.5
140	2170	2145.5	2140.5	2018.3	1971.7	1957.1	1946.2	1899.5	1913	1901	1893.2	1864
150	2200	2175.7	2167.6	2018.9	1982.5	1968.1	1955.6	1907.2	1928.7	1916.7	1907.6	1875.8
160	2230	2202.4	2192.9	2019.3	1993.1	1978.7	1964.9	1913.1	1945.5	1933.6	1923.2	1887.8
170	2250	2222.8	2216.9	2019.4	2003.4	1989.1	1974	1918.1	1963.8	1952	1940.2	1900.8
180	2250	2236.2	2233	2019.4	2013.5	1999.3	1983	1922.9	1984.2	1972.5	1959.3	1915.5
190	2250	2243.8	2242.3	2019.4	2023.5	2009.3	1991.8	1928.6	2007.4	1995.8	1981.4	1932.4
200	2250	2247.5	2246.9	2019.4	2033.2	2019	2000.5	1934.7	2035.3	2023.7	2008.1	1952.8

Table 2: Expected value calculation for (N2) with $N = (50, 50)$, $R = (25, 20, 35)$, $p = (.4, .3, .1)$ and $T = 10 - 200$: the first section considers homogeneous arrivals, and the next two are non-homogeneous arrivals with $\rho = (50, 50, 30)$ and $a = 2$.

For a 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 on-line fashion to provide an adaptive control mechanism where opportunity cost estimates are rapidly recalculated at each iteration.

7.1.3 Quality of approximation

Expected value computations for 2 leg networks. The results in Table 2 show expected value calculations for the case of the 2-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 optimal, whereas the bid-pricing policy appears to be off by a constant. When the time horizon is small, both policies perform optimally. It is not surprising that differences in performance are observed for intermediate stages when both some demand and capacity constraints become binding. In this range, it is not clear which policy is better.

Simulation Results For 2-leg networks we provide in Table 3 a comparison between simulation results and exact calculations for the DP and values of the CEC and BPC

(N2)								
T	LP	DP	CEC			BPC		
		EXP	EXP	SIM	(σ)	EXP	SIM	(σ)
50	975	975	975	973.3	(71)	975	973.2	(71)
100	1950	1897.5	1896.4	1904.05	(83)	1897.4	1903.75	(82)
150	2200	2175.5	2167.6	2174.9	(45)	2018.8	2021.2	(44)
200	2250	2247.5	2246.9	2246.4	(18)	2019.4	2023.3	(17)
300	2250	2250	2250	2250	(0)	2200	2200	(0)

Table 3: Simulations and expected value calculation for 2-leg network. **(N2)** $N = (50, 50)$, $R = (25, 20, 35)$ and $p = (.4, .3, .1)$.

(N3)					
T	LP	CEC	(σ)	BPC	(σ)
100	2300	2273.6	(140)	2268.7	(140)
200	3200	3155.8	(184)	3148.4	(181)
300	3600	3508.7	(178)	3476.9	(177)
500	4300	4158	(102)	4110	(103)
(N4.1)					
T	LP	CEC	(σ)	BPC	(σ)
100	3500	3462.05	(150)	3460.5	(154)
300	5800	5615.55	(78)	5635.35	(74)
600	6050	5951.25	(85)	5842.8	(94)
(N4.2)					
T	LP	CEC	(σ)	BPC	(σ)
100	4070	4087	(213)	4085.25	(216)
200	6125	5914.2	(241)	5856.95	(233)
300	11110	10859	(302)	10846	(309)

Table 4: Simulations for larger networks (top-down): **(N3)** $N = (60, 60, 100)$, $R = (10, 20, 50, 30)$ and $p = (.2, .3, .4)$; **(N4.1)** $N_i = 50$, $R = (20, 20, 50, 30, 60, 40, 45, 55)$, $p = (.2, .25, .05, .05, .1, .15, .1, .1)$; **(N4.2)** $N_i = 50$, $R^l = (20, 20, 50, 30, 60, 40, 45, 55)$, with $R^h = 2 \cdot R^l$ and $p = (.15, .02, .2, .01, .05, .015, .05, .01, .1, .01, .15, .015, .1, .01, .1, .01)$

policies. The high dimensionality of the problem does not allow for exact expected value calculations for larger networks. In Tables 4, we present results from 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 variance σ of the estimation. In all cases we observe that our CEC policy performs better than the BPC policy, by up to 2%.

7.1.4 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.

In order 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. In order to measure the robustness of the LP-value to bias and noise in demand forecasts, we want to compare $LP(n, D)$ and $E_\delta[LP(n, D + \delta)]$. From the dual formulation we have that:

$$E_\delta[LP(n, D + \delta)] - LP(n, D) = E_\delta[\mathrm{OC}_\delta(n, D)] \leq (R' - v' \cdot A)^+ \cdot E[\delta],$$

where v are the leg-shadow prices of $LP(n, D)$. 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 demonstrate that the CEC policy is surprisingly robust to certain forms of noise and bias in the demand data, much more so than the BPC policy.

Correlated Random Noise. At each point in time t we generate a multivariate random variable $\delta \sim N(b, \sigma)$, 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.

	Constant Rate					Log Rate				
T	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	1170	1175.5	1170	1170	1170	1170	1123.5	1170	1170	1170
70	1365	1366.4	1365	1365	1365	1365	1315.3	1365	1365	1365
80	1560	1605.6	1559.5	1559.5	1559.5	1560	1569	1559.5	1559.5	1559.5
90	1755	1770.3	1745.7	1745.2	1745.3	1755	1685.5	1745.7	1745.2	1745.1
100	1950	1977.8	1897.5	1895.8	1895.4	1950	1878.5	1897.5	1895.8	1894.6
110	2020	2086.1	1998.9	1996.1	1961.4	2020	2008.5	1998.9	1996	1974.7
120	2090	2115.2	2064.6	2061.7	1971.7	2090	2069.2	2064.6	2061.6	1967.2
130	2140	2139.4	2110.6	2107.1	1980.1	2140	2105.5	2110.6	2106.3	1981.2
140	2170	2172	2145.5	2139.6	1987.7	2170	2147.1	2145.5	2137.4	1987.9
150	2200	2197.8	2175.7	2166.9	1994.5	2200	2168	2175.7	2162.5	1993.5
160	2230	2227.8	2202.4	2192.8	1999.5	2230	2178.8	2202.4	2185.5	2030.2
170	2250	2250	2222.8	2216.8	2000.2	2250	2203.2	2222.8	2208.4	2151.1
180	2250	2250	2236.2	2233	2000.2	2250	2240.5	2236.2	2227.3	2208.9
190	2250	2250	2243.8	2242.3	2130.6	2250	2235.2	2243.8	2240.3	2223.4
200	2250	2250	2247.5	2246.9	2227	2250	2248.5	2247.5	2246.8	2230.6

Table 5: Computation with noisy forecasts: **(N2)** $N = (50, 50)$, $R = (25, 20, 35)$, $p = (.4, .3, .1)$, $b = (.07, -.05, .03)$, $\sigma = \frac{1}{5}\begin{pmatrix} .1 & -.02 & .04 \\ -.02 & .08 & .01 \\ .01 & .02 & .06 \end{pmatrix}$.

- Log-rate noise: generate δ and add $\delta^t = log(a + \rho \cdot \frac{t-a}{T}) \cdot \delta$ to the arrival rate at each time.

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: $LP_\delta(n, D^t) = LP(n, D^t + E[\delta^t])$.

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 b, and in each consecutive experiment subscripted (αb), 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.

7.1.5 Extensions

Squared 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 so-called squared heuristics CEC^2 and BPC^2. We observe that the squaring 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 2-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 on-line 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 Section 4.3, that assigns the same weights to all the trials and defines:

$$\mathrm{OC}_j^{MC}(n,t) = \frac{1}{r} \cdot \sum_{i=1}^{r} \mathrm{OC}_j(n, \hat{D}_i^{t-1}).$$

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

Table 6: Increasing bias in arrival rates. **(N2)** Data: $N = (19, 19)$, $R = (25, 20, 35)$, $p = (.3, .4, .1)$, $b = (.07, -.05, .03)$.

T	DP	LP	CEC	BPC	LPb	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$_\delta$	CEC$_\delta$	BPC$_\delta$
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

(N2.1)						
T	LP	DP	CEC	BPC	CEC^2	BPC^2
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	1170	1170	1170	1170	1170	1170
70	1365	1365	1365	1365	1365	1365
80	1560	1559.5	1559.5	1559.5	1559.5	1559.5
90	1755	1745.7	1745.4	1745.7	1745.6	1745.7
100	1950	1897.5	1896.4	1897.1	1897.2	1897.4
110	2020	1998.9	1996.6	1993.4	1998	1995.5
120	2090	2064.6	2061.6	2031.6	2063.2	2042.2
130	2140	2110.6	2107.3	2021.9	2108.9	2052.8
140	2170	2145.5	2140.5	2018.3	2143	2104.2
150	2200	2175.7	2167.6	2018.9	2172.9	2163.9
160	2230	2202.4	2192.9	2019.3	2199.6	2197.2
170	2250	2222.8	2216.9	2019.4	2220.6	2217.4
180	2250	2236.2	2233	2019.4	2234.8	2230.2
190	2250	2243.8	2242.3	2019.4	2243.1	2238.1
200	2250	2247.5	2246.9	2019.4	2247.2	2241.9
(N2.2)						
T	LP	DP	CEC	BPC	CEC^2	BPC^2
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

Table 7: Squared Heuristics. **(N2.1)**: $N = (50, 50)$, $R = (25, 20, 35)$, $p = (.4, .3, .1)$; **(N2.2)**: $N = (19, 19)$, $R = (25, 20, 35)$, $p = (.3, .4, .1)$

(N2.1)	LP	PI^{LP}	MC	CEC	BPC
EXP	2200	2191.2	2170	2162.8	2161.6
Std	0	44.48	50.45	52.41	54.19
LB	2200	2070	2035	2000	2000
UB	2200	2250	2250	2250	2240

P	AvgP	StdP	MinP	MaxP
CEC/BPC	1.0006	.0035	.9931	1.0074
CEC/PI^{LP}	.987	.0093	.9569	1
MC/PI^{LP}	.9903	.0084	.9645	1
MC/CEC	1.003	.0067	.9909	1.0225

(N2.2)	LP	PI^{LP}	MC	CEC	BPC
EXP	225	215	208	205	202.5
Std	0	10.55	16.53	18.7	19.76
LB	225	195	185	175	175
UB	225	225	225	225	225

P	AvgP	StdP	MinP	MaxP
CEC/BPC	1.01	.028	1	1.075
CEC/PI^{LP}	.9521	.0498	.8537	1
MC/PI^{LP}	.9664	.038	.9024	1
MC/CEC	1.016	.034	1	1.0857

(N2.3)	LP	PI^{LP}	MC	CEC	BPC
EXP	8000	7922	7785	7771	7766
Std	0	224.3	238.7	285.2	283.8
LB	8000	7500	7200	7050	7050
UB	8000	8600	8400	8500	8400

P	AvgP	StdP	MinP	MaxP
CEC/BPC	1.0007	.011	.975	1.02
CEC/PI^{LP}	.9808	.0166	.94	1
MC/PI^{LP}	.9828	.017	.939	1
MC/CEC	1.002	.0184	.9615	1.0425

(N4)	LP	PI^{LP}	MC	CEC	BPC
EXP	6800	n.a.	6694.4	6588	6568.8
Std	0	n.a.	214.44	213.73	215
LB	6800	n.a.	6260	6140	6160
UB	6800	n.a.	7065	6950	6930

P	AvgP	StdP	MinP	MaxP
CEC/BPC	1.0029	.0051	.9935	1.0117
MC/CEC	1.01624	.0127	.9908	1.0421

Table 8: **(N2)** $N = (50, 50)$, $R = (25, 20, 35)$, $p = (.3, .4, .1)$, $T = 150$, $NT = 100$, $r = 50$; **(N2.1)** $N = (5, 5)$, $R = (25, 20, 35)$, $p = (.3, .4, .1)$, $T = 20$, $NT = 10$, $r = 50$; **(N2.3)** $N = (20, 20)$, $R = (25, 20, 35)$, $p = (.2, .3, .5)$, $T = 50$, $NT = 50$, $r = 50$; **(N4)** $N = (50, 50, 100, 100)$, $R = (20, 20, 50, 30, 60, 40, 45, 55)$, $p = (.15, .15, .15, .15, .1, .1, .1, .1)$, $T = 200$, $NT = 50$, $r = 60$.

the ratio MC/PI^{LP} is surprisingly high, ranging between $96\% - 99\%$. The computational exercises presented in Table 8 show an improvement of up to 20% of MC over CEC.

7.2 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 5 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%.

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

Table 9: The expected revenue in 200 simulation runs as a function of the overbooking penalty. The cancellation probability was 0.01 and the fare of a two-leg itinerary was equal to the fare of a single-leg itinerary.

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.

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

Table 10: The expected revenue in 200 simulation runs as a function of the overbooking penalty. The overbooking penalty was \$130 and the fare of a two-leg itinerary was equal to the fare of a single-leg itinerary.

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.

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

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

8 Conclusions

In this paper, we have presented several models and algorithms for solving the stochastic and dynamic network revenue management problem (NRM). 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:

- They are uniquely defined at each state, whereas for additive bid-prices, there may be several dual optimal solutions;
- The induced 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[3] tend to be large. When the load factor is small both policies perform optimally. There is a critical range when the load factor is close to 1, 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.

[3] We use the term load factor here to express the ratio of expected demand to available capacity.

- 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 squaring 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.
  - From an implementation standpoint, the neighbor search procedure provides a more robust alternative to the usual practice based on overnight bid-price recalculation.

Acknowledgement

We would like to thank Eduardo Messmacher for implementing the algorithms with cancellation and overbooking.

References

[1] J. Alstrup, S. Boas, O.B.G. Madsen, and R.V.V. Vidal. Booking policy for flights with two types of passengers. *Eur. J. Opnl. Res.*, 27:274–288, 1986.

[2] P.P. Belobaba. *Air Travel Demand and Airline Seat Inventory Management.* PhD thesis, MIT, Cambridge, MA, 1987.

[3] P.P. Belobaba. Application of a probabilistic decision model to airline seat inventory control. *Oper. Res.*, 37:183–197, 1989.

[4] D.P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming.* Athena Scientific, Belmont, MA, 1998.

[5] D. Bertsimas, C. Teo, and R. Vohra. Greedy, randomized and approximate dynamic programming algorithms for facility location problems. working paper, 1999.

[6] D. Bertsimas and J. Tsitsiklis. *Introduction to Linear Optimization.* Athena Scientific, Belmont, MA, 1997.

[7] G.R. Bitran and S. Gilbert. Managing hotel reservations with uncertain arrivals. working paper. Weatherhead School of Management, Dept. of Operations Research, Case Western Reserve University, 1992.

[8] G.R. Bitran and S.V. Mondschein. An application of yield management to the hotel industry considering multiple day stays. *Opns. Res.*, 43(3):427–443, 1995.

[9] S.L. Brumelle and J.I. McGill. Airline seats allocation with multiple nested fare classes. *Oper. Res.*, 41:127–137, 1993.

[10] J. Buhr. Optimal sales limits for 2-sector flights. *Proceedings of the AGIFORS 22nd Annual Symposium*, pages 291–303, 1982.

[11] R.E. Chatwin. Optimal airline overbooking. Ph.D. Dissertation, Stanford University, Palo Alto, CA, 1992.

[12] D. Chen. Network flows in hotel yield management. Tech Report TR1225, Cornell University, School of Operations Research and Industrial Engineering, 1998.

[13] V. Chen, D. Gunther, and E. Johnson. A markov decision problem based approach to the airline ym problem. 1998.

[14] Z. Chi. *Airline Yield Management in a Dynamic Network Environment.* PhD thesis, MIT, Operations Research Center, 1995.

[15] T. Cook. Sabre soars. *ORMS Today*, 6:27–31, 1998.

[16] R.E. Curry. Optimal airline seat with fare classes nested by origins and destinations. *Trans. Sci*, 24(3):193–204, 1989.

[17] R.E. Curry. Real-time revenue management: Bid price strategies for origins/destinations and legs. *Scorecard*, 2(Q), 1992. Aeronomics Inc., Atlanta, GA.

[18] M. Diamond and R. Stone. Dynamic yield management on a single leg. Unpublished manuscript, Northwest Airlines, 1991.

[19] M. Dror, P. Trudeau, and S.P. Ladany. Network models for seat allocation of flights. *Trans. Res.*, 22(B):239–250, 1988.

[20] Y. Feng and G. Gallego. Optimal starting times for promotional fares and optimal stopping times for end-of-season sales. *Mgmt. Sci.*, 41:1371–1391, 1995.

[21] S. Fuchs. Managing the seat auction. *Airline Business*, 7:40–44, 1987.

[22] F. Glover, R. Glover, J. Lorenzo, and C. McMillan. The passenger-mix problem in scheduled airlines. *Interfaces*, 12:73–79, 1982.

[23] D. Gunther, V. Chen, and E. Johnson. Airline yield management: Optimal bid prices for single-hub problems without cancellations. 1999.

[24] S. Janakiram, S.S. Stidham, and A. Shaikevich. Airline yield management with overbooking, cancellations and no-shows. Tech. Report No. UNC/OR/TRE 94-9, Dept. of Operations Research, University of North Carolina, 1994.

[25] I. Karaesman and G.J. van Ryzin. Overbooking with substitutable inventory classes. submitted to Opns. Res., 1998.

[26] S.E. Kimes. The basics of yield management. *Cornell H.R.A. Quarterly*, 4:14–19, 1989.

[27] S. Ladany. Dynamic operating rules for motel reservations. *Dec. Sci*, 7:829–840, 1976.

[28] T.C. Lee and M. Hersh. A model for airline seat inventory control with multiple seat bookings. *Trans. Sci.*, 27:1252–265, 1993.

[29] B.A. Leida. A cost model of internet service providers: implications for internet telephony and yield management. Master's thesis, MIT, Dept. of Electrical Engineering, 1998.

[30] V. Liberman and U. Yechiali. On the hotel overbooking problem – an inventory system with stochastic cancellations. *Mgmt. Sci.*, 11:1117–1126, 1978.

[31] K. Littlewood. Forecasting and control of passengers. *Proceedings of the AGIFORS 12th Annual Symposium*, pages 95–128, 1972.

[32] S.A. Maragos. *Yield Management for the Maritime Industry.* PhD thesis, MIT, Dept. of Ocean Engineering, 1994.

[33] J. McGill and G.J. van Ryzin. Revenue management: Research overview and prospects. *Trans. Sci.*, 1999. to appear.

[34] I.C. Paschalidis and J.N. Tsitsiklis. Congestion dependent pricing of network services. Tech. Report, Boston University, Systems Group, Department of Manufacturing Engineering, 1998.

[35] I. Popescu. Applications of optimization in probability, finance and revenue management. Ph.D. Dissertation, Applied Mathematics Department and Operations Research Center, MIT, Cambridge, Massachusetts, 1999.

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

[37] M. Rothstein. An airline overbooking model. *Trans. Sci.*, 5:180–192, 1971.

[38] M. Rothstein. Hotel overbooking as a markovian sequential decision process. *Dec. Sci.*, 5:389–404, 1974.

[39] E. Schlifer and Y. Vardi. An airline overbooking policy. *Trans. Sci.*, 9:101–114, 1975.

[40] R.W. Simpson. Using network flow techniques to find shadow prices for market and seat inventory control. Memorandum M89-1, MIT Flight Transportation Laboratory, Cambridge, MA, 1989.

[41] J.and S. Stidham Jr. Subramanian and C. Lautenbacher. Airline yield management with overbooking, cancellations and no-show. Technical report, Department of Operations Research, University of North Carolina, Chapel Hill., 1997.

[42] K.T. Talluri and G.J. van Ryzin. An analysis of bid-price controls for network revenue management. *Mgmt. Sci.*, 44:1577–1593, 1998.

[43] B. Titze and R. Greisshaber. Realistic passenger behavior and the simple low-fare/high-fare seat allotment model. *Proceedings of the AGIFORS 23rd Annual Symposium*, pages 197–223, 1983.

[44] G.J. van Ryzin and J. McGill. Revenue management without forecasting or optimization: An adaptive algorithm for determining seat protection levels. *Management Science*, 1999. to appear.

[45] K. Wang. Optimum seat allocation for multi-leg flights with multiple fare types. *Proceedings of the AGIFORS 23rd Annual Symposium*, pages 225–246, 1983.

[46] L.R. Weatherford and S.E. Bodily. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking and pricing. *Opns. Res.*, 41:831–844, 1992.

[47] E.L. Williamson. Airline network seat control. Ph.D. Dissertation, MIT, Cambridge, MA, 1992.

[48] R.D. Wollmer. An airline reservation model for opening and closing fare classes. Unpublished Internal Report, McDonnell Douglas Corporation, Long Beach, CA, 1986.

[49] R.D. Wollmer. A hub-spoke seat management model. Unpublished Internal Report, McDonnell Douglas Corporation, Long Beach, CA, 1986.

[50] R.D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Opns. Res.*, 40:26–37, 1992.

[51] J.T. Wong, F.S. Koppelman, and M.S. Daskin. Flexible assignment approach to itinerary seat allocation. *Transpn. Res.*, 27(B):33–48, 1993.