

Pricing substitutable flights in airline revenue management

Dan Zhang^a, William L. Cooper^{b,*}

^a *Desautels Faculty of Management, McGill University, 1001 Sherbrooke Street West, Montreal, Quebec H3A 1G5, Canada*

^b *Graduate Program in Industrial and Systems Engineering, Department of Mechanical Engineering, University of Minnesota, 111 Church Street S.E., Minneapolis, MN 55455, United States*

Received 1 December 2005; accepted 6 October 2006

Available online 21 February 2008

Abstract

We develop a Markov decision process formulation of a dynamic pricing problem for multiple substitutable flights between the same origin and destination, taking into account customer choice among the flights. The model is rendered computationally intractable for exact solution by its multi-dimensional state and action spaces, so we develop and analyze various bounds and heuristics. We first describe three related models, each based on some form of pooling, and introduce heuristics suggested by these models. We also develop separable bounds for the value function which are used to construct value- and policy-approximation heuristics. Extensive numerical experiments show the value- and policy-approximation approaches to work well across a wide range of problem parameters, and to outperform the pooling-based heuristics in most cases. The methods are applicable even for large problems, and are potentially useful for practical applications.

© 2008 Elsevier B.V. All rights reserved.

Keywords: Pricing; Yield management; Markov decision processes; Customer choice

1. Introduction

The development of Internet distribution channels has helped create both opportunities and challenges in airline-ticket pricing. On one hand, it has allowed price changes to be made quickly and frequently with negligible costs. On the other, prices have become more visible to consumers because comparison shopping can be done with the click of a mouse. Old revenue management models that rely on the notion of “exogenous demand for a fare class” are becoming less appropriate, and consequently, it is important to develop operational models that incorporate customer choice.

From a customer’s viewpoint, flight schedule and price information is often readily available when making a purchase decision. For example, on November 22, 2006, the website of JetBlue Airways showed nine different flights spread throughout the day from 6:15 a.m. to 9:15 p.m. for a one-way trip on December 4, 2006 from New York (JFK) to Orlando (MCO). Prices differed across the flights. The earliest and latest flights were priced at \$79. One of the other seven flights was priced at \$124, and all the other flights were priced at either \$79 or \$99. Given such information on flight schedule and the price quotes, customers make their purchase decisions based on their own preferences: Do they mind taking a very early flight, or taking an evening flight and arriving late in the night? From our own experience as consumers, it is not hard to imagine that prices on the mid-day flights would affect choices of customers and effectively change demand for morning or evening flights. Given that consumers do typically choose among alternatives, how should an airline price its flights? This paper contains models to help answer this question.

* Corresponding author. Tel.: +1 612 624 4322; fax: +1 612 625 4344.

E-mail addresses: dan.zhang@mcgill.ca (D. Zhang), billcoop@me.umn.edu (W.L. Cooper).

We study a dynamic pricing problem for multiple substitutable flights between the same origin and destination, in which the airline's objective is to maximize the total expected revenue from customer bookings over the finite selling horizon by setting the prices of the flights. Customers choose among the flights (or decide not to purchase) based upon their own preferences and the prices of *all* the flights offered. The problem we consider is particularly relevant to low-cost airlines that sell many tickets on the Internet, fly point-to-point between select city pairs, and have multiple-flights scheduled for the same day between each pair. Low-cost airlines also typically use a simplified fare structure (in comparison to those employed by traditional carriers). At each point in the booking horizon, prices for each flight are visible to customers. In addition, the effect of capacity inflexibility is even more pronounced for many low-cost airlines in comparison to traditional carriers. For example, JetBlue owns only one type of airplane (Airbus A320), which makes capacity adjustments by contingent fleet assignment impossible. Given the inflexible capacity and the simplified fare structure, pricing is a crucial lever for matching demand and capacity. Other low-fare airlines that face these issues include European carriers such as Ryanair and EasyJet. Although Internet sales make up a smaller fraction of total business for major carriers (such as Northwest, Delta, etc.), they too face problems in which customers have a choice among multiple flights between a common origin and destination. More generally, our models are applicable to retailers that employ dynamic pricing strategies to sell substitutable products over a finite horizon.

Most models for airline pricing consider one-dimensional problems with a single flight, where customer purchase decisions are based only upon the price for that flight. Such one-dimensional problems do not account for the effects of consumer behavior that are captured in our model. In our multiple-flight setting, the fact that bookings for any particular flight are influenced by the price of tickets on all the flights makes solving the problem considerably more difficult than solving multiple single-dimensional problems each with just one flight and no consumer substitution effects.

1.1. Literature review

Dynamic pricing problems have been studied extensively in the economics, marketing, and operations literature. We review pricing research only in the revenue management context that is directly related to our model. In this body of literature, capacity (or inventory) is assumed to be fixed, or is prohibitively expensive to change during the selling horizon. For reviews of pricing models for revenue management, please refer to [Bitran and Caldentey \(2003\)](#) and [Talluri and van Ryzin \(2004b\)](#), and for a survey of the literature that considers both pricing and inventory decisions, see [Elmaghraby and Keskinocak \(2003\)](#). We also briefly describe some revenue management models where customer choice is modeled explicitly, although no pricing decisions are involved. We close by reviewing several studies aimed at deriving customer choice parameters in the airline context.

Dynamic pricing problems for a fixed stock of a single item sold in a finite selling horizon have attracted considerable attention in the revenue management literature. [Gallego and van Ryzin \(1994\)](#) formulate an intensity control model of the problem and derive several structural properties. They also study a heuristic policy based on a deterministic upper bound and prove that it is asymptotically optimal. [Zhao and Zheng \(2000\)](#) consider a similar problem with nonhomogeneous demand and show that dynamic pricing policies can have a significant impact on revenue when demand is nonhomogeneous. [Bitran and Mondschein \(1997\)](#) present a continuous-time model in the context of fashion retailing and compare it to a model with periodic pricing review, where price is allowed to change only at several pre-specified time points. They show that the loss in expected revenue from implementing an appropriate periodic pricing review policy is small.

[Gallego and van Ryzin \(1997\)](#) consider dynamic pricing problems where a set of resources is used to produce a set of products. They develop asymptotically optimal heuristic policies and apply their results to network revenue management problems. [Kleywegt \(2001\)](#) considers a deterministic optimal-control formulation of a pricing problem where multiple products are sold to multiple customer classes over time. [Lin and Li \(2004\)](#) develop bounds on the value function of a dynamic pricing problem for a line of substitutable products. [Liu and Milner \(2006\)](#) study a multi-item pricing problem with a common pricing constraint. They obtain an optimal policy for a deterministic version of the problem and propose heuristics for the stochastic version.

[Talluri and van Ryzin \(2004a\)](#) study a single-leg revenue management problem where customers choose among the open fare classes. They prove structural properties that greatly simplify the computation of an optimal policy. [Maglaras and Meissner \(2006\)](#) consider the pricing problem faced by a firm that owns a fixed stock of a resource, which is used to produce several different products. Customer choice among the products is modeled by joint price elasticity. They prove structural properties that reduce the decision problem to an equivalent one-dimensional problem, and propose several heuristic policies. [Iyengar et al. \(2004\)](#) and [van Ryzin and Liu \(2007\)](#) consider choice-based linear programming models for network revenue management. [van Ryzin and Vulcano \(2006\)](#) consider a network revenue management problem with customer choice behavior, and propose a simulation-based optimization approach to obtain virtual nesting controls. [Boyd and Kalllesen \(2004\)](#) discuss the impact of consumer purchase behavior on revenue management practice, distinguishing between two types of demand: yieldable, where demand is class-specific, and priceable, where demand is price-sensitive and not class-specific.

McFadden (2000) reviews the economics literature dealing with models and estimation for consumer choice in travel. Perhaps the most widely used choice models in practice are discrete-choice models (see, e.g., Ben-Akiva and Lerman, 1985). Train (2003) summarizes recent advances in discrete-choice theory and its applications, and discusses simulation-based methods to estimate choice probabilities for several discrete-choice models. Utility maximization is frequently used as a basis for deriving customer choice probabilities. Mahajan and van Ryzin (2001) point out that a number of choice models can be viewed as special cases of the utility maximization model.

Several recent works also focus on describing or fitting particular choice models in the revenue management context. Carrier (2003) considers how to model passenger preference on flight schedule, and reports results from an extensive simulation study. Algiers and Beser (2001) describe how to estimate customer choice probabilities for flights and fare classes using revealed preference and stated preference data. Andersson (1998) reports on a study of passenger choice in the context of seat inventory control. Talluri and van Ryzin (2004a) use a maximum likelihood method to estimate multinomial logit choice probabilities for fare classes on a single flight.

1.2. Overview

We pose the joint pricing problem for multiple substitutable flights between the same origin and destination as a Markov decision process (MDP). The MDP has multi-dimensional state and action spaces, and therefore suffers from the well-known curse of dimensionality. Since the problem is intractable, we develop and analyze a variety of bounds and heuristics. We begin by formulating three related problems, each based upon some notion of pooling. In addition to yielding bounds on the value function of the original problem, these “pooled problems” suggest various baseline pricing heuristics for the original problem. Among these are *single-price policies* that, for each time period, quote the same price for all open flights (the price is, however, updated as time progresses).

We also derive separable upper and lower bounds for the value function of the original problem. These separable bounds are based upon solutions of several corresponding one-dimensional MDPs. The bounds and the associated one-dimensional problems suggest two other families of heuristics, which we term value approximation and policy approximation. These heuristics have the advantage that they remain computationally tractable, even for very large problems with many flights.

Our numerical studies show that the value- and policy-approximation heuristics appear to work well, and to perform better than the pooling-based heuristics, especially when there is asymmetry among the flights in terms of demand load and customer preferences. Several other insights also emerge from the study. For instance, the results show that the revenue loss from instituting a single-price policy can be quite significant, even when the best possible single-price policy is used. This shortcoming of single-price policies underscores the importance of using sophisticated policies that allow different prices for different flights. Moreover, this observation potentially has relevance beyond airline revenue management. In fashion retailing, as described in Bitran et al. (1998), it may be required that products at different physical locations be priced identically. Our study indicates that such a requirement (typically made to allow simpler centralized pricing control or to protect against loss of customer good will) may result in a significant loss in revenue.

We also examine via numerical experiments policies that change prices only at pre-specified time points. The analysis reveals that for practically implementable choices of such time points, the revenue loss from these policies in comparison to an optimal policy is small. The main insight here is that much of the benefit from using sophisticated dynamic pricing policies can be obtained even if the airline does not exercise complete real-time control of prices. This observation has practical significance, since airlines may want (or be able) to change prices only at certain pre-specified times, such as after daily or semi-daily database updates.

Before proceeding, we compare this paper to Zhang and Cooper (2005), hereafter ZC, which considers a seat availability problem with multiple flights between the same origin and destination in which customers choose among the open flights. For some comments comparing pricing control and availability control, see pp. 176–177 of Talluri and van Ryzin (2004b), where it is argued that pricing, when possible, is the better option. ZC assumes that customers belong to different *classes* that arrive sequentially in distinct periods, and that the order of the classes is pre-determined. The fare for each class is the same on all the flights. The decisions involve the number of seats to open on each flight in each period. The problem in the present paper is related to that in ZC, but differs in the following aspects. First, the present paper considers pricing decisions rather than availability decisions, and prices for different flights can be different in the same period. Second, the present paper assumes that there is at most one arrival in each period as opposed to the sequential “block-demand” setup. Third, the present paper does not assume a pre-determined order of arrivals.

The methodological approach in the present work is also related to that in ZC. Both papers consider solution procedures that bound value functions of high-dimensional MDPs with sums of value functions of one-dimensional MDPs, and both consider value-approximation heuristics. However, to apply the bounds and value-approximation procedures one must identify and exploit the structure of the specific problem. Hence, the particulars of the methods are different in the two papers. Moreover, the increased complexity of the pricing problem expands the scope of methods one might

use. For instance, policy approximation does not have a counterpart in ZC, and questions regarding the frequency of price changes are not easily incorporated into the block-demand setup of our earlier work. In addition, the added complexity of the pricing problem leads us to consider three different pooling procedures in this paper, while only one form is considered in ZC.

To summarize, our main contributions are (1) formulation of the pricing problem for multiple flights with customer choice among the flights, (2) development of bounds for the value function of the MDP, (3) proposal of heuristic approaches to the problem, and (4) numerical testing that provides managerial insights and shows the proposed approaches to be potentially viable. The remainder of the paper is organized as follows. Section 2 reviews the single-flight case. Section 3 formulates the multi-flight problem. Section 4 considers various pooling models. Section 5 develops the separable bounds. Section 6 introduces the value- and policy-approximation heuristics. Section 7 reports numerical results. Section 8 provides a brief summary. Proofs are in the Appendix.

2. Preliminaries: pricing a single flight

In this section, we formulate a single-flight pricing problem as an MDP with one-dimensional state and action spaces. The model, which is a discrete-time analog of that in Zhao and Zheng (2000), is a building block for the multiple-flight case, and its solutions are, in part, the basis for heuristics we develop for the multi-flight case.

Consider a single-leg flight with capacity q . The booking horizon is divided into τ discrete-time periods. The earliest period is period τ , and the last period is period 1. In each period t , independent of everything else, there is one customer arrival with probability λ_t , and no customer arrival with probability $1 - \lambda_t$. In any time period, there is a single price in effect. The set of allowable prices is denoted by $\mathcal{P} = \{\rho_0, \rho_1, \dots, \rho_k\}$ with $\rho_0 > \rho_1 > \dots > \rho_k$. In period t , given a price $r \in \mathcal{P}$, an incoming customer buys a ticket with purchase probability $P_t(r)$ and makes no purchase with probability $1 - P_t(r)$. To model situations where the flight is closed or capacity is depleted, we assume $P_t(\rho_0) = 0$ for all t . The price ρ_0 is often called *null price* in the literature. Throughout, we do not consider overbooking. The objective is to maximize total expected revenue from bookings subject to the capacity constraint.

To formulate the MDP, let the state $x \in \{0, 1, \dots, q\}$ be the number of unsold seats at the beginning of a period. Given state x at time t , let $w_t(x)$ be the maximum expected revenue over time periods $t, \dots, 1$. Let $\mathcal{P}(x) = \mathcal{P}$ if $x > 0$, and $\mathcal{P}(x) = \{\rho_0\}$ otherwise. For any function $u(\cdot)$ of a single variable define $\Delta u(x) = u(x) - u(x - 1)$. For each t and x , the optimality equation is

$$w_t(x) = \max_{r \in \mathcal{P}(x)} \{ \lambda_t P_t(r) [r + w_{t-1}(x - 1)] + [1 - \lambda_t P_t(r)] w_{t-1}(x) \} = \max_{r \in \mathcal{P}(x)} \lambda_t P_t(r) [r - \Delta w_{t-1}(x)] + w_{t-1}(x). \tag{1}$$

We use the convention $w_t(-1) = 0$. The boundary conditions are $w_0(x) = 0$ for all x . A policy that specifies a maximizing action $r_t^*(x)$ in (1) for each x and t is optimal. The optimal price $r_t^*(x)$ given by an optimal policy satisfies

$$r_t^*(x) \geq \Delta w_{t-1}(x), \tag{2}$$

whenever $\lambda_t P_t(r_t^*(x)) > 0$. Inequality (2) says that, given state x , an optimal price in period t should be at least the marginal value of the x th remaining seat.

An important property of the single-flight model is that $\Delta w_t(x) \leq \Delta w_t(x - 1)$, which can be proved by induction. Using this fact, it can be shown that there is an optimal policy $\{r_t^*(x)\}$ such that for all t , we have $r_t^*(x') \leq r_t^*(x)$ if $x \leq x'$. In words, in any given time period, higher remaining inventory implies a lower optimal price. Hence, there is a threshold-type optimal policy; that is, for each fixed t , there exists a set of thresholds $\{I_j(t) : j = 0, 1, \dots, k\}$ with $0 = I_0(t) \leq I_1(t) \leq \dots \leq I_k(t) = q$, such that the price ρ_j is optimal if $I_{j-1} < x \leq I_j$ for $j = 1, \dots, k$. The price must be ρ_0 when $x = 0$.

3. Model formulation

In this section we formulate the multi-flight pricing problem that is the main focus of this paper. There are n flights between a single origin–destination pair. The booking horizon is divided into τ discrete-time periods, and time is counted backwards. To simplify notation, we reserve the symbols i and t for flights and times, respectively, where $i \in \{1, \dots, n\}$ and $t \in \{1, \dots, \tau\}$. The capacity of flight i is c^i . Let $c = (c^1, \dots, c^n)$. Customer arrivals are independent across time periods. In period t , there is one customer arrival with probability λ_t and no customer arrival with probability $1 - \lambda_t$. The prices of the flights are denoted by a vector $r = (r^1, \dots, r^n)$, where r^i is the price for flight i . The allowable prices for flight i are in the set $\mathcal{R}_i = \{\rho_0, \rho_{i,1}, \rho_{i,2}, \dots, \rho_{i,k_i}\}$, where k_i is a constant for each i and $\rho_0 > \rho_{i,1} > \rho_{i,2} > \dots > \rho_{i,k_i}$. We have added a null price ρ_0 to the price set to model cases when flight i is not offered. It is without loss of generality to take ρ_0 to be the same on all the flights. Let $\mathcal{R} = \mathcal{R}_1 \times \dots \times \mathcal{R}_n$.

Customers choose among the flights or purchase nothing. The choice of a particular customer depends on the price vector r . Given a price vector r in period t , and given that a customer arrives in period t , the probability that the customer

purchases a ticket on flight i is $P_t^i(r)$. We use $P_t^0(r)$ to denote the probability that an arriving customer does not make a purchase. For each t and r , the choice probabilities satisfy (i) $\forall i$, if $r^i = \rho_0$, then $P_t^i(r) = 0$; otherwise, $P_t^i(r) \geq 0$, and (ii) $\sum_{i=0}^n P_t^i(r) = 1$. In Section 7, we describe choice models that satisfy the two conditions.

To formulate the MDP, let the state $s = (s^1, \dots, s^n)$ be the vector whose i -th entry, $s^i \in \{0, 1, \dots, c^i\}$, is the number of unsold seats on flight i . Let $\mathcal{R}(s) = \{r : r^i \in \mathcal{R}_i \text{ if } s^i > 0 \text{ and } r^i = \rho_0 \text{ if } s^i = 0 \forall i\}$ be the set of allowable price vectors given state s . Note that when there is no remaining capacity on flight i , the only allowable price is ρ_0 ; i.e., the flight is closed. For each i let e^i be the n -vector with the i th component 1 and zeros everywhere else, and for any function $v(\cdot)$ of n variables define $\Delta_i v(s) = v(s) - v(s - e^i)$.

Let $v_t(s)$ be the maximum expected revenue from periods $t, \dots, 1$ given the state at time t is s . For each t and s , the optimality equation for the MDP is

$$\begin{aligned} v_t(s) &= \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i + v_{t-1}(s - e^i)] + [1 - \lambda_t + \lambda_t P_t^0(r)] v_{t-1}(s) \right\} \\ &= \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i - \Delta_i v_{t-1}(s)] \right\} + v_{t-1}(s). \end{aligned} \tag{3}$$

The boundary conditions are $v_0(s) = 0$ for all s , and we set $v_t(s) = 0$ if $s^i = -1$ for some i .

The MDP associated with (3) has n -dimensional state and action spaces. Although the backward induction algorithm can be applied to (3) to obtain an optimal policy, the computational effort can be overwhelming. In principle, an optimal policy specifies a vector of prices for each possible state s in each period. Even if we could compute such a policy, storage and implementation of the policy would be, for practical purposes, impossible, except when n is very small. For a 10-flight problem (e.g., like that mentioned in the Introduction) with 100 seats on each flight, there are 10^{20} states. If there are 1000 periods, then to store an optimal policy, we need to store a 1×10 price vector for each state in each period – so 10^{24} numbers need to be stored. Structural properties of optimal policies could potentially decrease the amount of information to be stored; however, it is not hard to find examples for which even the natural analog of property (2) does not hold. That is, for an optimal policy $\{\hat{r}_t(s)\}$, it is not necessarily the case that $\hat{r}_t^i(s) \geq \Delta_i v_{t-1}(s)$ for all i .

It can be shown that the value function is increasing in s and t . Under additional assumptions on choice probabilities, we are able to obtain some structural properties. Unfortunately, these properties together are not enough to allow practical computation and storage of an optimal policy.

We say that $P = \{P^1(r), \dots, P^n(r) : r \in \mathcal{R}\}$ satisfies *Condition S* (S is for substitutes) if for all $k \neq j$ and $z^j > r^j$, $P^k(r) \leq P^k(r^{-j}, z^j)$, where (r^{-j}, z^j) is the vector r with the j th component replaced by z^j . The condition says that when the price of one flight increases, the probability that a customer chooses other flights increases. Condition S is related to the concept of positive cross price elasticity, which is often used to characterize substitutable products in the economics literature. For choice probabilities that satisfy Condition S , an analog of property (2) holds.

Proposition 1. Fix state s and time t . Suppose P_t satisfies Condition S and $\hat{r}_t(s)$ is a maximizing action in (3). If $\lambda_t P_t^i(\hat{r}_t(s)) > 0$, then $\Delta_i v_{t-1}(s) \leq \hat{r}_t^i(s) \forall i$.

Proposition 2. Suppose sequences of choice probabilities $\{P_t\}$ and $\{R_t\}$ satisfy Condition S and $P_t^i(r) \geq R_t^i(r) \forall i, t, r$. Let $v_t^p(\cdot)$ and $v_t^r(\cdot)$ be the value functions associated with $\{P_t\}$ and $\{R_t\}$ respectively. Then $v_t^p(s) \geq v_t^r(s) \forall s, t$.

4. Pooling

The MDP of the previous section has n -dimensional state and action spaces, making it intractable. In this section, we consider pooling models with one-dimensional action space, state space, or both. For models with price pooling, we assume the set of allowable prices is the same for all flights. We denote this set by \mathcal{R}_0 , so $\mathcal{R}_i = \mathcal{R}_0$ for all i .

We first consider a model where at each time point a common price is quoted on all the flights with positive remaining capacity. To specify the *price pooling* model, let $\hat{\mathcal{R}}(s) = \{r \in \mathcal{R}(s) : r^i = r^j \in \mathcal{R}_0 \forall i, j \text{ with } s^i, s^j > 0\}$,

The value function $v_t^{PP}(s)$ of the price pooling model satisfies (3) with the action set $\mathcal{R}(s)$ in state s replaced by $\hat{\mathcal{R}}(s)$. We refer to the associated pricing policy as the price pooling (PP) heuristic. Although the MDP has a one-dimensional action space, it still has an n -dimensional state space, and hence the price pooling model is intractable for moderate n .

Next, we consider an MDP with a one-dimensional state space and n -dimensional action space, In the *inventory pooling* model the capacities of the n flights are regarded as perfect substitutes, and are assumed to form the capacity of a *single* flight. Prices on different flights in the original problem are viewed as prices for different classes on the pooled flight, and choice of flights in the original problem is re-cast as choice of fare classes in the pooled problem. The pooled model involves a single flight with capacity $c^p = \sum_{i=1}^n c^i$ and n fare classes. (Throughout, a superscript “ p ” indicates a scalar obtained by

summing entries of a vector.) Customers choose among the n fare classes. The fare of class i is $r^i \in \mathcal{R}_i$. Given price vector r in period t , a customer chooses class i with probability $P_t^i(r)$. Let $v_t^{\text{IP}}(x)$ be the value function and let $\tilde{\mathcal{R}}(x) = \{r : r^i \in \mathcal{R}_i \text{ if } x > 0 \text{ and } r^i = \rho_0 \text{ if } x = 0 \forall i\}$ be the set of allowable price vectors when the state is x . The optimality equations are

$$v_t^{\text{IP}}(x) = \max_{r \in \tilde{\mathcal{R}}(x)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i - \Delta v_{t-1}^{\text{IP}}(x)] \right\} + v_{t-1}^{\text{IP}}(x) \quad \forall t, x \tag{4}$$

with boundary conditions $v_0^{\text{IP}}(x) = 0 \forall x$.

The policy from the inventory pooling model can be used to control the multi-flight problem with customer choice. In particular, suppose $\{\hat{p}_t(x)\}$ is an optimal pricing policy for the inventory pooling model. The inventory pooling (IP) heuristic for the multiple-flight problem is the policy that in state s at time t sets the price on flight i to be $\hat{p}_t^i(s^p)$ if $s^i > 0$ and ρ_0 if $s^i = 0$.

The computations of the inventory pooling problem are potentially quite demanding, because of the multi-dimensional action space $\tilde{\mathcal{R}}(x)$. Without taking advantage of any structural properties, the maximization on the right-hand side of (4) requires evaluating $|\tilde{\mathcal{R}}(x)|$ possibilities. In the following two paragraphs, we briefly outline how the analysis of Talluri and van Ryzin (2004a) can be directly adapted to solve this problem. (Talluri and van Ryzin (2004a) assume that the price of each fare class is fixed and the decision is about which fare classes to open and which to close. Here the decision involves what prices to charge for the different fare classes; the price for, say, class- i can be assigned any value in \mathcal{R}_i , including the null price.)

Given a price vector r , let $W_t(r) = \lambda_t \sum_{i=1}^n P_t^i(r) r^i$ and $C_t(r) = \lambda_t \sum_{i=1}^n P_t^i(r)$. Here, $W_t(r)$ and $C_t(r)$ are the expected revenue in period t and the expected number of seats purchased in period t when price r is used in period t . We can now rewrite (4) as

$$v_t^{\text{IP}}(x) = \max_{r \in \tilde{\mathcal{R}}(x)} \{W_t(r) - \Delta v_{t-1}^{\text{IP}}(x) C_t(r)\} + v_{t-1}^{\text{IP}}(x). \tag{5}$$

For fixed t , a price vector r' is said to be *inefficient* if there exist probabilities $Q(\cdot)$ on \mathcal{R} with $\sum_{r \in \mathcal{R}} Q(r) = 1$ such that $C_t(r') \geq \sum_{r \in \mathcal{R}} Q(r) C_t(r)$ and $W_t(r') < \sum_{r \in \mathcal{R}} Q(r) W_t(r)$. Otherwise, r' is *efficient*. It can be shown that an inefficient price vector is never optimal in (5).

For fixed t , let $N \subseteq \mathcal{R}$ be the set of efficient price vectors. Suppose $N = \{r_k : k = 1, \dots, m\}$. Then for $r, r' \in N$, if $C_t(r) \leq C_t(r')$, then $W_t(r) \leq W_t(r')$. Therefore, we can assume without loss of generality that $C_t(r_1) \leq \dots \leq C_t(r_m)$ and $W_t(r_1) \leq \dots \leq W_t(r_m)$. It can be shown that a maximizing action in (4) is to choose a price vector $r_{k^*(x)} \in N$, where the optimal index $k^*(x)$ is increasing in x for fixed t . This simplifies computations for the pooled problem. However, the determination of the set of efficient price vectors N could itself be a formidable task when there are many fare classes (or equivalently, many flights in the original problem). For a 10-class problem with five price points for each class, we may need to evaluate $5^{10} \approx 10^7$ price vectors to determine N .

It is also possible to aggregate both state and action spaces by adding up the capacity of all the flights to form a single flight and assuming that a single price is quoted in each period. We call this the *inventory and price pooling* model. The model is a one-flight pricing model as discussed in Section 2. For $f \in \mathcal{R}_0$, let $Q_t(f) = \sum_{i=1}^n P_t^i(\vec{f})$, where \vec{f} is the n -vector with f in each entry. Let $v_t^{\text{IPP}}(\cdot)$ denote the value function, which is associated with capacity c^p , arrival probabilities $\{\lambda_i\}$, purchase probabilities $\{Q_t(\cdot)\}$ (see Section 2). As discussed in Section 2, such a model is easy to solve. The resulting policy is a reasonable choice for building a simple heuristic method for the multi-flight problem. In particular, suppose $\{f_t^*(x)\}$ is an optimal policy for the inventory and price pooling model. For each given time t and state s , let $r_t(s)$ be such that $r_t^i(s) = f_t^*(s^p)$ if $s^i > 0$ and $r_t^i(s) = \rho_0$ if $s^i = 0$. We call $\{r_t(s)\}$ the *inventory and price pooling (IPP) heuristic*.

The following summarizes the relationships among the value functions we have encountered.

Proposition 3. *Suppose $\mathcal{R}_i = \mathcal{R}_0 \forall i$. Then $v_t^{\text{PP}}(s) \leq v_t(s) \leq v_t^{\text{IP}}(s^p)$ and $v_t^{\text{IPP}}(s^p) \leq v_t^{\text{IP}}(s^p) \forall s, t$.*

5. Separable bounds

In this section, we provide separable upper and lower bounds for $v_t(s)$. The bounds are composed of value functions of several one-dimensional problems as described in Section 2. These bounds provide ingredients for computational approaches that take advantage of the relatively simple and nicely-structured solutions of one-dimensional problems.

For all i and t , let $\underline{P}_t^i(\cdot)$ and $\overline{P}_t^i(\cdot)$ be functions from \mathcal{R}_i to $[0, 1]$ that satisfy $\underline{P}_t^i(r^i) \leq P_t^i(r) \leq \overline{P}_t^i(r^i) \forall r$. In Section 7, we explain how to determine such $\underline{P}_t^i(\cdot)$ and $\overline{P}_t^i(\cdot)$ in certain situations. Our approach in the remainder of the present section is to use one-dimensional MDPs to generate lower [resp., upper] bounds for $v_t(s)$. To this end, consider n single flight (one-dimensional as in Section 2) problems indexed by i . Suppose the i th problem has capacity c^i , arrival probabilities $\{\lambda_i\}$, and

purchase probabilities $\{\underline{P}_t^i(r^i)\}$ [resp., $\{\bar{P}_t^i(r^i)\}$], and denote the value functions $\underline{v}_t^i(\cdot)$ [resp., $\bar{v}_t^i(\cdot)$]. The argument of $\underline{v}_t^i(\cdot)$ [resp., $\bar{v}_t^i(\cdot)$] is a scalar. We have

$$\underline{v}_t^i(x) = \max_{\rho \in \mathcal{R}_i(x)} \{\lambda_t \underline{P}_t^i(\rho)[\rho - \Delta \underline{v}_{t-1}^i(x)]\} + \underline{v}_{t-1}^i(x) \quad \forall t, x, \tag{6}$$

$$\bar{v}_t^i(x) = \max_{\rho \in \mathcal{R}_i(x)} \{\lambda_t \bar{P}_t^i(\rho)[\rho - \Delta \bar{v}_{t-1}^i(x)]\} + \bar{v}_{t-1}^i(x) \quad \forall t, x, \tag{7}$$

and $\underline{v}_0^i(x) = 0$ and $\bar{v}_0^i(x) = 0$.

Let $\{\underline{r}_t^i(x)\}$ be an optimal policy for the one-dimensional problem associated with $\underline{v}_t^i(\cdot)$ for each i . Let $\underline{r}_t(s) = (\underline{r}_t^1(s^1), \dots, \underline{r}_t^n(s^n))$, and let $\underline{\pi}$ be the policy that specifies action $\underline{r}_t(s)$ for the problem associated with $v_t(s)$ when the state at time t is s . Let $v_t^\pi(s)$ be the total expected revenue for periods $t, \dots, 1$ under policy $\underline{\pi}$ when the state is s at time t . We have the following proposition.

Proposition 4. $\sum_{i=1}^n \underline{v}_t^i(s^i) \leq v_t^\pi(s) \leq v_t(s) \leq \sum_{i=1}^n \bar{v}_t^i(s^i) \quad \forall s, t.$

For a given s , the relationship between the separable upper bound $\sum_{i=1}^n \bar{v}_t^i(s^i)$ and the inventory pooling bound $v_t^{\text{IP}}(s^p)$ depends upon problem specifics. When there is no customer choice among flights (i.e., $P_t^i(r)$ only depends on r^i for all i), the separable bounds are tight. On the other hand, if customers do not have preference on the flights but are shopping for price (e.g., $P_t^i(r) = 1$ if $i = \min\{j : j \in \arg \min_k r^k\}$ and $r^i \neq \rho_0$, and $P_t^i(r) = 0$ otherwise), then $v_t(s) = v_t^{\text{IP}}(s^p)$.

The separable upper [resp., lower] bound is composed of the value functions of n single-flight problems, each with one-dimensional state and action spaces. As discussed in Section 2, the optimal policy of a single-flight problem is of threshold-type. Therefore, only a few numbers need to be stored for each period to implement $\underline{\pi}$. As we will see in Section 6, certain other heuristic policies (based upon the separable bounds) for the multi-flight problem inherit a similar structure.

6. Value and policy approximation

Motivated by the analysis of the previous section, we next discuss various computationally feasible heuristic approaches for the multi-dimensional pricing problem of Section 3.

Let $\beta \in [0, 1]$. In value approximation, we approximate the value function $v_t(s)$ and choice probabilities $\{P_t^i(r)\}$ by, respectively

$$\tilde{v}_t(s) = \sum_{i=1}^n [\beta \bar{v}_t^i(s^i) + (1 - \beta) \underline{v}_t^i(s^i)] \quad \text{and} \tag{8}$$

$$\tilde{P}_t^i(r^i) = \beta \bar{P}_t^i(r^i) + (1 - \beta) \underline{P}_t^i(r^i) \quad \forall i. \tag{9}$$

From (8), it follows that $\Delta_i \tilde{v}_t(s) = \tilde{v}_t(s) - \tilde{v}_t(s - e^i) = \beta \Delta \bar{v}_t^i(s^i) + [1 - \beta] \Delta \underline{v}_t^i(s^i)$ for all i . Hence, the approximate marginal value of a seat on flight i is a weighted sum of the upper bound marginal value $\Delta \bar{v}_t^i(s^i)$ and lower bound marginal value $\Delta \underline{v}_t^i(s^i)$. An action is determined in state s at time t by solving

$$\max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n \tilde{P}_t^i(r^i) [r^i - \Delta_i \tilde{v}_{t-1}(s)] \right\} = \sum_{i=1}^n \max_{r^i \in \mathcal{R}_i(s^i)} \{ \lambda_t \tilde{P}_t^i(r^i) [r^i - \Delta_i \tilde{v}_{t-1}(s)] \}. \tag{10}$$

The maximization problem in (10) is motivated by the dynamic programming recursion (3).

It follows from the discussion at the end of Sections 2 and 5 that the policy on flight i as determined by solving (10) is of threshold type for each i . That is, there exists a set of thresholds $\{\tilde{I}_0^i(t), \dots, \tilde{I}_{k_i}^i(t)\}$ such that if the state vector is s at time t , the price on flight i is $\rho_j \in \mathcal{R}_i$ if $\tilde{I}_{j-1}^i(t) < s^i \leq \tilde{I}_j^i(t)$. Therefore, the policy on flight i in period t can be characterized by $k_i = |\mathcal{R}_i| - 1$ numbers (thresholds). Note that $\tilde{I}_0^i(t) = 0$ for all i and t . Consequently, to implement the policy from (10), the data storage requirements are quite manageable. For an n -flight problem, we need only store $\tau \sum_{i=1}^n k_i$ scalars (the thresholds) instead of $\tau \prod_{i=1}^n (c^i + 1)$ n -dimensional price vectors (the actions themselves) if no such structure exists.

The approach above is similar to the roll-out policies described in, e.g., Bertsekas and Tsitsiklis (1996). Applied to our context, a roll-out policy would select an action for state s at time t by maximizing the right side of (3), but with $v_{t-1}(\cdot)$ replaced by some approximation. In our approach, we also approximate the choice probabilities as in (9) to render the maximization tractable.

Similar to value approximation, we can obtain heuristics by using the optimal policies associated with the upper or the lower bound. By Proposition 4, the expected revenue from the policy $\underline{\pi}$ is at least as large as the lower bound. A potentially better heuristic can be obtained by “mixing” the policies of the upper and the lower bounds. To this end, let $(\bar{I}_0^i(t), \dots, \bar{I}_{k_i}^i(t))$ and $(\underline{I}_0^i(t), \dots, \underline{I}_{k_i}^i(t))$ be the thresholds in period t for $\bar{v}_t^i(\cdot)$ and $\underline{v}_t^i(\cdot)$, respectively. Given $0 \leq \beta \leq 1$, a threshold-type policy for flight i in period t can be determined by thresholds

$$([\beta \bar{T}_0^i(t) + (1 - \beta) \underline{T}_0^i(t)], \dots, [\beta \bar{T}_{k_i}^i(t) + (1 - \beta) \underline{T}_{k_i}^i(t)]), \tag{11}$$

where $\lfloor \cdot \rfloor$ denotes the integer part of a nonnegative real number.

The quality of the above methods depends on β . To find a good β , we use a search that varies β , determines the corresponding policies, evaluates the policies via simulation, and keeps the best one. The method is feasible since the simulations can be done quickly off-line (before the booking horizon), even for large problems. Below, we summarize the approach for value approximation (the procedure is essentially that described in Section 7.2 of Zhang and Cooper (2005)).

1. Initialize π^* and set $v^* = 0$.
2. Fix $0 \leq \underline{\beta} \leq \bar{\beta} \leq 1$ and $\delta > 0$. For $\beta = \underline{\beta}$ to $\bar{\beta}$ with step size δ do the following.
 - (a) Calculate $\tilde{v}_i(\cdot)$ and $\tilde{P}_i^i(\cdot)$ using (8) and (9), and plug the resulting values into (10) to obtain a corresponding policy π^β .
 - (b) Simulate the policy π^β for l replications and record the average total revenue from all the flights as \hat{v} . Then, \hat{v} is an estimator for $v_\tau^{\pi^\beta}(c)$.
 - (c) If $\hat{v} > v^*$, then $v^* = \hat{v}$ and $\pi^* = \pi^\beta$.

The policy π^* is the value-approximation policy. For policy approximation, replace step 2(a) by:

2(a') For $t = 1, \dots, \tau$, calculate thresholds using (11), and call the corresponding policy π^β .

7. Numerical experiments

In this section we describe our numerical study. To show the tightness of the bounds and the effectiveness of the heuristics, we test them on examples with two flights, and compare the simulated values with exact MDP values. We also consider examples with six and 12 flights to examine how the methods perform on relatively large problems. The price set is $\{\$150, \$200, \$250, \$300, \rho_0\}$ for each flight in all examples, where ρ_0 is the null price. Each flight has 80 seats unless noted otherwise. PA and VA are policy approximation and value approximation as discussed in Section 6. In the simulation step, we use $\underline{\beta} = 0, \bar{\beta} = 1, \delta = 0.1$, and $l = 100$. Because of the computational challenges mentioned in Section 4, we report results for IP only in the two- and six-flight examples and for PP only in the two-flight examples. In all tables, the values of MDP, LB, UB, PUB, and PP are computed by the backward induction algorithm for dynamic programs. Here, LB and UB are the separable lower and upper bounds, and PUB is the upper bound from inventory pooling. The expected revenue for each heuristic was estimated by an average over 5000 simulation runs for two-flight problems, and an average over 1000 simulation runs for other problems.

We used the following choice model. A customer arriving in period t assigns random utility vector $U_t = (U_t^1, \dots, U_t^n)$ to the n flights, where U_t is independent of everything else in the model. We suppress the t in the remainder of this section. For simplicity and to eliminate the need for “tie-breaking” in the choice process, we assume the distribution of U is such that $P(U^i - r^i \neq U^j - r^j) = 1$. The no-purchase utility is normalized to 0. Given price vector r , the consumer surplus for a customer is $U^i - r^i$ if the customer chooses flight i . We assume that customers are utility maximizers; that is, they will always choose an option with the highest consumer surplus. In period t , given there is an arrival and the price vector is r , the customer chooses flight i with probability

$$P_t^i(r) = P(U^i - r^i > (U^j - r^j)^+ \text{ for all } j : j \neq i). \tag{12}$$

In the above, we use the notation $(a)^+ = \max\{a, 0\}$. Mahajan and van Ryzin (2001) explain how many well-known choice models can be captured within this setup.

To apply the results of Section 5, let

$$\bar{P}_t^i(r^j) = P(U^i - r^j > 0), \tag{13}$$

$$\underline{P}_t^i(r^j) = P(U^i - r^j > (U^j - \rho_{k_j})^+ \text{ for all } j : j \neq i), \tag{14}$$

where ρ_{k_j} is the lowest price possible on flight j . It follows that $\underline{P}_t^i(r^j) \leq P_t^i(r) \leq \bar{P}_t^i(r^j)$, and hence we can apply Proposition 4 to bound the value function of the MDP. Also, note that the choice probabilities in (12) satisfy Condition S.

We used simulation to estimate $P_t^i(r)$ and $\underline{P}_t^i(r)$ in (12) and (14). Therefore, the values of MDP, LB, PUB, and PP may inherit errors from the simulation. We used two million samples of the utility vector U to ensure that the error is small. We used exact values of $\bar{P}_t^i(r^j)$ in (13), which we calculated directly from the marginal (one-dimensional) utility distributions.

7.1. Two-flight examples

In all two-flight examples, there are 1000 periods. We tested sets of examples with three different time-homogeneous arrival probabilities $\lambda_t = 0.3, 0.4$, and 0.5 . In all examples, the utility vectors are Normally distributed, but with negative values truncated to zero.

Table 1
Percentage difference from MDP value for two-flight examples where utility distribution is (truncated) bivariate Normal with mean (200, 200), standard deviation (100, 100), and correlation coefficient as specified in Correlation row

Correlation	Arrival rate								
	0.3			0.4			0.5		
	0.8	0	-0.8	0.8	0	-0.8	0.8	0	-0.8
MDP	33687.82	38747.57	40888.95	39249.94	42788.98	44243.12	42135.70	46015.50	47062.14
LB	-22.82%	-15.58%	-5.91%	-27.99%	-14.45%	-3.74%	-27.73%	-13.55%	-2.73%
UB	21.61%	5.73%	0.19%	12.81%	3.48%	0.08%	11.73%	2.31%	0.03%
PUB	0.34%	0.53%	0.49%	0.33%	0.32%	0.38%	0.26%	0.41%	0.64%
PP	-0.28%	-0.73%	-0.94%	-0.21%	-0.76%	-0.95%	-0.22%	-0.57%	-0.33%
IPP	-0.52%	-1.05%	-1.42%	-0.46%	-1.16%	-1.43%	-0.39%	-0.92%	-0.66%
IP	-0.52%	-1.15%	-2.40%	-0.46%	-4.86%	-4.15%	-0.41%	-0.92%	-0.74%
PA	-1.20%	-0.44%	-0.01%	-1.79%	-0.12%	-0.02%	-2.02%	-0.12%	-0.02%
VA	-0.72%	-0.14%	-0.01%	-0.30%	-0.10%	-0.08%	-0.42%	-0.04%	-0.02%

Table 1 reports results for cases where the utility distributions have mean (200, 200) and standard deviation (100, 100). The table shows MDP values, and for compactness expresses all other quantities in terms of percentage difference from the MDP value. Negative numbers represent quantities that are lower than the MDP value. Similar conventions are used in the other tables. For each arrival rate, we vary the correlation coefficient of the utility vector among 0.8, 0, and -0.8. Since the upper bound purchase probabilities depend only on marginal distributions, the UB values are the same for the same arrival rate. All the other expected revenue values are decreasing in correlation coefficient. This may be attributed to the fact that the choice probabilities are roughly decreasing in the correlation coefficient. The ordering on MDP values appears to reflect the results of Proposition 2, although we have no formal proof that its assumptions hold. On an intuitive level, if a customer assigns a low utility to one flight (and hence prefers not to purchase a ticket for it), he is likely to assign a high utility to the other (and therefore will purchase a ticket on the other flight), when the utility values are highly negatively correlated. On the other hand, when the utility values are highly positively correlated, if a customer assigns a low utility to one flight, he is likely to assign a low utility to the other also. Overall, in Table 1 all five heuristics perform well, yielding expected revenues within 2% of optimal in most cases.

Next, we report on cases where the utility distribution is not symmetric. Here, customers (as a whole) have a strong preference for one of the flights over the other. Table 2 reports results for cases where the utility distributions have mean (200, 180) and standard deviation (100, 80). Table 3 reports results for cases with mean (200, 160) and standard deviation (100, 60). A comparison of Tables 2 and 3 shows that PA and VA perform better than IPP and IP in most situations, with the performance gap increasing in asymmetry. The IPP policy performs poorly in almost all cases, with a revenue gap ranging from 2% to 19%. Interestingly, although the inventory pooling bound (PUB) is tighter than separable upper bounds in most cases, the policy implied by the bound (IP) does not perform very well as customer valuations become more asymmetric. The PP values are up to 10% below the MDP values. This shows that requiring the price to be the same for all flights can have a major impact on revenue, even if the best possible single-price policy is used.

Table 4 shows results for symmetric utility distributions and asymmetric capacity. The capacities of the two flights are 100 and 60, and the utility distributions are the same as those for Table 1. The results show that PA and VA perform better

Table 2
Percentage difference from MDP value for two-flight examples where utility distribution is (truncated) bivariate Normal with mean (200, 180), standard deviation (100, 80), and correlation coefficient as specified in Correlation row

Correlation	Arrival rate								
	0.3			0.4			0.5		
	0.8	0	-0.8	0.8	0	-0.8	0.8	0	-0.8
MDP	31826.14	36248.07	38097.73	36642.62	40073.83	41528.48	39339.74	43071.88	43852.99
LB	-18.79%	-13.89%	-4.69%	-22.44%	-11.85%	-3.63%	-25.71%	-12.92%	-1.85%
UB	19.98%	5.35%	0.23%	13.42%	3.71%	0.08%	11.50%	1.84%	0.02%
PUB	1.55%	0.53%	1.42%	0.87%	1.61%	0.82%	2.67%	0.56%	0.77%
PP	-0.83%	-1.76%	-0.83%	-1.99%	-0.02%	-1.20%	-3.45%	-6.11%	-5.93%
IPP	-4.77%	-7.52%	-6.58%	-5.72%	-6.55%	-7.91%	-4.96%	-8.74%	-9.35%
IP	-3.54%	-1.46%	-3.76%	-1.91%	-3.95%	-1.96%	-4.28%	-0.95%	-1.12%
PA	-0.69%	-0.18%	-0.03%	-1.97%	-0.13%	-0.02%	-1.13%	-0.04%	-0.01%
VA	-0.54%	-0.12%	-0.03%	-0.31%	-0.13%	-0.04%	-0.37%	-0.01%	-0.01%

Table 3

Percentage difference from MDP value for two-flight examples where utility distribution is (truncated) bivariate Normal with mean (200, 160), standard deviation (100, 60), and correlation coefficient as specified in Correlation row

Correlation	Arrival rate								
	0.3			0.4			0.5		
	0.8	0	-0.8	0.8	0	-0.8	0.8	0	-0.8
MDP	30282.96	33854.26	35816.90	34257.42	37562.28	38448.40	37188.34	39675.70	40307.05
LB	-28.18%	-10.06%	-4.71%	-28.24%	-10.95%	-1.83%	-26.44%	-10.45%	-1.08%
UB	18.55%	6.05%	0.23%	12.27%	2.39%	0.03%	8.42%	1.62%	0.03%
PUB	3.16%	1.82%	1.75%	5.11%	2.18%	1.80%	5.78%	2.89%	2.55%
PP	-4.57%	-5.68%	-6.56%	-3.89%	-6.05%	-5.86%	-7.59%	-9.43%	-9.51%
IPP	-10.75%	-14.66%	-16.90%	-11.60%	-15.19%	-17.32%	-12.11%	-16.61%	-18.57%
IP	-7.80%	-5.89%	-3.49%	-4.28%	-9.34%	-13.23%	-8.35%	-13.61%	-14.22%
PA	-0.53%	-0.13%	-0.05%	-0.26%	-0.12%	-0.00%	-0.98%	-0.01%	-0.00%
VA	-0.65%	-0.13%	-0.05%	-0.32%	-0.04%	-0.02%	-0.13%	-0.02%	-0.00%

Table 4

Percentage difference from MDP value for two-flight examples with asymmetric capacity of (100, 60)

Correlation	Arrival rate								
	0.3			0.4			0.5		
	0.8	0	-0.8	0.8	0	-0.8	0.8	0	-0.8
MDP	33466.69	37969.07	40213.50	38809.34	42713.48	43854.39	41969.84	45143.60	45710.57
LB	-22.38%	-15.86%	-7.45%	-27.42%	-15.28%	-3.70%	-28.94%	-13.75%	-2.26%
UB	20.48%	6.19%	0.27%	13.03%	2.70%	0.03%	8.92%	1.26%	0.01%
PUB	1.00%	2.59%	2.18%	1.47%	0.49%	1.27%	0.66%	2.35%	3.62%
PP	-1.51%	-3.06%	-3.79%	-0.89%	-3.73%	-3.72%	-1.13%	-1.61%	-0.72%
IPP	-2.19%	-4.19%	-5.95%	-1.35%	-5.31%	-6.12%	-1.76%	-3.20%	-3.13%
IP	-2.20%	-3.99%	-7.48%	-1.35%	-5.33%	-9.75%	-1.78%	-3.23%	-3.35%
PA	-1.20%	-0.24%	-0.02%	-1.49%	-0.25%	-0.07%	-1.74%	-0.01%	-0.05%
VA	-0.54%	-0.15%	-0.02%	-0.33%	-0.08%	-0.03%	-0.35%	-0.02%	-0.01%

The utility distribution is (truncated) bivariate Normal with mean (200, 200), standard deviation (100, 100), and correlation coefficient as specified in Correlation row.

Table 5

Percentage difference from best upper bound (the minimum of UB and PUB) for six-flight examples

	Arrival rate					
	0.4			0.6		
	A	B	C	A	B	C
Best UB	114783.00	121467.90	127691.70	127791.40	139178.00	144000.00
LB	-42.16%	-38.43%	-39.67%	-36.34%	-36.77%	-36.64%
UB	7.70%	10.70%	12.65%	-	-	-
PUB	-	-	-	6.21%	1.96%	0.00%
IPP	-18.41%	-10.72%	-2.74%	-17.00%	-9.47%	-0.02%
IP	-13.12%	-7.94%	-5.42%	-8.44%	-7.41%	-0.18%
PA	-5.43%	-6.28%	-6.18%	-2.96%	-2.94%	-0.01%
VA	-3.07%	-3.15%	-1.93%	-2.96%	-2.94%	-0.01%

The utility distribution for column A [respectively; B, C] is (truncated) multivariate Normal, with mean 200 [200, 200] on flights 1–3 and 160 [180, 200] on flights 4–6. The standard deviation is 100 [100, 100] for flights 1–3 and 60 [80, 100] for flights 4–6. The correlation coefficients are all 0.

than IPP and IP in all cases. Capacity asymmetry and utility distribution asymmetry apparently share a similar effect on the pooling heuristics: IP and IPP work better when there is low arrival rate, symmetric capacity, and symmetric utility distributions. PA and VA, on the other hand, appear to be robust with respect to parameter asymmetries, and work well in all cases. Also observe that with only capacity asymmetry, the PP value is up to 4% less than the MDP value, a significant revenue shortfall for many applications.

7.2. Six- and 12-flight examples

In six-flight examples there are 1500 periods. We tested a series of examples with arrival probabilities 0.4 and 0.6. Table 5 reports the results. In the table, columns A, B, and C correspond to the cases with highly-asymmetric, asymmetric, and symmetric utility distributions, respectively. The performance of IP and IPP deteriorates as the problems become more asymmetric, while the effects of asymmetry on PA and VA are less obvious. We are not able to compute MDP values because the state space is too large, so results are compared with the best upper bound, which is the minimum of the UB and PUB values. Note that VA appears to do better than PA. This observation also applies to 12-flight examples described below. For the cases with arrival rate 0.6, both PA and VA give the same policies – those associated with separable upper bounds. (When we use smaller stepsize, say, $\delta = 0.01$, the PA and VA policies differ, and give slightly different average revenues.)

To confirm our observations from the six-flight examples and to demonstrate the use of the heuristics for large problems, we consider several 12-flight examples with different arrival rates and utility distributions. The IP and PP heuristics, PUB, and MDP are not presented because of computational intractability. There are 2000 periods. The arrival probability in each period is either 0.6 or 0.9. Table 6 reports the results, which are similar to those in the six-flight examples. For the 12-flight examples, PA and VA again appear to perform well. The largest revenue gap is less than 9% from the separable upper bound. Comparing UB to PUB and MDP in the earlier tables suggests that much of the 9% gap may be explained by the fact that UB is a somewhat loose bound on MDP. Note that IPP performs poorly in some cases, with a maximum gap from UB of about 29%. Also, IPP performs worse when the utility distribution is not symmetric.

Tables 5 and 6 show that at the higher arrival rate, the policies PA and VA yield the same estimated expected revenue. This occurs because in both cases the algorithm described in Section 6 outputs a value of $\beta = 1$ (which corresponds to the policy associated with the separable upper bound) for both PA and VA, in which case the two policies are identical. (For a common choice of β , PA and VA differ from each other when $0 < \beta < 1$.) If we allow the algorithm to search over a finer grid of β values by taking (for instance) $\delta = 0.01$ rather than $\delta = 0.1$, then PA and VA will be different. To get a feel why $\beta = 1$ was chosen for both PA and VA in these problems with many flights and high arrival rate, note that the separable lower bound in these problems is quite loose (which makes low values of β worse), because it severely underestimates the choice probabilities for higher price vectors.

7.3. Periodic price changes

The heuristics PA and VA (as well as an optimal policy) prescribe prices for each state in each period. The policies require changing prices frequently, which is sometimes not desirable in applications. In this section, we consider policies where prices change only in pre-specified time periods. Bitran and Mondschein (1997) consider periodic pricing policies for a single type of item in fashion retailing, and show numerically that the revenue loss associated with such a policy is small, if the update interval is chosen appropriately. Our experiments suggest this is true in our context also.

We report results for two-flight examples with 1000 periods. We tested PA and VA with different price-change frequencies. We count setting the initial prices as one change, and the changes are distributed evenly through the horizon. For example, 40 changes means that the price is changed once every 25 periods. The prices are fixed between changes, unless the capacity of a flight is depleted, in which case the price on the flight is set to ρ_0 . To clarify, suppose β is fixed and $\{q_t(s)\}$ is the policy from VA (or PA). Suppose t' and t'' are two consecutive price-change times. If the state is s' at time t' , the price vector $q_{t'}(s')$ is used in periods $t', \dots, t'' - 1$, unless capacity is reached.

Good VA and PA pricing policies with periodic changes are determined by the simulation procedure described in Section 6, where in step 2(b), the policy π^β associated with a given β is a periodic pricing policy with a fixed price-change frequency. As before, the simulations are done off-line ahead of the selling horizon.

Table 6
Percentage difference from separable upper bound UB for 12-flight examples

	Arrival rate					
	0.6			0.9		
	A	B	C	A	B	C
UB	262576.70	286206.70	288000.00	265919.80	288000.00	288000.00
LB	-56.23%	-51.84%	-47.56%	-50.34%	-41.26%	-40.56%
IPP	-28.97%	-10.62%	-0.12%	-23.13%	-1.17%	-0.00%
PA	-7.94%	-8.49%	-0.06%	-1.29%	-0.48%	-0.00%
VA	-6.43%	-6.17%	-0.06%	-1.29%	-0.48%	-0.00%

The utility distribution for column A [respectively; B, C] is (truncated) multivariate normal, with mean 200 [200, 200] on flights 1–6 and 160 [180, 200] on flights 7–12. The standard deviation is 100 [100, 100] for flights 1–6 and 60 [80, 100] for flights 7–12. The correlation coefficients are all 0.

Table 7
The effect of price-change frequency

# changes	Arrival rate					
	0.3		0.4		0.5	
	PA	VA	PA	VA	PA	VA
1000	36205.76	36214.08	40007.26	40018.26	43053.31	43066.48
500	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%	-0.01%
200	-0.05%	-0.02%	-0.05%	-0.05%	-0.05%	-0.08%
100	-0.09%	-0.07%	-0.25%	-0.02%	-0.12%	-0.11%
50	-0.23%	-0.14%	-0.22%	-0.11%	-0.23%	-0.16%
40	-0.30%	-0.15%	-0.28%	-0.17%	-0.24%	-0.22%
25	-0.44%	-0.26%	-0.42%	-0.25%	-0.40%	-0.32%
20	-0.38%	-0.38%	-0.34%	-0.35%	-0.46%	-0.42%
10	-0.70%	-0.76%	-0.79%	-0.91%	-0.96%	-0.74%
4	-1.58%	-1.51%	-2.36%	-2.15%	-2.21%	-1.64%
2	-3.02%	-2.86%	-6.32%	-4.24%	-4.32%	-2.87%
1	-3.41%	-3.43%	-10.02%	-8.71%	-3.84%	-3.87%

The utility distribution is (truncated) bivariate Normal with mean (200, 180), standard deviation (100, 80), and correlation coefficient 0.

Table 7 shows the performance of these policies. The table shows that revenue loss is rather small, even with fairly infrequent price changes. If the price is changed more than 10 times, the revenue loss experienced is less than 1%. Note that the revenue loss is in terms of the best revenue achievable by PA (VA), which is shown to be very close to optimal in Section 7.1 (see also Table 2). It is most important to have more frequent price changes when the demand is moderate ($\lambda_t = 0.4$). When demand is low or high, even a fixed-price policy that does not change prices works reasonably well. In general, our experiments suggest that by making less-frequent price changes, an airline may be able to realize nearly all the benefit of the implementation of sophisticated pricing.

8. Summary

We developed a pricing model for substitutable flights where customers choose among the available flights. To overcome computational problems posed by the formulation’s multi-dimensional state and action spaces, we considered heuristics based on pooling ideas. We also derived easily-computable separable bounds for the value function of our model. Policies motivated by these bounds were shown numerically to be near optimal for a range of problem instances, and to dominate the policies from pooling in most cases. Our results suggest that pooling heuristics perform well for symmetric problems in which (a) customers, when viewed as a population, are mostly indifferent in their preferences over flights and (b) the flights have the same seating capacity. However, the pooling heuristics can perform poorly for asymmetric problems. The approaches motivated by the separable bounds do not suffer from such shortcomings, and remain implementable for large problems.

Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant Number DMI-0115385.

Appendix. Proofs

Proof of Proposition 1. Fix s and t , and let $\hat{r} = \hat{r}_t(s)$ be the maximizing action in (3). Let $A = \{i : \hat{r}^i < \Delta_i v_{t-1}(s)\}$, and $\tilde{r} = (\tilde{r}^1, \dots, \tilde{r}^n)$ where $\tilde{r}^i = \hat{r}^i$ if $i \notin A$ and $\tilde{r}^i = \rho_0$ otherwise. Denote by $\tilde{v}_t(s)$ the expected revenue when prices are \tilde{r} in period t and an optimal policy is used from $t - 1$ onward. Then

$$\tilde{v}_t(s) = \lambda_t \sum_{i=1}^n P_t^i(\tilde{r})[\tilde{r}^i - \Delta_i v_{t-1}(s)] + v_{t-1}(s) \geq \lambda_t \sum_{i=1}^n P_t^i(\hat{r})[\hat{r}^i - \Delta_i v_{t-1}(s)] + v_{t-1}(s) = v_t(s). \tag{15}$$

To see the inequality, note that $P_t^i(\tilde{r}) \geq P_t^i(\hat{r})$ for $i \notin A$ because P_t satisfies Condition S, and $\hat{r}^i < \Delta_i v_{t-1}(s)$ for $i \in A$. If there exists an $i \in A$ with $\lambda_t P_t^i(\tilde{r}) > 0$, then (15) is a strict inequality and hence $\tilde{v}_t(s) > v_t(s)$, contradicting the definition of $v_t(s)$. Therefore, no such i can exist. \square

Proof of Proposition 2. The proof is by induction on t . The statement in the proposition holds trivially for $t = 0$ by boundary conditions. Suppose the statement holds for $t - 1$.

Let $\tilde{v}_t(s)$ be the value function of an MDP with choice probabilities P_t, R_{t-1}, \dots, R_1 . We first show that $v_t^P(s) \geq \tilde{v}_t(s)$, and then show that $\tilde{v}_t(s) \geq v_t^R(s)$. By the inductive hypothesis, we have

$$v_t^P(s) = \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i + v_{t-1}^P(s - \epsilon^i)] + [1 - \lambda_t + \lambda_t P_t^0(r)] v_{t-1}^P(s) \right\} \\ \geq \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i + v_{t-1}^R(s - \epsilon^i)] + [1 - \lambda_t + \lambda_t P_t^0(r)] v_{t-1}^R(s) \right\} = \tilde{v}_t(s).$$

For fixed s , let \hat{r} be an optimal action in period t for the problem associated with $v_t^R(s)$. Then

$$v_t^R(s) = \lambda_t \sum_{i=1}^n R_t^i(\hat{r}) [\hat{r}^i - \Delta_i v_{t-1}^R(s)] + v_{t-1}^R(s) \leq \lambda_t \sum_{i=1}^n P_t^i(\hat{r}) [\hat{r}^i - \Delta_i v_{t-1}^R(s)] + v_{t-1}^R(s) \leq \tilde{v}_t(s).$$

In the above, the first inequality follows from Proposition 1 and the fact that $P_t^i(\hat{r}) \geq R_t^i(\hat{r})$. \square

Proof of Proposition 3. It is apparent that $v_t^{PP}(s) \leq v_t(s)$ and $v_t^{IPP}(s^p) \leq v_t^{IP}(s^p)$. It remains to prove $v_t(s) \leq v_t^{IP}(s^p)$. The statement is true for $t = 0$. Assume it holds for $t - 1$. Fix $s = (s^1, \dots, s^n)$ and let $s^p = \sum_{i=1}^n s^i$. By the inductive hypothesis and the fact that $\mathcal{R}(s) \subseteq \tilde{\mathcal{R}}(s^p)$, we have

$$v_t(s) = \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i + v_{t-1}(s - \epsilon^i)] + [1 - \lambda_t + \lambda_t P_t^0(r)] v_{t-1}(s) \right\} \\ \leq \max_{r \in \tilde{\mathcal{R}}(s^p)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i + v_{t-1}^{IP}(s^p - 1)] + [1 - \lambda_t + \lambda_t P_t^0(r)] v_{t-1}^{IP}(s^p) \right\} = v_t^{IP}(s^p).$$

This completes the proof. \square

Proof of Proposition 4. The second inequality follows from the definition of $v_t(s)$. We prove the first and third inequalities by induction. For $t = 0$, the inequalities hold. Suppose the first inequality holds for $t - 1$. Fix s and let $\underline{r} = \underline{r}_t(s)$. Then

$$v_t^{\pi}(s) = \lambda_t \sum_{i=1}^n P_t^i(\underline{r}) [\underline{r}^i + v_{t-1}^{\pi}(s - \epsilon^i)] + [1 - \lambda_t + \lambda_t P_t^0(\underline{r})] v_{t-1}^{\pi}(s) \\ \geq \lambda_t \sum_{i=1}^n P_t^i(\underline{r}) \underline{r}^i + \lambda_t \sum_{i=1}^n P_t^i(\underline{r}) \sum_{j \neq i} \underline{v}_{t-1}^j(s^j) + \lambda_t \sum_{i=1}^n P_t^i(\underline{r}) \underline{v}_{t-1}^i(s^i - 1) \\ + [1 - \lambda_t + \lambda_t P_t^0(\underline{r})] \sum_{i=1}^n \underline{v}_{t-1}^i(s^i) \\ = \lambda_t \sum_{i=1}^n P_t^i(\underline{r}) [\underline{r}^i - \Delta \underline{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \underline{v}_{t-1}^i(s^i) \tag{16}$$

$$\geq \lambda_t \sum_{i=1}^n P_t^i(\underline{r}^i) [\underline{r}^i - \Delta \underline{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \underline{v}_{t-1}^i(s^i) \tag{17} \\ = \sum_{i=1}^n \underline{v}_{t-1}^i(s^i).$$

The inequality (17) follows from the facts that $P_t^i(\underline{r}) \geq P_t^i(\underline{r}^i)$, and that $\underline{r}^i \geq \Delta \underline{v}_{t-1}^i(s^i)$ by (2).

Now we prove the third inequality in the proposition. Arguments like those that give (16) yield

$$v_t(s) \leq \max_{r \in \mathcal{R}(s)} \left\{ \lambda_t \sum_{i=1}^n P_t^i(r) [r^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i) \right\}. \tag{18}$$

Let $\tilde{r} = (\tilde{r}^1, \dots, \tilde{r}^n)$ be an action that maximizes the right side of (18). We have

$$v_t(s) \leq \lambda_t \sum_{i=1}^n P_t^i(\tilde{r}) [\tilde{r}^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i) \leq \lambda_t \sum_{i: \tilde{r}^i \geq \Delta \bar{v}_{t-1}^i(s^i)} P_t^i(\tilde{r}) [\tilde{r}^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i) \\ \leq \lambda_t \sum_{i: \tilde{r}^i \geq \Delta \bar{v}_{t-1}^i(s^i)} \bar{P}_t^i(\tilde{r}^i) [\tilde{r}^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i).$$

Let $\bar{r} = (\bar{r}^1, \dots, \bar{r}^n)$ where $\bar{r}^i \in \arg \max_{r^i \in \mathcal{R}_i(s^i)} \lambda_t \bar{P}_i^i(r^i)[r^i - \Delta \bar{v}_{t-1}^i(s^i)]$. From the definition of \bar{r} and the preceding inequality, we have

$$v_t(s) \leq \lambda_t \sum_{i: \bar{r}^i \geq \Delta \bar{v}_{t-1}^i(s^i)} \bar{P}_i^i(\bar{r}^i)[\bar{r}^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i) \leq \lambda_t \sum_{i=1}^n \bar{P}_i^i(\bar{r}^i)[\bar{r}^i - \Delta \bar{v}_{t-1}^i(s^i)] + \sum_{i=1}^n \bar{v}_{t-1}^i(s^i) = \sum_{i=1}^n \bar{v}_t^i(s^i).$$

Above, the second inequality follows from property (2) for one-flight problems. \square

References

- Algers, S., Beser, M., 2001. Modelling choice of flight and booking class – A study using stated preference and revealed preference data. *International Journal of Services Technology and Management* 2 (1/2), 28–45.
- Andersson, S., 1998. Passenger choice analysis for seat capacity control: A pilot project in Scandinavian Airlines. *International Transactions in Operational Research* 5 (6), 471–486.
- Ben-Akiva, M., Lerman, S., 1985. *Discrete Choice Analysis*. MIT Press, Cambridge, MA.
- Bertsekas, D.P., Tsitsiklis, J.N., 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.
- Bitran, G., Caldentey, R., 2003. An overview of pricing models for revenue management. *Manufacturing and Service Operations Management* 5 (3), 203–229.
- Bitran, G.R., Mondschein, S.V., 1997. Periodic pricing of seasonal products in retailing. *Management Science* 43 (1), 64–79.
- Bitran, G., Caldentey, R., Mondschein, S., 1998. Coordinating clearance markdown sales of seasonal products in retail chains. *Operations Research* 46 (5), 609–624.
- Boyd, A., Kallesten, R., 2004. The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management* 3 (2), 171–177.
- Carrier, E., 2003. Modeling airline passenger choice: Passenger preference for schedule in the passenger origin–destination simulator (PODS). Master's Thesis, Massachusetts Institute of Technology.
- Elmaghraby, W., Keskinocak, P., 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* 49 (10), 1287–1309.
- Gallego, G., van Ryzin, G., 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* 40 (8), 999–1020.
- Gallego, G., van Ryzin, G., 1997. A multi-product, multi-resource pricing problem and its applications to network yield management. *Operations Research* 45 (1), 24–41.
- Iyengar, G., Phillips, R., Gallego, G., Dubey, A., 2004. Managing flexible products on a network. CORC Technical Report TR-2004-01, Department of Industrial Engineering and Operations Research, Columbia University.
- Kleywegt, A.J., 2001. An optimal control problem of dynamic pricing. Working Paper, School of Industrial and Systems Engineering, Georgia Institute of Technology.
- Lin, K.Y., Li, F., 2004. Optimal dynamic pricing for a line of substitutable products. Working Paper, Grado Department of Industrial and Systems Engineering, Virginia Tech.
- Liu, B., Milner, J., 2006. Multiple-item dynamic pricing under a common pricing constraint. Working Paper, University of Toronto.
- Maglaras, C., Meissner, J., 2006. Dynamic pricing strategies for multi-product revenue management problems. *Manufacturing and Service Operations Management* 8 (2), 136–148.
- Mahajan, S., van Ryzin, G., 2001. Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49 (3), 334–351.
- McFadden, D., 2000. Disaggregate behavioral travel demand's RUM side: A 30-year retrospective. Technical Report, Department of Economics, University of California, Berkeley.
- Talluri, K., van Ryzin, G., 2004a. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50 (1), 15–33.
- Talluri, K., van Ryzin, G., 2004b. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston.
- Train, K.E., 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, UK.
- van Ryzin, G., Liu, Q., 2007. On the choice-based linear programming model for network revenue management. Forthcoming in *Manufacturing and Service Operations Management*.
- van Ryzin, G., Vulcano, G., 2006. Computing virtual nesting controls for network revenue management under customer choice behavior. Forthcoming in *Operations Research*.
- Zhang, D., Cooper, W.L., 2005. Revenue management for parallel flights with customer-choice behavior. *Operations Research* 53 (3), 415–431.
- Zhao, W., Zheng, Y.-S., 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Science* 46 (3), 375–388.