Joint network pricing and resource allocation

Dimitris Bertsimas * Sanne de Boer †

June 2002

Abstract

We study a joint pricing and resource allocation problem in a network with applications to production planning and airline revenue management. We show that the objective function reduces to a convex optimization problem for certain types of demand distributions, which is tractable for large instances. We propose several approaches for dynamic pricing and resource allocation. Numerical experiments suggest that coordination of pricing and resource allocation policies in a network while taking into account the uncertainty of demand can lead to significant revenue gains. Our strongest approach for the dynamic problem is an elegant combination of linear and dynamic programming.

1 Introduction

We study a joint pricing and resource allocation problem in a network with applications to production planning and airline revenue management. Suppose a company has a finite supply of resources which it can use to produce multiple products with different resource requirements. The demand for each product is uncertain and depends on its price. The company has to decide how to allocate its resources for production and how to price its finished products to maximize its expected revenue. The application of this problem that motivated this research is airline pricing, where the products are combinations of origin, destination and fare class, and the resources are seats on flights.

Our model extends the classical single-product, single-period newsvendor problem in three ways. 1) We study a multi-product resource allocation problem with capacity constraints, instead of a re-stocking decision for a single product. 2) Prices are included as decision variables, not given exogeneously. 3) We consider a multi-period planning horizon, which allows dynamic pricing. To highlight the contributions of this paper, we briefly review the relevant literature on airline revenue management, newsvendor problems and dynamic pricing.

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

†Operations Research Center, Massachusetts Institute of Technology, E40-130, Cambridge, MA 02139. sanne@mit.edu. Research partially supported by an MIT graduate student fellowship and the MIT-Singapore Alliance.

1.1 Literature review

Airline revenue management

The focus of the literature in airline revenue management is on seat inventory control for a given set of prices. Models to determine static booking limits for a single-leg flight were studied by Littlewood (1972), Belobaba (1987, 1989), Wollmer (1992), Brumelle and McGill (1993) and Robinson (1995). Continuous-time booking policies, where each booking request is accepted or rejected depending on the remaining capacity of the flight, were studied by Lee and Hersh (1993), Chatwin (1996, 1998), Lautenbacher and Stidham (1999) and Subramanian et al. (1999). These policies are typically based on a dynamic programming model that includes extensions of the basic problem such as overbooking, cancellations and no-shows. During the mid-eighties, American Airlines recognized the importance of network (origin-destination based) RM, aimed at getting the right mix of local and connecting traffic (Smith and Penn, 1988). Exact DP models are intractable in this case. Mathematical programming based heuristics for this problem were developed by Glover (1982), Wollmer (1986), Simpson (1989), Curry (1990), Williamson (1992), Chen et al. (1998), Talluri and van Ryzin (1998, 1999), Ciancimino et al. (1999), Bertsimas and Popescu (2000) and De Boer et al. (2002). Alternative solutions are proposed in van Ryzin (1998), Bratu (1999) and Bertsimas and de Boer (2002), who develop a simulation-based approximate dynamic programming method. Weatherford (1998) is a general discussion of airline revenue management. McGill and van Ryzin (1999) give an extensive literature review of this field that includes most of the work cited above.

To our knowledge, few papers address the joint pricing and seat allocation problem. Kuyumcu and Garcia-Diaz (2000) study the problem of an airline having defined a large number of booking classes with corresponding fares and restrictions, only a limited number of which can be put in the market due to CRS restrictions. In effect, this limits the pricing flexibility to a given discrete set of allowable prices. The authors propose a large-scale binary IP formulation of this problem. They develop a special-purpose algorithm to solve this model efficiently, but its performance on large practical instances is unclear.

Weatherford (1997) studies the joint pricing and inventory control problem for a single-leg flight with multiple fare classes. Demand is modeled by Gaussian distributed random variables whose expectations are functions of the fares. The author considers both partitioned inventory control with no diversion, and nested inventory control with and without diversion. For nested inventory control, the expected revenue as a function of the booking limits and prices can only be expressed analytically for small instances (2 or 3 classes) under the assumption that lower fare classes book first. The function has to be evaluated numerically and can be optimized using standard non-linear optimization algorithms. This approach does not scale well to problems with more fare classes or to a network environment. No structural results about the expected revenue function are given.

Newsvendor problem

The extension of the classical single-product, single-period newsvendor problem to a capacitated multiple-product inventory problem, aptlty named the *newsstand problem*, was studied by Lau and Lau (1995, 1996). They propose a non-linear optimization algorithm to determine the optimal single-period production plan for given prices and demand distributions. No structural results about the objective function of this problem are given.

The extension of the classical newsvendor problem with price as a decision variable was studied by Mills (1959, 1962), Karlin and Carr (1962), Zabel (1970), Young (1978), Lau and Lau (1988), Polatoglu (1991) and Dana and Petruzzi (2001). Petruzzi and Dada (1999) review the literature in this field including some extensions. Karlin and Carr (1962), Nevins (1966), Zabel (1972), Thomas (1974), Thowsen (1975), Petruzzi and Dada (1999) and Federgruen and Heching (1999) consider the single-product, multi-period combined pricing and inventory control problem, which is typically solved by dynamic programming. To our knowledge, the case of multiple products has not been studied within this framework.

Dynamic pricing

Gallego and van Ryzin (1994) and Zhao and Zheng (2000) model demand as a controlled continuous-time stochastic point process with intensity a known, decreasing function of the price. They consider the problem of optimally pricing a given inventory of a single product over a finite planning period before it perishes or is sold at salvage value. There are no re-ordering decisions. Chatwin (2000) considers the special case that the price has to be chosen from a finite set of allowable prices. Feng and Gallego (2000) extend the model by allowing the demand intensity to depend on sales-to-date. The model of You (1999) allows group bookings.

Gallego and van Ryzin (1997) and Paschalidis and Tsitsiklis (2000) extend this type of model to the dynamic pricing of multiple products whose production draws from a shared supply of resources. Gallego and van Ryzin consider a finite horizon joint pricing and resource allocation problem. The continuous-time pricing problem cannot be solved exactly, but the authors propose two heuristics based on a deterministic version of the problem that are asymptotically optimal as the scale of the problem (both expected demand and available resources) increases. Paschalidis and Tsitsiklis address the pricing of network services, which they formulate as a finite-state, continuous-time, infinite-horizon average reward problem. Similarly to Gallego and van Ryzin, the authors show that an asymptotically optimal pricing policy can be determined using a static, deterministic model. In Section 4, we will argue that this result may depend on demand following a Poisson process. Our periodical model allows more general types of demand distributions.

Kleywegt (2001) gives an optimal control formulation of the multi-product dynamic pricing problem. Kachani and Perakis (2002) propose a deterministic fluid model of dynamic pricing and inventory management for non-perishable products in capacitated and competitive make-to-stock manufacturing systems. The shortcoming of this type of approach is that the model ignores randomness, which as we shall see in Section 5.4 does affect the optimal pricing policy.

1.2 Contributions and outline

The main contribution of this work is the proposed joint pricing and resource allocation model in a network. To our knowledge, this is a new extension of the newsvendor model, which we believe has a promising application in airline pricing. We show that our model is convex for certain types of demand distributions thus tractable for large instances.

We test several heuristics to solve the dynamic pricing and resource allocation problem. The strongest approach is an elegant combination of linear and dynamic programming. We derive simple bounds on the optimal expected revenue of this problem that depend on the coefficient of variation of demand. Extensive simulation experiments show the value of centralized dynamic pricing based on a stochastic model of demand. In contrast, Gallego and van Ryzin (1997) and Paschalidis and Tsitsiklis (2000) found that a static pricing policy based on a deterministic model of demand is close to optimal. Our results suggest that this is only true when demand follows a Poisson process.

The outline of this paper is as follows. In Section 2 we present the single-period model and show that it can be reformulated as a convex optimization problem for certain types of demand distributions. In Section 3 we propose three heuristic multi-period extensions of this model that allow dynamic pricing. In Section 4 we derive bounds on the optimal expected revenue of the joint pricing and resource allocation problem that explicitly depend on the level of demand variability. In Section 5 we present numerical results that show the value of centralized dynamic pricing and compare the heuristics of Section 3. In Section 6 we investigate the impact of the inventory policy on the optimal prices. Section 7 summarizes our conclusions.

2 The Model

2.1 Assumptions

To model the single-period pricing and resource allocation problem, we make the following assumptions:

1. The production takes place at the beginning of the period.
2. Demand for each product is uncertain. Its expectation only depends on its price.
3. Excess demand is lost. Unsold products have no salvage value.

In Section 3 we consider multi-period planning horizons. This is an important extension, since dynamic pricing is an effective revenue management tool in many practical applications.

Assumption 1 models a make-to-stock inventory policy. This implies that actual sales do not depend on the order in which demand materializes, which allows a static formulation of the problem. The role of the inventory policy is further discussed in Section 6.

By Assumption 2, demand for a product is a random variable whose distribution only depends on its price. This simplification was necessary to show that our model has certain desirable properties that allow efficient solution methods. The model can easily be extended to include cross-elasticities of demand, but further research is required to find out the potential revenue gain and whether this model would be computationally tractable.

Assumption 3 applies to the airline pricing problem and simplifies the objective function. The model can easily be extended to incorporate penalties for shortages, salvage value of excess production or inventory costs, in which case the analysis carries through as in Section 2.3.

2.2 Model formulation

We now introduce some notation that is used throughout the paper. Let x_i and p_i be the production level and price of product i respectively. Let $\mu_i(p_i)$ be the expected demand for product i at this price. Let A_{ij} be the amount of resource j required for the production of one unit of product i. Let n be the number of products the business can produce and m be the number of resources it can utilize for this. The demand $D_i(p_i)$ for good i at price p_i is a random variable defined by

$$D_i(p_i) = \mu_i(p_i) Z_i \qquad (i = 1, ..., n) \tag{1}$$

for some continuous random variable $Z_i \geq 0$ with $E[Z_i] = 1$, $Var[Z_i] < \infty$ and pdf $f_i(z_i)$. These random variables are not required to be independent. This type of multiplicative model was orginally proposed by Karlin and Carr (1962) and has since been used frequently in the supply chain literature (e.g. Bernstein and Federgruen, 2002). Separating the price-dependent expected demand function $\mu_i(p_i)$ from the demand uncertainty modeled by Z_i significantly simplifies the analysis of the model. Note that (1) implies that the coefficient of variation of demand is independent of price.

The joint resource allocation and pricing problem outlined above is modeled by:

$$\begin{array}{ll} max_{\mathbf{x},\mathbf{p}} & E\left[\sum_i p_i \min\{x_i, D_i(p_i)\}\right] \\ s.t. & \mathbf{Ax} \leq \mathbf{c} \\ & \mathbf{x} \geq \mathbf{0} \end{array} \tag{2}$$

Here and in the remainder of this paper, bold face notation denotes vectors and matrices.

2.3 Structural results

In this section we investigate under which assumptions Model (2) is well-behaved, i.e. poses a concave optimization problem that can be solved efficiently. Since in practice problems of this type can be large, this is an important issue.

The objective function of Model (2) is separable and there are no constraints connecting the prices. As a result, the price of each product can be set separately and optimally as a function of its production level. Thus the price variables can be elimated from the optimization problem and the objective function can be reformulated as a (perhaps only implicit) function of the production level. This is a typical approach in newsvendor type pricing problems, cf. Petruzzi and Dada (1999). We will show that although the expected revenue may not be jointly concave in x and p, as suggested by the crossterm $x_i p_i$ in the objective function, the expected revenue at the optimal prices as a function of the production level can be shown to be concave under certain regularity conditions.

In the following, we assume that the expected demand functions $\mu_i(p_i)$ are non-negative, non-increasing and twice continuously differentiable, and that the *riskless profit functions* $r_i(p_i) = p_i \mu_i(p_i)$ are strictly concave (terminology by Mills, 1959). For example, these conditions are satisfied by the linear expected demand function $\mu_i(p_i) = a_i - b_i p_i$ for $a_i, b_i > 0$.

Because of the separable objective function, we can limit the analysis to the case of a single product and drop the index i. Let the expected revenue function $R(x,p)$ be defined as

$$\begin{aligned} R(x,p) &= pE[\min\{x, \mu(p)Z\}] \\ &= p\mu(p)E[\min\{y(x,p), Z\}] \end{aligned} \tag{3}$$

where $y(x,p) \triangleq x/\mu(p)$ is an auxiliary variable that simplifies notation. For convenience, the dependence of y on x and p is sometimes surpressed. We first need to show that the optimal price at a given production level is well-defined.

Theorem 1: For $x > 0$, the optimal price $p(x) = argmax_p R(x,p)$ is a continuously differentiable function of the production level x and the unique solution of

$$E[\min\{y, Z\}] = \frac{-p\mu'(p)}{\mu(p)} \int_0^y z f(z) dz \tag{4}$$

Proof: See Zabel (1970). ■

Optimality condition (4) can be solved numerically. Let p_∞ denote the price that maximizes $r(p)$, i.e. the optimal price in the case of infinite capacity. Then $p_\infty \leq p(x)$, which was first shown by Karlin and Carr (1962). Let $R(x) = R(x, p(x))$ be the expected revenue at the optimal price as a function of the production level. We now derive sufficient conditions for $R(x)$ to be concave.

Theorem 2: If $log\ f(z)$ is concave or if $Z \sim Lognormal$, then for $x > 0$, $R(x)$ is concave and

$$\frac{\delta p(x)}{\delta x} \leq 0$$

Proof: See Young (1978). ■

The sufficient condition of the theorem holds for a wide class of probability distributions, for instance the Uniform, Beta, Gamma, Weibull or truncated Normal distribution (e.g. Barlow and Proschan, 1996). Note that if Z were allowed to take negative values, even with small probability, the result would not hold. For very small x, a further decrease in production level may lead expected sales to become negative, thus requiring the optimal price to go down to stimulate demand. We were able to construct a numerical example with $Z \sim Normal$ for which $R(x)$ was not concave.

To conclude this section, we work out an example that illustrates our approach.

Example: Let $Z \sim \text{Uniform}(0,2)$ and let $\mu(p) = \alpha - \beta p$ for some $\alpha, \beta > 0$. Then $f(z) = \frac{1}{2}$ and $F(z) = \frac{z}{2}$ for $0 \leq z \leq 2$. Both are 0 otherwise. Note that $p_\infty = \frac{\alpha}{2\beta}$, thus $\mu(p_\infty) = \frac{\alpha}{2}$ is an upperbound on the expected demand. By (1) and $Z \leq 2$, we therefore have $D(p) \leq \alpha$ for all p, so we only have to consider $x \leq \alpha$. For $0 \leq y \leq 2$ we have

$$E\left[\min\{y, Z\}\right] = y - \frac{1}{4}y^2 \tag{5}$$

Note that for $y > 2$ we have $E\left[\min\{y, Z\}\right] = 1$ and $p = p_\infty$, which implies

$$y = \frac{x}{\mu(p_\infty)} \leq \frac{\alpha}{\alpha/2} = 2$$

Contradiction, so we don't have to consider this case. Using (5), optimality condition (4) becomes

$$y - \frac{1}{4}y^2 = \frac{\beta p}{\alpha - \beta p}\frac{1}{4}y^2$$

which is solved by

$$p(x) = \frac{\alpha}{\beta} - \frac{\sqrt{\alpha x}}{2\beta}$$

Thus $p(x)$ is indeed decreasing in x. Working out the equations, this gives

$$R(x) = \frac{x}{\beta}\left(\alpha - \sqrt{\alpha x} + \frac{x}{4}\right)$$

From this, it follows that

$$\frac{\partial R(x)}{\partial x} = \frac{1}{\beta}\left(\alpha - \frac{3}{2}\sqrt{\alpha x} + \frac{x}{2}\right) \quad \frac{\partial^2 R(x)}{\partial x^2} = \frac{1}{2\beta}\left(1 - \frac{3}{2}\sqrt{\frac{\alpha}{x}}\right) \leq 0$$

since $x \leq \alpha$. Thus $R(x)$ is indeed a concave function of x. It is now easy to calculate the Hessian of $-R(x,p)$ and find an example for which this matrix is not PSD, which shows $R(x,p)$ is not jointly concave in x and p. This example motivates the approach we have taken in this paper.

2.4 Solution algorithms

It is now clear how we can solve Problem (2) using an iterative non-linear optimization algorithm. First, the problem is reformulated as

$$\begin{array}{ll} max_{\mathbf{x}} & \sum_i R_i(x_i) \\ s.t. & \mathbf{A}\mathbf{x} \leq \mathbf{c} \\ & \mathbf{x} \geq \mathbf{0} \end{array} \tag{6}$$

Given a feasible allocation $\mathbf{x} = (x_1, ..., x_n)$ we can calculate the optimal price vector $\mathbf{p}(\mathbf{x}) = (p_1(x_1), ..., p_n(x_n))$ by solving optimality condition (4) for all products $i = 1, ..., n$. From this, the first and the second order derivative of the objective function at this point follow (e.g. Young, 1978). This information can be used to determine a new and improved feasible allocation, for instance using the Frank-Wolfe method (e.g. Bertsekas, 1999). These steps are repeated until no significant further improvement is made, at which point the algorithm terminates. Given the conditions of Theorem 2, this will be the optimal production plan, from which the optimal prices follow.

For reasons that will become clear in Section 3, in the numerical experiments of Section 5 we have used a simpler approach. In all our test cases $Z \sim Erlang$, thus the functions $R_i(x_i)$ in the objective function of (6) are concave. We have approximated these by piecewise linear functions defined by a finite set of discretization points. This is possible because the domain of the functions $R_i(x_i)$ is bounded by the resource and non-negativity constraints. It is easy to see that using this piecewise linear approximation of the concave objective function, Problem (6) can be rewritten as a linear optimization problem. Both the accuracy and the number of variables of this approximation depend on the number of discretization points. We ran some numerical tests to make sure that our approximation was sufficiently accurate. Further increasing the number of discretization points did not significantly affect the optimal objective function value.

To conclude this section, we give some results that can speed up the numerical experiments of Section 5. If $Z \sim Erlang(k, 1/k)$, it is easy to see that

$$\int_0^y zf(z)dz = P(Y \geq k+1)$$

for some random variable $Y \sim Poisson(yk)$. Thus

$$\begin{aligned} E[\min\{y, Z\}] &= \int_0^y zf(z)dz + yP(Z \geq y) \\ &= P(Y \geq k+1) + yP(Y \leq k-1) \end{aligned}$$

These results are useful for solving (4) numerically, e.g. by the bi-section method.

3 The multi-period problem

In this section we consider a multi-period extension of the pricing and resource allocation problem in a network, where the (finite) planning horizon is divided into T periods. At the beginning of each period, the company needs to decide its prices and production levels given the remaining supply of resources. We assume that the company implements a make-to-order inventory policy. To optimally allocate its resources, the company limits how much of each good it is willing to sell during each period, but the supply of resources is only affected by actual sales. As a result, there is no inventory of finished products.

In the following we use the notation from Section 2.2, with the convention that the index t refers to the situation at the beginning of the period. For $t = 1, \ldots, T$, the dynamic programming formulation of the problem outlined above is:

$$\begin{aligned} J^t(\mathbf{c}^t) = \quad & max_{\mathbf{x}^t, \mathbf{p}^t} E\left[\sum_{i=1}^{n} min\{D_i^t(p_i^t), x_i^t\}p_i^t + J^{t+1}(\mathbf{c}^t - \mathbf{A} * \mathbf{min}\{\mathbf{D}^t(\mathbf{p}^t), \mathbf{x}^t\})\right] \\ s.t. \quad & \mathbf{A}\mathbf{x}^t \leq \mathbf{c}^t \\ & \mathbf{x}^t, \mathbf{p}^t \geq \mathbf{0} \end{aligned} \tag{7}$$

subject to the initial condition

$$J^{T+1}(\mathbf{c}^{T+1}) = 0$$

In principle, this model can be used to determine the optimal dynamic pricing and resource allocation policy. However, the state vector of the model is the remaining supply of each of the resources. Since the number of resources may be large, the curse of dimensionality of dynamic programming renders this approach numerically intractable for practical purposes. We therefore need to develop heuristic solutions. In this section we consider three such methods, which are all based on solving a mathematical programming problem similar to (6). In our opinion, this is the only tractable way to solve the resource allocation issue that is inherent to the pricing problem. All methods are understood to be implemented on a rolling-horizon basis, i.e. the model is re-optimized at the beginning of each period given the remaining supply of resources.

A note on dynamic pricing

For a better understanding of the performance of each of these heuristics, we need to understand how dynamic pricing can increase revenues. First, price changes can be used to compensate for normal statistical fluctuations of demand. For instance, when sales are behind expectations the company may lower its prices to stimulate demand. Second, when the willingness to pay of its customers changes over time, the company should charge higher prices in periods of price-inelastic demand to maximize its profits. For instance, some people are willing to pay a premium in the beginning of the fashion season to wear the latest fashion, while others prefer to wait until an end-of-season sale. Zhao and Zheng (2000) showed that although the second effect is more important than the first, its impact on expected revenue can still be significant.

3.1 The Demand Aggregation Method

The first method is to calculate the optimal fixed price and resource allocation policy for the remaining part of the planning horizon. Let

$$D^t(p^t) = \mu^t(p^t) * Z^t \qquad (t = 1, ..., T)$$

for some random variables Z^t with $E[Z^t] = 1$ and $Var[Z^t] = \sigma_z^2$. Assume that $\mu^t(p) = \mu(p)$ and $Z^t \sim Z$ for all $t = 1, ..., T$, i.e. demand is time-homogeneous. Let $\overline{D}^t(p)$ be the aggregate demand over periods t to T at a constant price p. Then it is easy to see that

$$\begin{aligned} \overline{D}^t(p) &= \mu(p) \sum_{i=t}^{T} Z^i \sim \overline{\mu}^t(p) \overline{Z}^t \qquad (8) \\ \overline{\mu}^t(p) &= \sum_{i=t}^{T} \mu^i(p) = (T - t + 1)\mu(p) \end{aligned}$$

for some random variable $\overline{Z}^t$ with $E[\overline{Z}^t] = 1$ and $Var[\overline{Z}^t] = \sigma_z^2/(T - t + 1)$. Hence the aggregate demand to come fits into the framework of Model (1). The random variables $\overline{D}^t(p)$ can be used as input for Model (2) to obtain an approximate solution $\mathbf{x}^t$ and $\mathbf{p}^t$ to the Bellman equation (7). We refer to this heuristic as the Demand Aggregation Method (DAM).

This approach is less suitable when demand is not homogeneous over time, since the aggregate demand distribution may not fit exactly into the framework of Section 2.2. There may not be an aggregate expected demand function $\overline{\mu}^t(p)$ and a random variable $\overline{Z}^t$ with unit expectation and variance independent of p such that (8) holds. The expectations of the random variables on either side of (8) can easily be matched, but for the variance this may be impossible to do exactly. However, this is still a reasonable approximation to the problem when a constant price over time is indeed close to optimal. A more fundamental issue arises when the price elasticity of demand varies over time. DAM will lead to a price in period t that is presumably a good fixed price over the remaining periods. However, this is not necessarily the best price for period t if the price can be adjusted later on. We therefore expect DAM to perform poorly in this case.

3.2 The Demand Disaggregation Method

The second method is to determine the production level and price (x_i^k, p_i^k) for all periods to come $(k = t, ..., T)$ at the beginning of period t. To simplify the model, we assume that the production x_i^k is used exclusively to satisfy the demand in period k that materializes at price p_i^k. This problem can be formulated as:

$$\begin{array}{lll} max_{\mathbf{x},\mathbf{p}} & E\left[\sum\limits_{k=t}^{T}\sum\limits_{i=1}^{n} p_i \min\{x_i^k, D_i^k(p_i^k)\}\right] & \\ s.t. & \mathbf{A}\left(\sum\limits_{k=t}^{T}\mathbf{x}^k\right) \leq \mathbf{c}^t & \\ & \mathbf{x}^k \geq \mathbf{0} & \end{array} \tag{9}$$

The solution of (9) can be transformed to a heuristic solution $(\overline{x}_i^k, \overline{p}_i^k)$ of the Bellman equation (7) by taking

$$\overline{x}_i^t = \sum_{k=t}^{T} x_i^k \quad \overline{p}_i^t = p_i^t \qquad (i = 1, \ldots, n)$$

In the following, we refer to this heuristic as the Demand Disaggregation Method (DDM). Note that (9) is an extension of Model (2) with more decision variables, whereas for DAM only the model input is affected in the form of the demand distributions.

The strength of DDM is that this method explicitly takes into account that the company can set different prices for the same product in different periods, so that it can benefit from a time-varying willingness to pay by dynamic pricing. In that case, we expect DDM to outperform DAM. Its primary weakness is that the underlying model assumes that any unsold production for period t cannot be used to satisfy excess demand in period $t + 1$. Unlike in the DP formulation, unused resources do not carry over to the next period. We feel that this will lead to suboptimal production levels, and we therefore expect DAM to outperform DDM when the price elasticity of demand is constant over time and the effect of dynamic pricing is second-order.

3.3 Dynamic Programming Based Mathematical Programming

It is important to note that although both DAM and DDM are meant to be used dynamically, i.e. the prices and resource allocation are recalculated at the beginning of each period given the state of the system, the underlying models assume that the resulting policy is static. They do not take into account that the decisions can periodically be adjusted in response to how demand materializes. This may lead to overly conservative pricing decisions. Dynamic programming is the only optimization technique that incorporates this effect but due to the curse of dimensionality this approach is intractable when the dimension of the state space is even moderately sized. We are therefore looking for a way to combine the efficient mathematical programming solution of the resource allocation problem with a truly dynamic pricing model.

Following an idea succesfully applied in Bertsimas et al. (2001), we propose the following heuristic that we will henceforth refer to as the Dynamic Programming based Method (DPM). At the beginning of period t, fix the total production level x_i^t of each product, but allow dynamic pricing over the remaining part of the planning horizon. The expected revenue of each product under the optimal dynamic pricing policy can be calculated by solving the following dynamic programming model:

$$J_i^t(x_i^t) = max_{p_i^t \geq 0} E\left[p_i^t min\{D_t(p_i^t), x_i^t\} + J_i^{t+1}(max\{x_i^t - D_i^t(p_i^t), 0\})\right] \tag{10}$$

subject to the initial condition

$$J_i^{T+1}(x_i^{T+1}) = 0$$

The state is one-dimensional, thus this problem is tractable. The initial condition implies that the product has no salvage value at the end of the planning horizon, which is what we have assumed throughout this paper. Note that the state-space is continuous, hence the value functions $J_i^t(x_i^t)$ cannot be calculated exactly. We have used a piecewise linear approximation instead.

We can then obtain a heuristic solution to the Bellman equation (7) by solving

$$\begin{array}{ll} \max & \sum_{i=1}^{n} J_i^t(x_i^t) \\ \text{s.t.} & \mathbf{A}\mathbf{x}^t \leq \mathbf{c}^t \\ & \mathbf{x}^t \geq \mathbf{0} \end{array} \tag{11}$$

and letting

$$p_i^t = \arg\max_{p_i^t \geq 0} E\left[p_i^t min\{D_t^i(p_i^t), x_i^t\} + J_i^{t+1}(max\{x_i^t - D_t(p_i^t), 0\})\right] \qquad (i = 1, \ldots, n)$$

Given the piecewise linear approximation of the value functions, it is well known (e.g. Bertsimas and Tsitsiklis, 1997) that Problem (11) can be reformulated as a mixed integer program with binary variables. If the value functions are concave, Problem (11) can be reformulated as an ordinary linear optimization problem that can be solved efficiently. In our numerical experiments in Section 5, this always turned out to be the case. For the continuous-time dynamic pricing problem with Poisson demand, Gallego and van Ryzin (1994) were able to prove concavity of the value function under certain regularity conditions. Although thusfar we have not been able to prove this result for the multiplicative demand model, these findings suggest that it may very well hold in this case as well.

We feel that the strenght of DPM is that it explicitly incorporates the dynamic pricing decisions in a network setting. It may capture a significant part of the effect that potential price changes down the line should have on the current pricing decision. We expect this to be most important when the price elasticity of demand varies over time, in which case the price changes will be significant. However, the method does not take into account that the resource allocation can be updated in each period as well. The negative effect of this on the performance of DPM should be more significant for larger problems with many types of products. A possible improvement of DPM that we have not further considered here is to formulate the value functions in terms of the aggregate production level of all products with the same resource requirements. For instance, in the airline setting this would be all applicable fares for a particular itinerary. The state of the DP would still be one-dimensional, but now both the prices and the exact use of the resources allocated to this group of products can be updated periodically.

Note that the expected revenue function $R(x)$ of Section 2.3 can actually be seen as the value function of a single-period DP problem. Although it can be evaluated numerically at every point of its domain, we have chosen to work with a piecewise linear approximation of this function as well. Since this function is provably concave, the approximation can be taken arbitrarily close. As we have pointed out, Problem (6) can then be solved as a linear optimization problem. More importantly, the single-period Problem (6) and the multi-period heuristics DAM, DDM and DPM then all fit into the framework of Problem (11) with a piecewise linear objective function. This facilitates the development of generic computer software to solve both the static and the dynamic pricing and resource allocation problem.

4 Some theoretical bounds

In this section,we derive lower and upper bounds on the optimal expected revenue of the multi-period pricing and resource allocation model. The lower bound explicitly depends on the level of demand variability. These bounds will show the intuitive result that when the demand variability decreases, the expected revenue of the optimal dynamic policy can be approximated with increasing accuracy by a static policy based on a deterministic model of demand. Based on this result, we argue that the finding by Gallego and van Ryzin (1997) and Paschalidis and Tsitsiklis (2000) that the latter policy becomes close to optimal when the scale of the problem increases depends on demand being modeled by a Poisson process.

In the following, we assume that in each period t ($t = 1, ..., T$) and for each product i ($i = 1, ..., n$) we have a random demand $D_i^t(p_i^t)$ whose distribution depends on the price p_i^t. The results in this section do not require the multiplicative demand model (1), unless explicitly stated otherwise. Let

$\mu_i^t(p_i^t)$ be the expectation of $D_i^t(p_i^t)$ and let $\sigma_i^t(p_i^t)$ denote its standard deviation. Let $\rho_i^t(p_i^t) = \frac{\sigma_i^t(p_i^t)}{\mu_i^t(p_i^t)}$ be the coefficient of variation of demand. For the moment, we assume that we have a partitioned inventory control policy, i.e. we have allocations x_i^t which for each period t is the maximum amount of product i that we are willing to sell. A joint pricing and inventory control policy $\{x_i^t, p_i^t\}$ is considered to be static if it is determined at the beginning of the first period and not updated afterwards. A policy is considered to be dynamic if x_i^t and p_i^t are determined at the beginning of period t as a function of the remaining supply of resources.

Theorem 3: Let

$$\begin{array}{rll} J_{LB} = & max_{\{p_i^t\}} & \sum_{t,i} p_i^t \mu_i^t(p_i^t)\left(1 - \frac{\rho_i^t(p_i^t)}{2}\right) \\ & s.t. & \mathbf{A}\left(\sum \mu^t(p^t)\right) \leq \mathbf{c} \\ & & \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{array} \tag{12}$$

then

$$J_{LB} \leq J^*_{static} \leq J^*_{dynamic}$$

where J^*_{static} and $J^*_{dynamic}$ denote the expected revenue of the optimal static and dynamic joint pricing and inventory control policy respectively.

Proof: See appendix. ■

When $D_i^t(p_i^t) \geq 0$ and $\rho_i^t(p_i^t) \geq 1$ for some product i and period t, the lower bound can be strengthened by replacing the corresponding terms in the objective function of Model (12) by $p_i^t \mu_i^t(p_i^t)\left(1 - \frac{\rho_i^t(p_i^t)^2}{1+\rho_i^t(p_i^t)^2}\right)$, using Theorem 3 of Bertsimas and Popescu (1999). However, this would not lead to any additional insights.

The tightness of the bound depends on the coefficients of variation of demand. It may also depend on the scale of the problem, that can be changed by multiplying both the expected demand functions μ and the supply of resources c by some scaling factor α. For instance, for the additive model, $D_i^t(p_i^t) = \mu_i^t(p_i^t) + Z_i^t$ for some random variable Z with expectation 0 and standard deviation $\sigma_z^{i,t}$. In that case, $\rho_i^t(p_i^t) = \frac{\sigma_z^{i,t}}{\mu_i^t(p_i^t)}$, thus the lowerbound becomes more tight when the scale of the problem increases. For the multiplicative model, $D_i^t(p_i^t) = \mu_i^t(p_i^t) Z_i^t$ for some random variable Z with expectation 1 and standard deviation $\sigma_z^{i,t}$. The coefficients of variation of demand in this model do not depend on the scale, since $\rho_i^t(p_i^t) = \sigma_z^{i,t}$. An especially nice case is $\sigma_z^{i,t} = \sigma_z$, when the demand variation only occurs in the objective function as the multiplication factor $(1 - \sigma_z/2)$.

Theorem 4: Let $\mathbf{Z}$ be the vector of random variables that govern all the randomness in the problem, i.e. we assume $D_i^t(p_i^t) = D_i^t(p_i^t, Z_i^t)$ for some given functions $D_i^t(p_i^t, Z_i^t)$ that are non-increasing in p_i^t. Let $J^*_{dynamic}(\mathbf{z})$ be the revenue generated by the optimal dynamic pricing and resource allocation policy when $\mathbf{Z} = \mathbf{z}$. Define

$$\begin{array}{rll} F(\mathbf{z}) = & max_{\{p_i^t\}} & \sum_{t,i} p_i^t D_i^t(p_i^t, z_i^t) \\ & s.t. & \mathbf{A}\left(\sum D^t(p^t, z^t)\right) \leq \mathbf{c} \\ & & \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{array}$$

then

$$J^*_{dynamic} = E\left[J^*_{dynamic}(\mathbf{z})\right] \leq E\left[F(\mathbf{z})\right]$$

Proof: See appendix. ■

Note that Theorem 4 proves the intuitive result that if we have perfect foresight of the randomness in the problem, we can always do better than any non-anticipatory policy. Also note that the result did not depend on the inventory control policy that was applied, as long as it does not violate the resource constraints.

Theorem 5: If $F(z)$ is a concave function of z, then

$$J^*_{dynamic} \leq F(E\left[\mathbf{z}\right])$$

Proof: The result follows directly from Theorem 4 and Jensen's inequality. ■

Theorem 6: For multiplicative demand functions, i.e. $D^t_i(p^t_i, z^t_i) = \mu(p^t_i)z^t_i$, $F(\mathbf{z})$ is concave in $\mathbf{z}$ if the expected demand functions $\mu(\cdot)$ are convex in p.

Proof: See appendix. ■

Theorem 7: Let

$$\begin{array}{lll} J_{UB} = & max_{\{p^t_i\}} & \sum_{t,i} p^t_i \mu^t_i(p^t_i) \\ & s.t. & \mathbf{A}\left(\sum \mu^t(p^t)\right) \leq \mathbf{c} \\ & & \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{array}$$

Then for the multiplicative demand model with convex expected demand functions,

$$J_{LB} \leq J^*_{static} \leq J^*_{dynamic} \leq J_{UB}$$

and

$$\lim_{\{\rho^t_i \to 0\ \forall i,t\}} J_{LB} = \lim_{\{\rho^t_i \to 0\ \forall i,t\}} J^*_{static} = \lim_{\{\rho^t_i \to 0\ \forall i,t\}} J^*_{dynamic} = J_{UB}$$

Proof: The first result follows directly from Theorems 3, 5 and 6. The second result follows from the definition of J_{LB} and from J_{UB} being independent of the level of demand variability. ■

This theorem shows the intuitive result that if the randomness is taken out of the problem, the optimal dynamic policy can be approximated by a static policy based on a deterministic model of demand. Note that none of the proofs of these theorems depends on the inventory control policy of the dynamic model, as long as it does not violate any of the resource constraints. This means that the theorem also holds when we consider the optimal expected revenue of a make-to-order inventory policy by a manufacturer or of nested seat inventory control by an airline.

As we have seen, for the multiplicative demand model these bounds do not depend on the scale of the problem. This suggests that a policy based on a stochastic model will outperform a policy based on a deterministic model and that a static policy will be outperformed by a dynamic policy, even when the scale of the problem increases. This was confirmed by numerical experiments reported in Section 5. By contrast, Gallego and van Ryzin (1997) and Paschalidis and Tsitsiklis (2000) find that a static pricing policy based on a deterministic model of demand is asymptotically optimal. Our bounds offer an intuitive explanation for this apparent contradiction. Under Poisson arrivals, we know that $\rho_i^t(p_i^t) = \frac{1}{\sqrt{\mu_i^t(p_i^t)}}$. Thus intuitively, the lowerbound becomes more tight at rate $\sqrt{\alpha}$. Since the deterministic model gives an upperbound in this case as well, as shown by Gallego and van Ryzin, the solutions will converge when the scale of the problem increases. In a sense, this scaling takes out the randomness of the problem. The result above shows that this result does not hold in general and depends on the demand distribution.

5 Numerical experiments

5.1 Problem generation

Since we did not have real data to run numerical experiments, we have written a computer program that automatically generates test problems according to prespecified demand, network and fare settings. The examples were intended to be realistic and allow sensitivity analysis with respect to the network load factor, demand variability and other factors that may affect the simulation results. The basic features of the problem generation program are presented below.

All test cases are motivated by the airline revenue management problem. We have only considered single-hub networks. The number of spokes N, the number of fare classes F and number of periods T can be varied. The airline offers tickets from each spoke to the hub and to every other spoke and from the hub to each spoke. All trips are one-way. We only consider one bank of flights from the spokes into the hub and one connecting bank from the hub back to the spokes.

At the beginning of each period, the prices and resource allocation are reoptimized. Demand for each product in each period is independent and modeled by the Erlang distribution with a constant coefficient of variation. This parameter allows changing the overall demand variability in the network. The default setting was $CV^2 = 33\%$. The expected demand functions are linear in the prices.

Since price is a decision variable we cannot determine the network demand factor beforehand. We therefore introduce the notion of DF_∞, which we define as the network demand factor with all prices set to p_∞. DF_∞ is an input parameter for the problem generation program, which then calculates the demand levels at p_∞ on each leg and sets the leg capacities accordingly. The default setting was $DF_\infty = 2$.

The demand parameters are the same for each fare product in the same period, except for small random perturbations that we used to break the symmetry. We only distinguish fare products by their number of legs and by their fare class. The parameters are based on values used by Weatherford (1997) based on his experience with the airline industry.

In a realistic airline network, only about 30% of the passengers on a leg are local while the others connect at the hub. The average fare for a connecting flight is typically higher. If there are multiple fare classes, demand for the highest fare class is the least price-sensitive. Typically, the demand for the higher fare classes is also lower than the demand for the lower fare classes. All this was modeled as well, again for p_∞.

The time-dimension is added in three different ways designed to identify which of the heuristics of Section 3 should be used when. The first case is that of time-homogeneous demand (HD), where both the demand intensity and the price elasticity are constant over time. The second case (NHD, $F > 1$) applies to an airline segmenting the market by offering multiple well-differentiated fare classes. We model that the lower fare classes tend to book first, i.e. the higher the fare class, the later the peak in demand intensity. This reflects that price-inelastic business travelers tend to book later than price-elastic leisure travelers, as is the case in practice. As a result demand is non-homogeneous, but we assume that the price elasticity of demand for each fare class is constant over time. The third case (NHD, $F = 1$) applies to an airline that at any given time only offers one fare. Demand is non-homogeneous since its price elasticity decreases over time as business travelers start making up a larger part of the market. In response, the fare should go up significantly when getting close to flight departure.

5.2 The value of centralized pricing

In this section we show that the gain of a centralized pricing policy, i.e. optimized over the network as a whole rather than for each product individually, can be significant. We consider two single-period test problems for which the optimal prices and resource allocation can be determined using Model (6). The results are compared to the suboptimal pricing policy p_∞. *Given these prices*, the resource allocation is optimized by solving a stochastic programming model similar to (6). Here as well as in the following sections, all comparisons are based on 1000 iterations of a demand simulation program that approximates the expected revenue of any given pricing and resource allocation policy.

Our first example shows that as one would expect, the gain of optimal pricing increases with the network demand factor. The results for $N = 5$ and $F = 2$ are reported in Table 1. We have repeated this experiment for $N = 10$ and $N = 15$ but the results were very similar, presumably because of the symmetric problem generation method.

Table 1: Revenue gain optimal over benchmark pricing policy ($N = 5, F = 2, T = 1$)

DF_∞	0.50	0.85	1.20	1.55	1.90	2.25	2.60	2.95	3.30	3.65	4.00
LF	0.46	0.62	0.71	0.75	0.79	0.81	0.82	0.84	0.84	0.86	0.86
Gain	0.4%	4.0%	9.4%	14.3%	18.1%	20.5%	21.9%	22.2%	22.3%	21.9%	22.0%

Our second example shows that the benefit of optimal pricing decreases with the demand uncertainty in the network. The results for $N = 5$ and $F = 2$ are given in Table 2. The explanation may be that when there is more uncertainty, demand is less likely to be at the level for which the price that maximizes the *expected* revenue is indeed optimal.

Table 2: Revenue gain optimal over benchmark pricing policy ($N = 5, F = 2, T = 1$)

CV^2 (1/k)	0.10	0.14	0.17	0.20	0.25	0.33	0.50	1.00
Gain	21.4%	20.9%	20.4%	20%	19.6%	19.0%	18.1%	17.1%

The benchmark pricing policy p_∞ is not a good one, since we know it systematically sets the prices too high. However, that may be similar to what is going on in practice when each point of sale (say each spoke) can set its own profit-maximizing prices without regard to the profit maximizing pricing strategy for the network as a whole. The point we wanted to make in this section is that the impact of an optimal pricing policy can be significant and as high as 22%.

5.3 The value of dynamic pricing

In this section we show that the value of dynamic pricing can be significant by comparing the best static to the best dynamic policy. The results for $N = 5$ and $T = 5$ are reported in Table 4. The best static pricing policy in all cases was given by Model (6). In the case of stationary demand, this policy is optimal. The best dynamic pricing policy when the price elasticity did not vary over time (HD, $F = 1$ and NHD, $F = 2$) was DAM, which was marginally better than the DP based method. The latter performed best when the willingness to pay varied over time (NHD, $F = 1$). We will further investigate these differences in the examples of Section 5.4.

Table 4: Revenue gain best dynamic over best static pricing policy

CV^2 (1/k)	0.10	0.14	0.17	0.20	0.25	0.33	0.50	1.00
HD, F=1	2.8%	3.3%	3.5%	3.9%	4.3%	4.8%	5.6%	7.1%
NHD, F=1	36.3%	36.4%	36.5%	36.6%	36.7%	37.1%	37.6%	38.6%
NHD, F=2	3.8%	4.3%	4.6%	5.0%	5.4%	6.0%	6.8%	8.1%

The impact of dynamic pricing increases with the demand variance, since when there is more demand uncertainty the ability to compensate for statistical fluctuations in demand by price changes becomes more useful. Note that even when the price elasticity does not change over time, the value of dynamic pricing can be significant. The improvements found here are of the same order of magnitude as those reported by Federgruen and Heching (1999).

Comparison of dynamic pricing methods

Now that we have established the value of dynamic pricing, we need to investigate which of the methods discussed in Section 3 determines the best such policy. Since the DP based method is the most sophisticated, we expect this method to do best. Therefore, all results below are reported as the gain of DPM over the two other methods. Tables 5-7 compare the performance of the three heuristics for different levels of demand variability and network size.

Table 5: Sensitivity gain DP over DAM and DDM for demand uncertainty ($N = 1, T = 5$)

Problem	CV^2 (1/k)	0.10	0.14	0.17	0.20	0.25	0.33	0.50	1.00
HD, F=1	DAM	0.02%	0.03%	0.05%	0.04%	0.07%	0.11%	0.14%	0.26%
HD, F=1	DDM	0.33%	0.39%	0.38%	0.39%	0.44%	0.49%	0.49%	0.40%
NHD, F=1	DAM	13.07%	12.50%	12.42%	11.73%	11.07%	11.07%	11.55%	10.95%
NHD, F=1	DDM	0.64%	1.06%	1.23%	1.55%	1.79%	2.03%	1.49%	2.36%
NHD, F=2	DAM	0.06%	0.07%	0.08%	0.08%	0.13%	0.19%	0.23%	0.27%
NHD, F=2	DDM	0.20%	0.25%	0.28%	0.42%	0.46%	0.35%	0.35%	0.45%

Table 6: Sensitivity gain DP over DAM and DDM for demand uncertainty ($N = 5, T = 5$)

Problem	CV^2 (1/k)	0.10	0.14	0.17	0.20	0.25	0.33	0.50	1.00
HD, F=1	DAM	(0.02%)	(0.03%)	(0.03%)	(0.04%)	(0.03%)	(0.05%)	(0.08%)	(0.13%)
HD, F=1	DDM	0.09%	0.10%	0.10%	0.11%	0.11%	0.10%	0.05%	(0.04%)
NHD, F=1	DAM	15.1%	15.0%	14.9%	14.8%	14.7%	14.7%	14.8%	15.0%
NHD, F=1	DDM	0.27%	0.36%	0.47%	0.53%	0.52%	0.68%	0.05%	0.28%
NHD, F=2	DAM	(0.03%)	(0.02%)	(0.01%)	(0.01%)	(0.02%)	(0.01%)	0.01%	0.00%
NHD, F=2	DDM	0.11%	0.18%	0.20%	0.22%	0.23%	0.17%	0.23%	0.42%

Table 7: Gain DP over DAM and DDM for different network sizes ($T = 5$)

Problem	N	1	2	3	4	5	6	8	10
HD, F=1	DAM	0.10%	0.04%	(0.03%)	(0.03%)	(0.07%)	(0.07%)	(0.08%)	(0.09%)
HD, F=1	DDM	0.50%	0.29%	0.13%	0.12%	0.06%	0.06%	0.03%	0.01%
NHD, F=1	DAM	11.33%	12.68%	13.87%	14.60%	14.75%	14.85%	15.02%	15.18%
NHD, F=1	DDM	2.01%%	1.52%%	1.12%%	0.75%	0.68%	0.65%	0.54%	0.47%
NHD, F=2	DAM	0.14%	0.08%	0.04%	0.02%	0.00%	(0.01%)	(0.02%)	(0.03%)
NHD, F=2	DDM	0.36%	0.26%	0.22%	0.19%	0.17%	0.16%	0.15%	0.14%

These simulation results confirm our expectations about the performance of the heuristics for different types of test cases. When the price elasticity of demand is constant over time, DAM slightly but consistently outperforms DDM while the performance of DPM is roughly comparable. When the price-elasticity varies over time, DAM performs poorly and DPM significantly outperforms DDM. When the demand uncertainty increases, the performance of DPM relative to the other heuristics seems to improve somewhat, while the opposite is true when the network size increases. We can conclude that DPM method is in general somewhat more robust than the other two heuristics, but only when the price elasticity of demand varies over time this method is clearly preferable.

5.4 The value of a stochastic model

We now compare the DP based method with a simpler approach based on a deterministic model that ignores the randomness of demand. The results are given in Table 9. We also report the revenue gap between DPM and the upper bound J_{UB}, which is in fact the objective value of the deterministic model.

Table 9: Gain DP over deterministic model and difference with UB ($N = 5$,$T = 5$)

Problem	CV^2 (1/k)	0.10	0.14	0.17	0.20	0.25	0.33	0.50	1.00
HD, F=1	Det	0.30%	0.35%	0.38%	0.41%	0.45%	0.51%	0.60%	0.85%
HD, F=1	UB	2.80%	3.42%	3.75%	4.19%	4.79%	5.83%	7.72%	12.76%
NHD, F=1	Det	0.74%	1.01%	1.14%	1.26%	1.53%	1.89%	2.63%	4.46%
NHD, F=1	UB	3.90%	4.80%	5.35%	6.05%	6.93%	8.52%	11.16%	18.14%
NHD, F=2	Det	0.13%	0.19%	0.21%	0.25%	0.32%	0.43%	0.63%	1.34%
NHD, F=2	UB	2.13%	2.80%	3.11%	3.55%	4.28%	5.29%	7.14%	12.78%

The results show that the revenue gain of using a stochastic model can be significant. Note that the advantage of the deterministic approach is that the calculation of the optimal prices may not require solving the first-order optimality condition (4) numerically, but in fact this can be done "off-line" when setting up Model (6). Also note that the UB is not very tight and therefore may not be very useful.

We repeated this experiment for different scales of the problem (α ranging from 1 to 10) but this had no significant effect on the results in Table 9. The revenue of each method and the UB just increased proportionally. Earlier we showed that the revenue of the deterministic model and DPM converge when the demand uncertainty is taken out of the problem. These two results combined show that the finding by Gallego and van Ryzin (1997) and Paschalidis and Tsitsiklis (2000) is not the direct result of the scaling itself, but more of the corresponding reduction in randomness when demand is modeled by a Poisson process.

6 The role of inventory control

Thusfar we have assumed that the resource allocation determines how much of each product can be sold. For airline revenue management, this corresponds to a partitioned seat inventory control policy. In practice, airlines implement nested seat inventory control, where seats allocated to discount classes are available for sale to full-fare passengers as well. The problem with modeling such a policy is that who these seats are eventually sold to depends in part on the arrival order of booking requests. Because of this, it may be hard to express the expected revenue function in closed form. By contrast, the seat partitioning allows a static mathematical programming formulation of the pricing problem. In this section, we investigate the impact of the inventory policy on the optimal prices by working out a small example.

Consider a single-leg flight with two fare classes, of which class 1 has the highest fare. We use the multiplicative demand model (1) with Z exponentially distributed with parameter $\lambda = 1$, and assume that class 2 booking requests come in first. Let b denote the booking limit for class 2. Let

AD_i denote the accepted demand for class i. Then

$$E\left[AD_1|Z_2=z_2\right]=\begin{cases}\mu_1 E_{Z_1}\left[\min\left\{\frac{c-\mu_2 z_2}{\mu_1},Z_1\right\}\right] & \text{if } \mu_2 z_2\leq b\\ \mu_1 E_{Z_1}\left[\min\left\{\frac{c-b}{\mu_1},Z_1\right\}\right] & \text{otherwise}\end{cases}$$

Using that

$$E\left[\min\left\{y,Z\right\}\right]=1-e^{-y}\qquad y\geq 0 \tag{13}$$

we have

$$E\left[AD_1|Z_2=z_2\right]=\begin{cases}\mu_1\left[1-\exp\left(-\frac{c-\mu_2 z_2}{\mu_1}\right)\right] & \text{if } z_2\leq\frac{b}{\mu_2}\\ \mu_1\left[1-\exp\left(-\frac{c-b}{\mu_1}\right)\right] & \text{otherwise}\end{cases}$$

Thus by the law of iterated expectations

$$\begin{aligned}E\left[AD_1\right] &= E_{Z_2}\left[E\left[AD_1|Z_2=z_2\right]\right]\\ &= \mu_1\left(\int_0^{\frac{b}{\mu_2}}\left(1-e^{-\frac{c-\mu_2 z_2}{\mu_1}}\right)e^{-z_2}dz_2+\int_{\frac{b}{\mu_2}}^{\infty}\left(1-e^{-\frac{c-b}{\mu_1}}\right)e^{-z_2}dz_2\right)\\ &= \mu_1\left(1+\frac{\mu_2}{\mu_1-\mu_2}e^{-\left(\frac{c-b}{\mu_1}+\frac{b}{\mu_2}\right)}-\frac{\mu_1}{\mu_1-\mu_2}e^{-\frac{c}{\mu_1}}\right)\end{aligned}$$

for $\mu_1\neq\mu_2$. For $\mu_1=\mu_2$, we get

$$E\left[AD_1\right]=\mu_1\left(1-e^{-\frac{c}{\mu_1}}\right)-be^{-\frac{c}{\mu_1}}$$

Note that for $\mu_1=0$, $E\left[AD_1\right]=0$. By (13), we also have

$$E\left[AD_2\right]=\mu_2\left(1-e^{-\frac{b}{\mu_2}}\right)$$

For $\mu_2=0$, this gives $E\left[AD_2\right]=0$. The expected revenue as a function of p_1, p_2 and b is

$$ER(p_1,p_2,b)=p_1 E\left[AD_1\right]+p_2 E\left[AD_2\right]$$

Note that the first-order optimality condition for the nested booking limit is (Littlewood, 1972)

$$p_1 P(D_1\geq c-b)=p_2$$

which gives

$$b=c-\mu_1\ln\left(\frac{p_1}{p_2}\right)$$

thus b can be eliminated from the revenue optimization problem.

For this small example, we have compared the optimal nested (b, $\overline{p}$) and partitioned (x, p) pricing and seat inventory control policy. Here x denotes the seat allocation to class 2 from the partitioned model, which heuristically can be used as a nested booking limit as well. The results are reported in Table 10. ΔR_1 is how much higher the optimal expected revenue of the nested policy is compared to the optimal expected revenue of the partitioned policy. ΔR_2 is the optimality gap from using the booking limit and prices generated by the partitioned model for nested booking control. ΔR_3 is the optimality gap from using the optimal nested booking limit with the prices from the partitioned model for nested booking control. The first scenario is the base case. In scenarios 2 and 3 the price sensitivity of class 1 demand is varied. In scenarios 4 and 5 the volume of demand is increased. The optimization was done numerically in Excel.

Table 10: Comparison of a nested and a partitioned pricing model

α_1	β_1	α_2	β_2	$\overline{p}_1$	$\overline{p}_2$	b	p_1	p_2	x	ΔR_1	ΔR_2	ΔR_3
60	0.075	120	0.25	453.2	315.5	90.6	485.9	312.9	57.3	9.2%	4.5%	0.3%
60	0.05	120	0.25	679.5	324.8	80.8	707.9	320.5	49.8	7.7%	3.7%	0.2%
60	0.1	120	0.25	338.1	309.7	97.7	373.2	308.2	62.6	10.1%	5.1%	0.6%
120	0.15	240	0.5	513.3	368.9	85.8	536.1	354.6	51.3	10.3%	4.8%	0.4%
240	0.3	480	1	575.4	413.0	77.7	581.5	392.5	42.4	8.7%	4.3%	0.7%

The point to note is that the prices from the partioned model are good input for nested booking control, but the partioned seat allocation is not. The same result was reported by Weatherford (1997) for a similarly small example, assuming Gaussian demand.

The results suggest a two-stage approach to pricing and resource allocation in a nested or make-to-stock setting where substitution of certain products is allowed. First, use a partioned model such as (6) to heuristically determine the price of each product. Then, given these prices, optimize or heuristically determine the optimal inventory policy. For example, an approach for calculating nested booking limits for airline network RM is developed in Bertsimas and De Boer (2002). However, the examples considered here and by Weatherford are very limited. Further research is required to validate this approach.

7 Conclusions

In this paper we have developed a new model for joint pricing and resource allocation in a network, with an important application to airline pricing. We have shown that for certain types of demand distributions this problem is convex, thus tractable for large instances. Based on this model, we have developed and tested several heuristics for dynamic pricing and resource allocation. The strongest approach, in particular when the price elasticity of demand varies over time, is an elegant combination of linear and dynamic programming. Numerical experiments and theoretical bounds on the optimal expected revenue suggest that the extent to which a dynamic policy based on a stochastic model will outperform a simple static policy based on a deterministic model depends on the level of demand variability. When demand is modeled by a Poisson process, the variability decreases with the scale of the problem, in which case the latter approach is asymptotically optimal.

APPENDIX

Proof of Theorem 3: The second inequality follows from each static policy being an admissible dynamic policy as well. To see the first inequality, notice that the optimal static policy is obtained by solving

$$\begin{array}{lll} J^*_{static} = & max_{x^t,p^t} & E\left[\sum_{t,i} p_i^t \min\{x_i^t, D_i^t(p_i^t)\}\right] \\ & s.t. & \mathbf{A}\left(\sum x_t\right) \leq \mathbf{c} \\ & & \mathbf{x}^t, \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{array}$$

Note that for all feasible price vectors $\mathbf{p}^t$ such that $\mathbf{A}\left(\sum \mu^t(p^t)\right) \leq \mathbf{c}$, $\left\{\mathbf{x}^t = \boldsymbol{\mu}^t(\mathbf{p}^t), \mathbf{p}^t\right\}$ is a feasible solution to this problem. Let $J_{static}(\boldsymbol{\mu}^t(\mathbf{p}^t), \mathbf{p}^t)$ denote the expected revenue of such a policy. For any random variable D with mean μ and standard deviation σ, and for any real number d, we have the inequality

$$E\left[(D-d)^+\right] \leq \frac{\sqrt{\sigma^2 + (d-\mu)^2} - (d-\mu)}{2}$$

where $x^+ \triangleq \max(x, 0)$ (e.g. Gallego, 1992). Using that

$$E\left[\min\{x_i^t, D_i^t(p_i^t)\}\right] = E\left[D_i^t(p_i^t)\right] - E\left[\left(D_i^t(p_i^t) - x_i^t\right)^+\right]$$

(sales equal demand minus lost sales), this gives us

$$\min\{x_i^t, D_i^t(p_i^t)\} \geq \mu_i^t(p_i^t) - \frac{\sqrt{\left(\sigma_i^t(p_i^t)\right)^2 + \left(x_i^t - \mu_i^t(p_i^t)\right)^2} - \left(x_i^t - \mu_i^t(p_i^t)\right)}{2}$$

Thus

$$\begin{aligned} J_{static}(\boldsymbol{\mu}^t(\mathbf{p}^t), \mathbf{p}^t) &= E\left[\sum_{t,i} p_i^t \min\{\mu_i^t(p_i^t), D_i^t(p_i^t)\}\right] \\ &\geq \sum_{t,i} p_i^t \left(\mu_i^t(p_i^t) - \frac{\sigma_i^t(p_i^t)}{2}\right) \\ &= \sum_{t,i} p_i^t \mu_i^t(p_i^t) \left(1 - \frac{\rho_i^t(p_i^t)}{2}\right) \end{aligned}$$

For all feasible p^t, we have $J_{static}(\boldsymbol{\mu}^t(\mathbf{p}^t), \mathbf{p}^t) \leq J^*_{static}$. J_{LB} is the maximum such lowerbound. The inequality of the theorem follows. ■

Proof of Theorem 4: First, define

$$\begin{array}{lll} G(\mathbf{z}) = & max_{d^t,p^t} & E\left[\sum_{t,i} p_i^t \min\{d_i^t, D_i^t(p_i^t, z_i^t)\}\right] \\ & s.t. & \mathbf{A}\left(\sum d^t\right) \leq \mathbf{c} \\ & & \mathbf{d}^t, \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{array}$$

Then $F(\mathbf{z}) = G(\mathbf{z})$. To see this, consider the optimal solution $(\mathbf{d}^*, \mathbf{p}^*)$ of the problem defining $G(\mathbf{z})$. If for some i, $D_i^t(p_i^{t*}, z_i^t) < d_i^{t*}$, we could set $d_i^{t*} = D_i^t(p_i^{t*}, z_i^t)$ without affecting the objective function or the feasibility of the solution, since A is assumed to be non-negative. If, on the other hand, $D_i^t(p_i^{t*}, z_i^t) > d_i^{t*}$, we could increase p_i^{t*} by some small $\varepsilon > 0$ such that $D_i^t(p_i^{t*} + \varepsilon, z_i^t) \geq d_i^{t*}$, which would increase the objective function by the non-negativity of d_i^{t*} while keeping the solution feasible. This would contradict the optimality of $(\mathbf{d}^*, \mathbf{p}^*)$ though, so we know that this case can not occur. As a result, we know that we can always choose $d_i^{t*} = D_i^t(p_i^{t*}, z_i^t)$ and effectively eliminate the variables d_i^{t*}, which gives the problem defining $F(\mathbf{z})$. Now let $p(\mathbf{z})$ be the optimal dynamic pricing policy when $\mathbf{Z} = \mathbf{z}$ and let $\mathbf{d}(\mathbf{z})$ be the corresponding amount sold of each product. Then we know that $\mathbf{A} \sum_t \mathbf{d}^t(\mathbf{z}) \leq \mathbf{c}$ (since the resource constraints are not allowed to be violated) and that $\mathbf{d}(\mathbf{z}) \leq \mathbf{D}(\mathbf{p}(\mathbf{z}), \mathbf{z})$ (since we cannot sell more than we had demand for). Note that

$$\begin{aligned} J^*_{dynamic}(\mathbf{z}) &= \mathbf{p}(\mathbf{z})'\mathbf{d}(\mathbf{z}) \\ &= \sum_{t,i} p(z)_i^t \min\{d(z)_i^t, D_i^t(p(z)_i^t, z_i^t)\} \\ &\leq G(\mathbf{z}) = F(\mathbf{z}) \end{aligned}$$

where the last inequality follows from $(\mathbf{p}(\mathbf{z}), \mathbf{d}(\mathbf{z}))$ being a feasible solution to the problem defining $G(\mathbf{z})$ with objective value $J^*_{dynamic}(\mathbf{z})$. This is true for all $\mathbf{z}$, so it also holds in expectation. The result of the theorem follows. ■

Proof of Theorem 6: For the multiplicative demand model, we have

$$\begin{aligned} F(\mathbf{z}) = \quad & max_{\{p_i^t\}} && \textstyle\sum_{t,i} p_i^t \mu(p_i^t) z_i^t \\ & s.t. && \textstyle\sum_{i,t} a_{ij} \mu(p_i^t) z_i^t \leq c_j \qquad j = 1, ..., m \\ & && \mathbf{p}^t, \boldsymbol{\mu}^t(\mathbf{p}^t) \geq \mathbf{0} \end{aligned}$$

We are done when we show that

$$F(\lambda \mathbf{z}^1 + (1-\lambda)\mathbf{z}^2) \geq \lambda F(\mathbf{z}^1) + (1-\lambda) F(\mathbf{z}^2)$$

for two vectors $\mathbf{z}^1, \mathbf{z}^2 \geq 0$ and some scalar $\lambda \in (0;1)$. Let $\mathbf{p}^1$ and $\mathbf{p}^2$ solve the problem for $\mathbf{z}^1$ and $\mathbf{z}^2$ respectively. Then

$$\lambda F(\mathbf{z}^1) + (1-\lambda) F(\mathbf{z}^2) = \lambda \sum_{t,i} p_{i,t}^1 \mu(p_{i,t}^1) z_{i,t}^1 + (1-\lambda) \sum_{t,i} p_{i,t}^2 \mu(p_{i,t}^2) z_{i,t}^2$$

Now let $\overline{p}_{i,t} = \frac{\lambda z_{i,t}^1}{\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2} p_{i,t}^1 + \frac{(1-\lambda) z_{i,t}^2}{\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2} p_{i,t}^2$. We have

$$\begin{aligned} \lambda z_{i,t}^1 p_{i,t}^1 \mu(p_{i,t}^1) + (1-\lambda) z_{i,t}^2 p_{i,t}^2 \mu(p_{i,t}^2) &= \left(\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2\right) * \\ & \left(\frac{\lambda z_{i,t}^1}{\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2} p_{i,t}^1 \mu(p_{i,t}^1) + \frac{(1-\lambda) z_{i,t}^2}{\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2} p_{i,t}^2 \mu(p_{i,t}^2) \right) \\ &\leq \left(\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2\right) \overline{p}_{i,t} \mu(\overline{p}_{i,t}) \end{aligned}$$

by the assumed concavity of $p\mu(p)$. Thus

$$\sum_{t,i} \left(\lambda z_{i,t}^1 + (1-\lambda) z_{i,t}^2\right) \overline{p}_{i,t} \mu(\overline{p}_{i,t}) \geq \lambda F(z^1) + (1-\lambda) F(z^2)$$

Also,

$$\begin{aligned}\mu(\overline{p}_{i,t}) &= \mu\left(\frac{\lambda z^1_{i,t}}{\lambda z^1_{i,t}+(1-\lambda)z^2_{i,t}}p^1_{i,t}+\frac{(1-\lambda)z^2_{i,t}}{\lambda z^1_{i,t}+(1-\lambda)z^2_{i,t}}p^2_{i,t}\right)\\ &\le \frac{\lambda z^1_{i,t}}{\lambda z^1_{i,t}+(1-\lambda)z^2_{i,t}}\mu(p^1_{i,t})+\frac{(1-\lambda)z^2_{i,t}}{\lambda z^1_{i,t}+(1-\lambda)z^2_{i,t}}\mu(p^2_{i,t})\end{aligned}$$

using that $\mu(p)$ is assumed to be convex. Thus

$$\begin{aligned}\sum_{i,t} a_{ij}\mu(\overline{p}_{i,t})\left(\lambda z^1_{i,t}+(1-\lambda)z^2_{i,t}\right) &\le \sum_{i,t} a_{ij}\left(\lambda z^1_{i,t}\mu(p^1_{i,t})+(1-\lambda)z^2_{i,t}\mu(p^2_{i,t})\right)\\ &= \lambda\sum_{i,t} a_{ij}z^1_{i,t}\mu(p^1_{i,t})+(1-\lambda)\sum_{i,t} a_{ij}z^2_{i,t}\mu(p^2_{i,t})\\ &\le \lambda c_j+(1-\lambda)c_j = c_j\end{aligned}$$

using that $z^1, z^2, A \ge 0$. Thus $\overline{p}_{i,t}$ is a feasible solution for the problem defining $F(\lambda\mathbf{z}^1+(1-\lambda)\mathbf{z}^2)$ with objective value at least as large as $\lambda F(\mathbf{z}^1)+(1-\lambda)F(\mathbf{z}^2)$. Thus $F(\lambda\mathbf{z}^1+(1-\lambda)\mathbf{z}^2) \ge \lambda F(\mathbf{z}^1)+(1-\lambda)F(\mathbf{z}^2)$, which is what we needed to show. ■

BIBLIOGRAPHY

R. E. Barlow and F. Proschan, *Mathematical Theory of Reliability*, SIAM Classics in Applied Mathematics, Philadelphia, PA, 1996.

P. P. Belobaba, Air Travel Demand and Airline Seat Inventory Management, Ph.D. thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 1987.

P. P. Belobaba, "Application of a Probabilistic Decision Model to Airline Seat Inventory Control," *Opns. Res.* **37,** 183-197 (1989).

F. Bernstein and A. Federgruen, "A General Equilibrium Model for Decentralized Supply Chains with Price- and Service- Competition," working paper, Graduate School of Business, Columbia University, New York, NY, 2002.

D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, MA, 1999.

D. Bertsimas and S.V. de Boer, "A Stochastic Booking-Limit Policy for Airline Network Revenue Management," submitted to *Operations Research*, 2002.

D. Bertsimas, J. Hawkins and G. Perakis, Optimal Bidding in Online Auctions, submitted to *Operations Research*, 2001.

D. Bertsimas and I. Popescu, On the Relation Between Option and Stock Prices: A Convex Optimization Approach, Working Paper, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, 1999.

D. Bertsimas and I. Popescu, Revenue Management in a Dynamic Network Environment, Working Paper, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, 2000.

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

S. Bratu, Network Value Concept in Airline Revenue Management, Master's thesis, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, 1999.

S. L. Brumelle and J. I. McGill, "Airline Seat Allocation with Multiple Nested Fare Classes," *Opns. Res.* **41,** 127-137 (1993).

R. E. Chatwin, "Multi-Period Airline Overbooking with Multiple Fare Classes," *Naval Res. Logistics* **43,** 603-612 (1996).

R. E. Chatwin, "Multi-Period Airline Overbooking with a Single Fare Class," *Opns. Res.* **46,** 805-819 (1998).

R. E. Chatwin, "Optimal Dynamic Pricing of Perishable Products with Stochastic Demand and a Finite Set of Prices," *Eur. J. Oper. Res.* **125,** 149-174 (2000).

V. C. P. Chen, D. Günther, and E. L. Johnson, A Markov-Decision Problem Based Approach to the Airline YM Problem, Working Paper, The Logistics Institute, Georgia Institute of Technology, Atlanta, GA, 1998.

A. Ciancimino, G. Inzerillo, S. Lucidi, L. and Palagi, "A Mathematical Programming Approach for the Solution of the Railway Yield Management Problem," *Transp. Sci.* **33,** 168-181 (1999).

R. E. Curry, "Optimal Airline Seat Allocation with Fare Classes Nested by Origins and Destinations," *Transp. Sci.* **24**, 193-204 (1990).

J. D. Dana and N.C. Petruzzi, "Note: The Newsvendor Model with Endogeneous Demand," *Management Sci.* **47,** 1488-1497 (2001).

S. V. de Boer, R. Freling, and N. Piersma, "Mathematical Programming for Network Revenue Management Revisited," *Eur. J. Oper. Res.* **137,** 72-92 (2002).

A. Federgruen and A. Heching, "Combined Pricing and Inventory Control," *Opns. Res.* **47,** 454-475 (1999).

Y. Feng and G. Gallego, "Perishable Asset Revenue Management with Markovian Time Dependent Demand Intensities," *Management Sci.* **46,** 941-956 (2000).

G. Gallego, "A Minmax Distribution Free Procedure for the (Q,R) Inventory Model," *O. R. Letts.* **11,** 55-60 (1992).

G. Gallego and G. van Ryzin, "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Sci.* **40,** 999-1020 (1994).

G. Gallego and G. van Ryzin, "A Multi-Product Dynamic Pricing Problem and its Applications to Network Yield Management," *Opns. Res.* **45,** 24-41 (1997).

F. Glover, R. Glover, J. Lorenzo, and C. McMillan, "The Passenger Mix Problem in the Scheduled Airlines," *Interfaces* **12,** 73-79 (1982).

S. Kachani and G. Perakis "A Fluid Model of Dynamic Pricing and Inventory Management for Make-To-Stock Manufacturing Systems", submitted to *Management Science*, 2002.

S. Karlin and C. R. Carr, "Prices and Optimal Inventory Policy," in *Studies in Applied Probability and Management Science*, K. J. Arrow, S. Karlin and H. Scarf (Eds.), Stanford University Press, Stanford, CA, 1962.

A. J. Kleywegt, An Optimal Control Problem of Dynamic Pricing, Working Paper, School of Industrial Engineering and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 2001.

A. Kuyumcu and A. Garcia-Diaz, "A Polyhedral Graph Theory Approach to Revenue Management in the Airline Industry," Computers & Industrial Engineering **38,** 375-396 (2000).

A. H. L. Lau and , "The Newsboy Problem with Price Dependend Demand Distribution," *IIE Trans.* **20,** 168-175 (1988).

H. S. Lau and A. H. L. Lau, "The Multi-Product Multi-Constraint Newsboy Problem: Applications, Formulations and Solution," *J. Opns. Mgmt.* **13,** 153-162 (1995).

H. S. Lau and A. H. L. Lau, "The Newsstand Problem: A Capacitated Multiple-Product Single-Period Inventory Problem," *Eur. J. Oper. Res.* **94,** 29-42 (1996).

C. J. Lautenbacher and S. J. Stidham, "The Underlying Markov Decision Process in the Single-Leg Airline Yield Management Problem," *Transp. Sci.* **33,** 136-146 (1999).

T. C. Lee and M. Hersh, "A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings," *Transp. Sci.* **27,** 252-265 (1993).

K. Littlewood, "Forecasting and Control of Passenger Bookings," in *AGIFORS Symposium Proc.* **12,** Nathanya, Israel, 1972.

J. I. McGill and G. J. van Ryzin, "Revenue Management: Research Overview and Prospects," *Transp. Sci.* **33,** 233-256 (1999).

E. S. Mills, "Uncertainty and Price Theory," *Quart. J. Econom.* **73,** 116-130 (1959).

E. S. Mills, *Price, Output and Inventory Policy*, John Wiley, New York, 1962.

A. J. Nevins, "Some Effects of Uncertainty: Simulation of a Model of Price," *Quart. J. Econom.* **80,** 73-87 (1966).

I. Ch. Paschalidis and J. N. Tsitsiklis, "Congestion-Dependent Pricing of Network Services", *IEEE/ACM Transactions on Networking* **8,** 171-184 (2000).

L. H. Polatoglu, "Optimal Order Quantity and Pricing Decisions in Single Period Inventory Systems," *Internat. J. Production Econom.* **23,** 175-185 (1991).

N. C. Petruzzi and M. Dada, "Pricing and the Newsvendor Problem: A Review with Extensions," *Opns. Res.* **47,** 183-194 (1999).

L. W. Robinson, "Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes," *Opns. Res.* **43,** 252-263 (1995).

R. W. Simpson, Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control, Memorandum M89-1, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 1989.

B. C. Smith and C. W. Penn, "Analysis of Alternative Origin-Destination Control Strategies," in *AGIFORS Symposium Proc.* **28,** New Seaburry, MA, 1988.

J. Subramanian, C. J. Lautenbacher, and S. J. Stidham, "Yield Management with Overbooking, Cancellations and No Shows," *Transp. Sci.* **33,** 147-167 (1999).

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

K. T. Talluri and G. J. van Ryzin, "A Randomized Linear Programming Method for Computing Network Bid Prices," *Transp. Sci.* **33,** 207-216 (1999).

L. J. Thomas, "Price and Production Decisions with Random Demand," *Opns. Res.* **22,** 513-518 (1974).

G. T. Thowsen, "A Dynamic, Nonstationary Inventory Problem for a Price/Quantity Setting Firm," *Naval Res. Logistics Quart.* **22,** 461-476 (1975).

G. J. van Ryzin, "What is Nesting on a Network," presented to AGIFORS Reservations and Yield Management Study Group, Melbourne, Australia, 1998.

L. R. Weatherford, "A Tutorial on Optimization in the Context of Perishable-Asset Revenue Management for the Airline Industry," in *Operations Research in the Airline Industry*, G. Yu (ed.), Kluwer Academic Publishers, 68-100, 1998.

L. R. Weatherford, “Using Prices More Realistically as Decision variables in Perishable-Asset Revenue Management Problems,” *J. Comb. Opt.* **1,** 277-304 (1997).

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

R. D. Wollmer, A Hub-Spoke Seat Management Model, unpublished company report, Douglas Aircraft Company, McDonnel Douglas Corporation, Long Beach, CA, 1986.

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).

P. S. You, "Dynamic Pricing in Airline Seat Management for Flights with Multiple Legs,” *Transp. Sci.* **33,** 192-206 (1999).

L. Young, "Price, Inventory and the Structure of Uncertain Demand,” *New Zealand Oper. Res.* **6,** 157-177 (1978).

E. Zabel, "Monopoly and Uncertainty,” *Rev. Econom. Studies* **37,** 205-219 (1970).

E. Zabel, "Multi-Period Monopoly under Uncertainty,” *J. Econom. Theory* **5,** 524-536 (1972).

W. Zhao and Y.S. Zheng, “Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand,” *Management Sci.* **46,** 375-388 (2000).