Research Article

Efficient frontiers in revenue management

Received (in revised form): 15th October 2011

Robert Phillips

Columbia University School of Business, Nomis Solutions, New York, New York, USA

Robert Phillips is Professor of Professional Practice at the Columbia University Business School and Founder and Chief Science Officer of Nomis Solutions, a company specializing in pricing and revenue optimization solutions for financial services companies.

Correspondence: Robert Phillips, Columbia University School of Business, Nomis Solutions, 408 Uris Hall, 3022 Broadway, New York, New York 10027-6902, USA
E-mail: rp2051@columbia.edu

ABSTRACT We consider the problem of generating the efficient frontier (or Pareto set) between two business goals in a pricing and revenue management context. We show that, under standard conditions on the demand function, the efficient frontier between revenue and profit will be continuous, bounded, downward-sloping and concave when pricing a single product. For the single-leg revenue management problem, we show that the efficient frontier between any two goals that are linear in load, such as revenue, load factor and operating contribution, can be efficiently generated using a weighted-sum (or scalarization) approach. We give some numerical examples of the weighted sum approach applied to the discrete-time single leg revenue management problem, as well as applied to an Expected Marginal Seat Revenue heuristic. We discuss possible extensions to a general choice model and to a full network.

Journal of Revenue and Pricing Management advance online publication, 20 September 2011; doi:10.1057/rpm.2011.26

Keywords: efficient frontier; revenue management; pricing; airline industry

BACKGROUND AND INTRODUCTION

Despite the common characterization of a firm as a short-run profit maximizer, managers often care about goals other than pure short-run profitability. There are at least two reasons for this. One is that the pursuit of long-term profitability might require sacrificing some profitability in the short run. An example of such behavior might be giving preferential availability or pricing to customers considered to have high life-time value, even if it means foregoing some short-term profit. Another example would be a company that prices low upon entry into a new market in order to build a base of customers with the belief that higher market share can be translated into higher margins sometime in the future. A second reason that a company might deviate from short-term profit maximization is that stock analysts follow a number of metrics other than profitability in judging a stock and may 'punish' the stock if some of these metrics suffer. Many of these actions have been studied under the general label of 'earnings management' in which companies manage costs or sales in order to meet analyst or 'market' expectations of various sorts – even at the cost of short-run profitability. As one example, an electronics wholesaling firm sought to maximize profit from a particular line of business subject to the constraint that the average margin for that line of business did not fall beyond a certain level. The rationale given by the CEO for that constraint was that securities analysts scrutinized the average margin for that business carefully and, 'if they felt that we were buying market share at the expense of margin, they would downgrade our shares'.

When a company faces two competing goals in a market it may use an *efficient frontier* as a mechanism for visualizing the tradeoffs between the two goals. An efficient frontier represents the set of all *undominated* (or efficient) combinations of two goals (or metrics) that a firm could achieve. For example a firm might be interested in the efficient frontier between revenue and profit in a particular market. In this case, a point on the efficient frontier would represent a combination of revenue and profit that (a) could be achieved by the firm under some set of decisions and (b) is efficient in the sense that there is no other achievable outcome with both higher revenue and higher profit. Visualizing the efficient frontier can help the company make tradeoffs – if the efficient frontier between revenue and profit is flat, then the company can increase revenue substantially with little effect on profit. If, on the other hand, the efficient frontier is steeply declining, then additional revenue can only be 'purchased' through substantial reductions in profit.

This article discusses the use and calculation of efficient frontiers in pricing and revenue management. We start by introducing the concept of the efficient frontier and presenting conditions under which it will be well behaved. We then analyze the case of an efficient frontier between revenue and profit in the case of a single price being offered to a market. We then consider the classic revenue management problem in which the underlying decisions are how much availability to offer to different booking classes over time. We first consider the problem of calculating efficient frontiers in the dynamic single-resource multi-class revenue management problem. We show that a *weighted-sum* approach will generate the stochastically undominated outcomes on the efficient frontier. We introduce a category of *load-linear* business goals that includes revenue, profit, load and 'customer value' and present a dynamic programming approach that generates points on the efficient frontier for any two load-linear goals. We then present an approach to estimating the efficient frontier using the Expected Marginal Seat Revenue (EMSR) heuristic. Finally, we discuss extensions to general consumer choice models and to large multi-resource networks.

PREVIOUS WORK

The concept of the efficient frontier was introduced by Markowitz (1959) to illustrate the tradeoff between risk and return in a portfolio of investment options. The efficient frontier is still a standard tool in portfolio analysis. Following its introduction in financial engineering, the concept of the efficient frontier has been used in a variety of business settings to enable visualization of the tradeoffs between competing business goals. Thomas *et al* (2005) use an efficient frontier in the context of a commercial lender setting prices to trade-off market share and profit to illustrate that, if a firm is pricing its loans 'correctly', increased market share can only be purchased at the expense of reduced profit. Other efficient frontiers have been used to visualize tradeoffs between fill rates and inventory levels in supply chains (Lee and Billington, 1995), revenue and revenue risk in bandwidth provisioning and route selection in telecommunications networks (Mitra and Wang, 2005), revenue and recreational quality in forest management (Tóth *et al*, 2006) and in many other contexts. Multi-criterion optimization and efficient frontiers are widely used in economics (Luenberger, 1994), particularly in the theory of the firm (Kumbhakar and Knox Lovell, 2000). Boyd and Vandenberghe (2004) examine continuous multi-dimensional efficient frontiers (which they call *optimal trade-off curves*) in terms of convex optimization.

The generation of an efficient frontier is closely related to multi-objective programming and the generation of Pareto frontiers in design optimization.[1] In particular, goal programming (Charnes and Cooper, 1961; Lee, 1972]) enables the solution of models with lexicographic objective functions. The approach that we use in this article is the *weighted sum* approach to multicriteria optimization first proposed by Zadeh (1963). We note that this approach goes by several different names in the literature. For example, Geoffrion (1968) calls it ‘vector maximization’ and describes conditions under which it generates points on the efficient frontier. Boyd and Vandenberghe (2004) call the same approach ‘scalarization’ and discuss its applicability to generating efficient frontiers. The weighted sum approach has been widely used in multicriteria linear programming and is a popular approach to multicriteria optimization in design engineering (Das and Dennis, 1997; Koski, 1988).

To our knowledge, the first paper to address the specific application of efficient frontiers to revenue management is Levin *et al* (2008) who calculate efficient frontiers between expected revenue and the probability that revenue will be above a certain minimum value. Walczak (2010) proposed an approach to estimating points on the efficient frontier between load and revenue for a flight.[2]

THE EFFICIENT FRONTIER

Consider a set of business decisions represented by a vector $x \in \mathbb{R}^n$ where x could be either continuous or discrete. We call a value of x a *policy* and assume there is a region Ω, such that $x \in \Omega$ means that x is feasible. A policy x could represent the prices for one or more products, production levels, allocation of resources among competing activities, booking limits on a flights, and so on. Now consider two different functions of x, $V^1(x)$ and $V^2(x)$, representing two different goals. We establish the following definitions:

1. The pair $(V^1(x),\ V^2(x)) \in \mathbb{R}^2$ for $x \in \Omega$ is termed an *outcome.*
2. $(V^1(x),\ V^2(x)) \in \mathbb{R}^2$ is an *undominated* (or *efficient*) *outcome* if there exists no $x' \in \Omega$ such that $V^1(x') > V^1(x)$ and $V^2(x') \geqslant V^2(x)$ or $V^1(x') \geqslant V^1(x)$ and $V^2(x') > V^2(x)$.
3. $x \in \Omega$ is an *undominated* (or *efficient*) *policy* if its outcome is undominated.
4. The efficient frontier is the set of all undominated outcomes.

We note that all definitions are with respect to the functions $V^1(x)$ and $V^2(x)$ and the feasible policy space Ω. For some pair of objective functions and a particular policy space, we can denote the efficient frontier as a set $S(V^1, V^2, \Omega)$. In situations in which the context is clear, we will suppress the dependence on V^1, V^2 and Ω and simply write S.

It is well known that points on the efficient frontier can be generated by solving the optimization problem:

$$x^* = \arg\max V^1(x) \quad \text{s.t. } V^2(x) \geqslant v^2, x \in \Omega \tag{1}$$

In particular, assume that $(v^1, v^2) \in S(V^1, V^2, \Omega)$. Then, it is clear that $v^1 = V^1(x^*)$ where x^* is the solution of Problem 1. Now, assume that x^* is the solution to (1) for some value of v^2 and let $\hat{v} = V^2(x^*)$. Then x^* is also the solution to problem (1) with the modified constraint $V^2(x) \geqslant \hat{v}$ and $(V^1(x^*), V^2(x^*))$ is clearly on the efficient frontier.

As the numbering of the functions is arbitrary, the same result holds if the functions V^1, V^2 are switched in Problem (1). Note that the point $(V^1(x^*), v^2)$ may not be on the efficient frontier as it is not necessarily the outcome of any feasible policy. Finally, we note that the efficient frontier can also be considered as a function F where $v^2 = F(v^1)$ means that (v^1, v^2) is an undominated outcome. The next sections consider the question of when F is well behaved. We first consider the case of the efficient frontier between revenue and profit in the case of a single continuous price. We then consider the efficient frontier in the context of the revenue management problem of setting booking controls on a flight leg.

THE EFFICIENT FRONTIER WITH A SINGLE PRICE

To gain insight, we first consider the simplest possible case: a single product being sold at a single price with demand a deterministic function of the price. We assume a continuously differentiable demand curve $d(p)$ and a unit cost of $c \geqslant 0$. We assume that the demand curve is strictly decreasing, that is, $d'(p) < 0$. The revenue associated with price p is $R(p) = pd(p)$ and the profit is $\Pi(p) = (p-c)d(p)$. The hazard rate is defined as $h(p) = -d'(p)/d(p)$. Note that the hazard rate (also known as the failure rate) is equal to the elasticity divided by the price.

We are interested in the efficient frontier between revenue and profit. The first lemma provides a condition under which the efficient frontier is well behaved.

Lemma 1: Let $d(p) > 0$ be a continuously differentiable, strictly decreasing demand function with strictly increasing hazard rate, $h(p)$. Then, there exist unique maximizers p^R and p^π of revenue and profit respectively with $p^R < p^\pi$.

Proof: Define $s(p,x) = (p-x)d(p)$. Then, $\partial s(p,x)/\partial p = d(p)[1-h(p)(p-x)]$. $\partial s(p,x)/\partial p$ is strictly decreasing in p from $\partial s(p,x)/\partial p|_{p=x} = d(x) > 0$ with $\lim_{p\to\infty} \partial s(p,x)/\partial p < 0$. This limit holds because $h(p)$ is increasing and p is increasing to infinity. Thus, for any x, there exists a unique $p^*(x)$ such that $\partial s(p^*(x), x)/\partial p = 0$ and $p^*(x)$ is a maximizer as the sign of $\partial s(p,x)/\partial p$ changes from positive to negative at $p^*(x)$. We note that $p^*(x) - x = 1/\mathrm{h}[\mathrm{p}^*(x)]$. Because the right side is strictly decreasing in $p^*(x)$, it is easy to see that $p^*(x)$ must be strictly increasing in x. Thus $p^R < p^\pi$. □

We note that the existence of a unique maximizer of revenue for demand functions with increasing hazard rates is well known – see Lariviere (2006) and the references therein.

We now show that revenue is decreasing and profit is increasing for $p \in [p^R, p^\pi]$.

Lemma 2: The revenue function $R(p) = pd(p)$ is strictly decreasing in the interval $p \in [p^R, p^\pi]$ and the profit function $\Pi(p) = (p-c)d(p)$ is strictly increasing in the same interval.

Proof: Let $\in (p^R, p^\pi)$, then:

$$\begin{aligned} R'(p) &= d(p)(1 - ph(p)) \\ &< d(p)\left(1 - p^R h\left(p^R\right)\right) < 0 \end{aligned}$$

where the first inequality follows from the fact that $h(p)$ is increasing and the second inequality follows from the fact that $p^R = 1/h(p^R)$. Similarly,

$$\begin{aligned} \Pi'(p) &= d(p)(1 - (p-c)h(p)) \\ &> d(p)(1 - (p-c)h(p^\pi)) \\ &> d(p)\left(1 - \frac{(p-c)}{(p^\pi - c)}\right) \\ &> 0 \end{aligned}$$

where the first inequality follows from the fact that $h(p)$ is increasing and the second inequality follows from the fact that $h(p^\pi) = 1/(p^\pi - c)$ and $p^\pi \geqslant c$. □

With the results of these two lemmas, we can fully characterize the efficient frontier.

Proposition 1: Under the conditions of Lemma 1, the efficient frontier is the set of ordered pairs $S = \{(R(p),\ \Pi(p)): p \in [p^R,\ p^\pi]\}$.

Proof: We first show that there is no $p \notin [p^R, p^\pi]$ such that $(R(p), \Pi(p))$ is on the efficient frontier. By Lemma 1, p^R is the global maximizer of $R(p)$ and p^π is the global maximizer of (p) therefore $R'(p) < 0$ and $\Pi'(p) < 0$ for $p > p^\pi$, implying that $(R(p),\ \Pi(p))$ is dominated by $(R(p^\pi),\ \Pi(p^\pi))$ for all $p > p^\pi$. By a similar argument, we must have $R'(p) > 0$ and $\Pi'(p) > 0$ for $p < p^R$, implying that $(R(p), \Pi(p))$ is dominated by $(R(p^R),\ \Pi(p^R))$ for all $p < p^R$. Thus, $(R(p),\ \Pi(p))$ is on the efficient frontier only if $p \in [p^R, p^\pi]$. To show that $p \in [p^R, p^\pi]$ implies that $(R(p),\ \Pi(p))$ is on the efficient frontier, we first note that $(R(p^R), \Pi(p^R))$ and

$(R(p^{\pi}), \Pi(p^{\pi}))$ are clearly undominated and hence on the efficient frontier. For $p \in [p^R, p^{\pi}]$, by Lemma 2, $R'(p) < 0$ and $\Pi'(p) > 0$, therefore $(R(p), \Pi(p))$ cannot be dominated by any $(R(q), \Pi(q))$ with $q \in [p^R, p^{\pi}]$. □

The implication of Proposition 1 is that, under the conditions of Lemma 1, the efficient frontier will show profit as a continuous, downward sloping function of revenue in the domain $[R(p^{\pi}), R(p^R)]$. An example is shown in Figure 1 for the case of the linear demand curve $d(p) = (1000 - 10p)^+$ and cost $c = \$20$. In this case, $p^{\pi} = \$60$ and $p^R = \$50$. We note that $R(p^{\pi}) = \$24\,000$, $R(p^R) = \$25\,000$, $\Pi(p^R) = \$15\,000$, and $\Pi(p^{\pi}) = \$16\,000$. Thus, the domain of the efficient frontier is [\$24 000, \$25 000] and the range is [\$15 000, \$16 000].

As noted in the preceding section, the fact that efficient frontier function is decreasing is a consequence of its definition. It will hold for any pair of objective functions and any policy space. However, continuity and boundedness of the efficient frontier require the stronger assumptions.

The efficient frontier in Figure 1 is concave. Concavity is a desirable property of an efficient frontier. If the efficient frontier were not concave, then the seller could improve both profit and revenue by employing a mixed strategy of charging price p_1 to some fraction α of customers chosen randomly and some other price p_2 to the remaining customers. As this behavior is rarely if ever employed by companies with mixed goals, it suggests that 'real-world' efficient frontiers are concave. Let $R_d(d)$ represent revenue as a function of demand. As the price-response curve, $d(p)$ is strictly downward sloping, $R_d(d)$ is well defined for all $d = d(p)$ with p in the relevant range. Proposition 2 gives conditions on the demand function under which the efficient frontier is guaranteed to be concave.

Proposition 2: If the demand function, $d(p)$ is continuous, differentiable and downward-sloping and $R_d(d)$ is concave for all d such that $d = d(p)$ for some $p \in [p^R, p^{\pi}]$, then the efficient frontier will be concave.

Proof: Let $\Pi(R)$ denote the profit corresponding to revenue R. We can write $\Pi(R) = R - cd(R)$, where, with some abuse of notation, $d(R)$ denotes the demand corresponding to R. Note that $\Pi''(R) = -cd''(R)$, thus a necessary and sufficient condition for the efficient frontier to be concave is that $d''(R) \geqslant 0$. Taking the second derivative using the Implicit Function Theorem, we find that $d''(R) = -R_d''(d)/[R_d'(d)]^3$. However, as $\Pi'(R) = 1 - cd'(R) < 0$ we must have that $d'(R) > 0$ and, hence $R_d'(d) > 0$. This implies that a necessary and sufficient condition for concavity is that $R_d''(d) < 0$. □

Some discussion of the concavity condition on $R_d(d)$ and its relationship with other common demand function assumptions in revenue

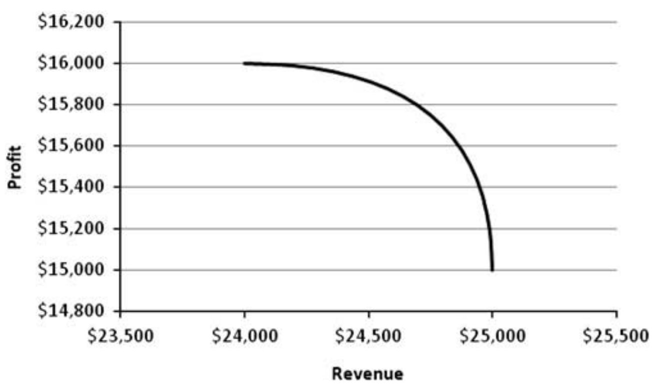

Figure 1: Efficient frontier with revenue and profit for linear demand example described in the text.

management increasing hazard rate and concavity of the revenue function with respect to price can be found in Ziya *et al* (2004).

In the case of a single price, estimating the efficient frontier is straightforward: determine p^R and p^π after which points on the frontier can be calculated as pairs $(R(p), \Pi(p))$ for any price p between p^R and p^π. We note that, in the case of a seller offering a single product at a single price, the frontier is not only efficient, it uniquely describes the attainable revenue and profit pairs – that is, points inside the efficient frontier are not attainable. As a seller offers multiple products at multiple prices through multiple channels, possibly with complex business constraints that link them, interior points will become feasible and the efficient frontier becomes the boundary of a set and, hence, much more difficult to compute. A number of pricing optimization providers such as Nomis Solutions and others provide the capability for a user to view the 'efficient frontier' for a particular set of market conditions (Britting, 2006).

With this general background, we can now consider the problem of computing an efficient frontier for a constrained resource being sold with many different prices – in other words, efficient frontiers in revenue management.

CALCULATING THE EFFICIENT FRONTIER FOR A SINGLE RESOURCE

In the single-resource revenue management problem, a perishable product with constrained capacity can be sold at a number of different prices to different buyers differentiated by time of booking, geography, and so on. The amount of product available to buyers in different segments at any moment can be controlled by the seller. In the 'classical' revenue management problem, there are a set of booking classes, each with an associated fixed price, and the seller determines how much product to offer at each price at each time to each customer segment. Given that the seller does not know how many customers for each booking class will arrive in the future, he needs to make his decisions under uncertainty. Discussions of the revenue management problem and approaches used to solve it can be found in Talluri and van Ryzin (2005) and Phillips (2005).

We consider the dynamic single-resource revenue management problem with n booking classes, indexed from 1 to n. Booking requests occur over time periods $t = 1, 2, \ldots, T$ with T being the time of departure. A booking request for fare class i occurs at in time period t with probability p_{it} and no booking request occurs with probability p_{0t}. We assume that the probability of more than one booking request in a period is negligibly small, so that, for all t, $p_{0t} + p_{1t} + \cdots + p_{nt} = 1$. Let f_i denote the fare for class i. We assume that fares are ordered so that $f_1 > f_2 > \cdots > f_n$ and that there are no no-shows or cancellations. The decision facing the revenue manager is which booking requests to accept and which to reject in each time period. This is a standard model for dynamic single-resource revenue management formulated originally by Lee and Hersh (1993).

Common business goals that a manager might pursue for a particular flight include maximizing expected revenue, maximizing expected profit and maximizing expected load factor.[3] Revenue, profit and load factor share the property that they are weighted functions of the loads in each booking class, that is, for a particular flight, they can be calculated as $w^T x$, where $x \in \mathbb{R}^n$ is the vector of the loads in each booking class and $w \in \mathbb{R}^n$ is a vector of weights. We call any business goal that can be calculated as a weighted sum of expected loads a *load-linear goal*. We note that for any load-linear goal, $E[w^T x] = w^T E[x]$. The weightings corresponding to the three goals of maximizing revenue, maximizing profit and maximizing load are:

$$w_i^R = f_i, \quad w_i^P = f_i - c_i, \quad w_i^L = 1$$

where the superscript R denotes 'revenue', P denotes 'profit' and L denotes 'load'.

We note that the class of load-linear goals does not include all performance measures that might be of managerial interest. For example, it does not include the probability that revenue will exceed a specified threshold – one of the goals considered by Levin *et al* (2008). However, the class of load-linear goals does include other measures that can be of managerial interest. For example, an airline might wish to maximize the total 'lifetime value' of bookings on a flight where the lifetime value of a customer may differ from both his fare and his incremental profit. An efficient frontier could be used to help understand the tradeoff between maximizing expected total lifetime value and maximizing expected revenue or profit from a given flight. Minimizing total expected incremental passenger cost for a flight is also a load-linear goal.

Let $w^A = (w_1^A, w_2^A, \ldots, w_n^A)^T$ be the weighting vector for some load-linear goal. Then, the Bellman equation to maximize $E[w^T x]$ is:

$$V_t^A(s) = V_{t+1}^A(s) + \max_{u_{it}} \sum_{i=1}^{n} p_{it} u_{it}^A \left[w_i^A - \Delta_{t+1}^A(s) \right] \quad (2)$$

The state variable $s \geqslant 0$ is the number of remaining unbooked seats and $\Delta_{t+1}^A(s) = V_{t+1}^A(s) - V_{t+1}^A(s-1)$. $u_{it}^A \in \{0, 1\}$ are the controls for each period such that a booking request of type i in period t is accepted if $u_{it}^A = 1$ and rejected if $u_{it}^A = 0$. The boundary conditions are:

$$V_{T+1}^A(s) = 0, \quad s = 0, 1, \ldots, C$$

$$V_t^A(0) = 0, \quad t = 1, \ldots, T$$

where $C > 0$ is the initial capacity. Under any of the three pure objective functions, the goal is to find the feasible policy $\wp \in \Omega$ that maximizes $V_1^A(C)$ for $A = R$, P or L, where a policy consists of a set of values $u_{it}^A(s) \in \{0, 1\}$ for $i = 1, 2, \ldots, n$; $T = 1, 2, \ldots, T$; and $s = 0, 1, 2, \ldots, C$ and Ω denotes the set of feasible policies.

Now, assume that we wish to calculate points on the efficient frontier for two load-linear goals – for concreteness, let us assume that we wish to calculate the efficient frontier between revenue and load. Note that the efficient frontier is now discrete as the underlying policy space is discrete. One possibility is to sequentially solve a set of constrained dynamic programming problems of the form:

$$\wp^* = \arg\max_{\wp \in \Omega} \left[V^R(\wp) \right] \quad (3)$$

$$\text{subject to } V^L(\wp) \geqslant \ell, \quad \wp \in \Omega$$

where ℓ specifies the minimum expected load factor and $V^L(\wp)$ and $V^R(\wp)$ are the expected load factor and revenue respectively that will be achieved over the entire time horizon given policy $\wp$. Then, as noted in the Section 'The Efficient Frontier', the efficient frontier $S(V^R, V^L, \Omega)$ is the set of all outcomes $(V^R(\wp^*), V^L(\wp^*))$ for $\wp^*$ that solve problem (3) for all achievable values of ℓ.

We note that the efficient frontier function $v^L = F(v^R)$ defined by S is discrete and, as with all efficient frontiers, downward sloping. However, it may not be concave. Consider a three-class example with three periods before departure and one seat currently unbooked. We assume that here is a possibility of, at most, one booking in each period and the booking probabilities and fares are as shown in Table 1. The points on the efficient frontier (v^R, v^L) for this example are:

- (\$100, 1) – achieved by opening class 2 in period 1,

Table 1: Fares and booking request probabilities for the example discussed in the text

Booking class	*Fare*	*Probability of booking request*		
		Period 1	*Period 2*	*Period 3*
1	\$500	0	0	0.4
2	\$100	1	0	0
3	\$70	0	0.5	0
No booking request		0	0.5	0.6

- (\$135, 0.7) – achieved by opening classes 1 and 3 but keeping class 2 closed.
- (\$200, 0.4) – achieved by only opening class 1.

We note that the frontier in this case is not concave as, starting at an expected load factor of 1, increasing expected revenue by \$35 requires a reduction of 0.3 in expected load factor, while an additional \$65 can be achieved from a second reduction of 0.3.

This counterexample notwithstanding, both computational experience and intuition suggest that the efficient frontier in most real-world examples should be concave. Otherwise, as noted before, airlines could achieve efficient results by using mixed strategies. It is tempting to conjecture that the efficient frontier for load-linear goals is concave in the limit as time goes to infinity.

The weighted sum approach

Define $w_i^{AB}(\alpha) = \alpha w_i^A + (1-\alpha)w_i^B$ with $0 \leqslant \alpha \leqslant 1$ as a *weighted sum* coefficient. For the goals of revenue, load and profit, we have three possible weighted sum coefficient vectors:

$$w_i^{RL}(\alpha) = \alpha w_i^R + (1-\alpha)w_i^L = \alpha(f_i - 1) + 1$$
$$w_i^{RP}(\alpha) = \alpha w_i^R + (1-\alpha)w_i^P = f_i + (1-\alpha)c_i$$
$$w_i^{PL}(\alpha) = \alpha w_i^P + (1-\alpha)w_i^L = \alpha(f_i - c_i - 1) + 1$$

The *weighted sum objective function* for any pair of goals, A and B is $[w^{AB}(\alpha)]^T x$. We note that, if A and B are load-linear objective functions, then the weighted-sum objective function combining A and B will also be load-linear. For any combination of two policies and any value of α, we set up the dynamic program to maximize the corresponding weighted-sum objective function in the obvious way:

$$V_t^{AB}(s) = V_{t+1}^{AB}(s) + \max_{u_{it}\in\{0,1\}} \sum_{i=1}^{n} p_{it}u_{it} \times \left[w_i^{AB}(\alpha) - \Delta_{t+1}^{AB}(s)\right] \qquad (4)$$

$$V_{T+1}^{AB}(s) = 0, \quad s = 0, 1, \ldots, C$$
$$V_t^{AB}(0) = 0, \quad t = 1, \ldots, T$$

There are two questions that we can ask about the dynamic program specified in (4). First, does it generate points on the efficient frontier? Secondly if it does, will it generate every point on the efficient frontier for some value of α?

To answer these questions, we establish some notation. For any policy $\wp \in \Omega$, $V^k(\wp)$ is the corresponding objective function value for goal k. As before, the outcome of a policy is the two-vector $(V^1(\wp), V^2(\wp))$. We say that an outcome $(v^1, v^2) \in \mathcal{R}^2$ is achievable if there is some policy $\wp \in \Omega$ such that $(V^1(\wp), V^2(\wp)) = (v^1, v^2)$. In addition, we say that an outcome $v \in \mathcal{R}^2$ is *stochastically achievable* if there is a value $\alpha \in [0,1]$ such that $v = \alpha v(\wp_i) + (1-\alpha)v(\wp_j)$ for some $\wp_i$, $\wp_j \in \Omega$. We note that a stochastically achievable outcome v can be achieved in expectation by choosing policy $\wp_i$ with probability α and choosing policy $\wp_j$ with probability $1-\alpha$. Let S be the set of outcomes on the efficient frontier and let Z be the set of all *undominated stochastically achievable outcomes* – that is the set of stochastically achievable outcomes that are not dominated by some other stochastically achievable outcome. We define the set $T = S \cap Z$ as the set of deterministic outcomes that are not dominated by any stochastically achievable outcome. If the efficient frontier function is concave, then $T = S$, if the efficient frontier function is continuous and concave then $T = S = Z$. Finally, we define U as the set of outcomes that are achievable by maximizing a weighted sum objective function for some value of α. That is,

$$U = \{\left(v^1(\wp), v^2(\wp)\right) | \wp = \max_{\wp\in\Omega}\left[\alpha v^1(\wp) + (1-\alpha)v^2(\wp)\right], \alpha \in [0,1]\}$$

The following result shows that the set of outcomes from maximizing weighted sum objective functions is equal to the set of (stochastically) undominated deterministic outcomes.

Proposition 3: If no three points in T are co-linear, then $U = T$.

Proof: Theorem 1 in Geoffrion (1968) shows that $U \subseteq T$ We now show that $T \subseteq U$. For any point $(v_i^1, v_i^2) \in T$, define the function g by $v_i^2 = g(v_i^1)$. Note that the function g is concave by the fact that each point in T is stochastically undominated. By the definition of concavity, there must be a line, say $v^2 = b_i - m_i v^1$ through (v_i^1, v_i^2) such that $b_i - m_i v_j^1 > v_j^2$, with $m_i \geqslant 0$ for all $j \neq$ i. The strict inequality is justified by the assumption that no three points in T are co-linear. We note that:

$$\frac{m_i}{(m_i+1)} v_i^1 + \frac{1}{(m_i+1)} v_i^2 = \frac{b_i}{(m_i+1)}$$
$$> \frac{m_i}{(m_i+1)} v_j^1 + \frac{1}{(m_i+1)} v_j^2$$

for all $j \neq$ i. Thus, $(v_i^1, v_i^2) = U(\alpha)$ for $\alpha = m_i/(m_i+1)$. □

The implications of Proposition 3 are illustrated in Figure 2. Essentially, the weighted-sum approach generates all of the points on the 'convex hull' of the upper-right boundary of the feasible region. Points on the efficient frontier that are not on this hull – such as point C in Figure 2 – are stochastically dominated and will not be generated by the weighted sum approach. This phenomenon is also noted by Boyd and Vandenberghe (2004). Finally, note that the requirement that no three points on the efficient frontier be co-linear is strictly a technical one. If three points on the efficient frontier are co-linear, then for one of them, there is a single value of α for which that point will give the weighted-sum objective value as the other two, but there is no guarantee that the solution algorithm will recommend the policy that generates that point rather than one of the other two.

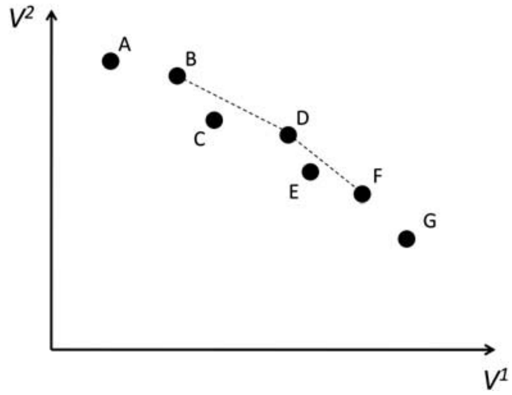

Figure 2: Points on a hypothetical efficient frontier. Points C and E are on the efficient frontier. However, neither one would be generated by the weighted sum approach because point C is stochastically dominated by the combination of points B and D and point E is stochastically dominated by the combination of points D and F.

Corollary If the efficient frontier function is concave, then the weighted-sum approach will generate all points on the efficient frontier.

We note that Proposition 3 and its corollary are completely general. They do not require any of the specific properties of the load-linear goals – in particular, they do not require linearity. However, load-linear goals have one desirable property – the optimal policies for all load-linear goals are nested, that is, in each state at each time, if it is optimal to have a class i open then it is optimal to have all classes with $w_j^A(\alpha) \geqslant w_i^A(\alpha)$ open and if it is optimal to have class I closed it is optimal to have all classes with $w_j^A(\alpha) \leqslant w_i^A(\alpha)$ also closed. As the weighted-sum objective function preserves load-linearity, this property is also true of the weighted-sum dynamic program in (4). The nesting property drastically reduces the space of possible policies and enables much more efficient computation.

Numerical example

Consider the case of a flight with three booking classes with associated fares $f_1 = \$1000$, $f_2 = \$750$ and $f_3 = \$150$. We assume that demand

Table 2: Booking probabilities by class and period for the dynamic example described in the text

Period range	*Booking probability in each period*			
	No booking (p_0)	*Class 1 booking (p_1)*	*Class 2 booking (p_2)*	*Class 3 booking (p_3)*
1–100	0.9	0.0027	0.0243	0.073
101–200	0.9	0.0095	0.0905	0
201–300	0.9	0.1	0	0
	Mean	11.22	11.48	7.29
	Variance	10.21	10.60	6.79

Table 3: Expected revenue and expected load factor as a function of α for three different capacity levels for the dynamic example discussed in the text

α	*c=10*		*c=20*		*c=30*	
	Revenue	*Load factor*	*Revenue*	*Load factor*	*Revenue*	*Load factor*
1	$9401.51	9.64	$16 680.04	19.11	$19 931.70	25.16
0.9	$9400.57	9.66	$16 679.55	19.13	$19 878.38	26.25
0.8	$9397.97	9.68	$16 678.53	19.13	$19 768.85	26.87
0.7	$9362.93	9.78	$16 657.94	19.19	$19 631.82	27.29
0.6	$9343.80	9.81	$16 649.90	19.21	$19 487.71	27.56
0.5	$9329.09	9.83	$16 624.96	19.23	$19 352.50	27.73
0.4	$9254.01	9.89	$16 488.72	19.34	$19 222.83	27.83
0.3	$9192.53	9.92	$16 200.66	19.49	$19 124.15	27.88
0.2	$9116.15	9.94	$15 662.68	19.66	$19 051.65	27.90
0.1	$8690.92	9.97	$14 547.23	19.86	$18 982.09	27.92
0	$3675.11	10	$11 203.68	19.96	$18 853.15	27.93

arrives independently over 300 periods with at most one booking request per period. The probability of a request for a booking in each class in each period is shown in Table 2 along with the mean and variance of total bookings in each class. The expected total number of bookings across all classes is 30 with a variance of 27. Note that while high-fare bookings tend to arrive later, this example does not assume a strict booking order.

Table 3 shows the revenues and load factors achieved for different values of α calculated using the weighted-sum dynamic programming approach for capacities of 10, 20 and 30. As a practical matter, all of the fares were divided by $1000 before calculation, so that the weighted sum value used for a booking request with fare f_i is $\alpha f_i/1000+1-\alpha$. Figure 3 shows the efficient frontiers between load factor and revenue for the three capacity levels.

Note that the efficient frontier in Figure 3 is quite flat when capacity is close to the mean demand ($C=30$). In this case, additional load factor can be purchased relatively cheaply by increasing the bookings from early booking low-fare demand with little effect on expected revenue from later-booking high fares. When capacity is quite scarce relative to demand ($C=10$), revenue is initially quite flat, but drops off very steeply. In this case, some additional

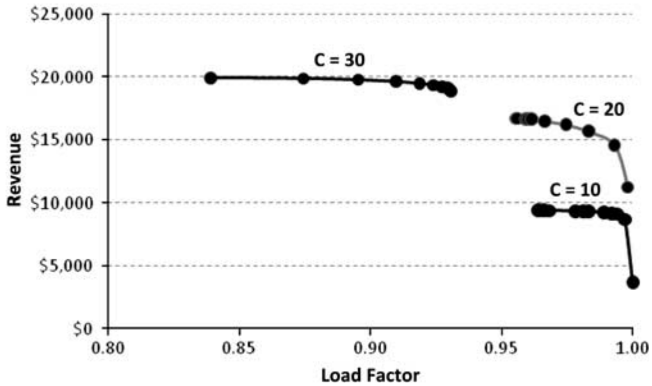

Figure 3: Efficient frontiers between load factor and total revenue computed using the weighted-sum approach. Efficient frontiers are shown for capacity levels (C) of 10, 20 and 30. For details on the demand and fare assumptions, see the text.

load factor can be purchased reasonably cheaply by leaving the relatively high 'middle fare' of $750 open a bit longer. However, reaching the highest expected load requires a drastic loss in expected revenue as it requires leaving all seats open to the lowest fare bookings of $150 in every period. The intermediate capacity case (C = 20) has a relatively high cost of additional load factor across the board.

We note that one endpoint of the efficient frontier between load factor and revenue is the maximum expected revenue achievable on the flight and that the other endpoint is the expected revenue from a policy of 'first come, first served'. The value of revenue management is the difference between these two values, that is, $V^R(1) - V^R(0)$. Note that this measure is very similar to the 'Revenue Opportunity Model' (ROM) for measuring revenue management effectiveness introduced by Smith *et al* (1992) and described in Chapter 7 of Phillips (2005). For the three capacity levels shown in Table 2, that value is $5726.40 for C = 10; $5476.36 for C = 20; and $1078.55 for C = 30. As expected, the expected value of revenue management decreases as capacity increases. However, it is notable that small increases in load factor can be purchased more cheaply in the case in which C = 10 than in the case when C = 20, even though the expected value of revenue management is higher. This suggests that the curvature of the efficient frontier – not just the endpoints – should be important to management in choosing which flights to manage in which fashion.

An EMSR-based heuristic

Although the dynamic-programming-based approach using load-linear goals enables exact computation of points on the efficient frontier, it suffers from the 'curse of dimensionality' and would be difficult to compute on a routine basis. In fact, it is not commonplace for airlines to use the dynamic program with the Bellman Equation in equation (2) for the same reason – it is simply too complex and time-consuming to run on a routine basis for every flight. Instead, many airlines use some version of the EMSR heuristic introduced by Belobaba (1989). Fortunately, we can combine the weighted goal approach with EMSR by defining α-weighted goals as below.

Profit and Revenue: $w_i^{RP}(\alpha) = f_i - \alpha c_i$
Revenue and Load: $w_i^{RL}(\alpha) = \alpha(f_i - 1) + 1$
Profit and Load: $w_i^{PL}(\alpha) = \alpha(f_i - c_i - 1) + 1$

When $\alpha = 1$, the EMSR heuristic will set booking limits based on the first goal, when $\alpha = 0$ it will set booking limits based on the second goal, for intermediate values of α, it will set booking limits based on the weighted average of the two goals. The efficient frontier can be traced by determining the booking limits for different values of α, and using simulation to calculate the expected revenue.

We illustrate the approach with an example based on an aircraft with a capacity of 100 seats and four fare classes with fares as shown in Table 4. Bookings occur in strict fare order – that is all class 4 bookings occur before all class 3 bookings which occur before all class 2 bookings and so on. We assume that the demands for the individual classes are independent normal distributions. We consider the two demand cases shown in Table 4. Case 1 is adapted from Talluri and van Ryzin (2005), Table 2 . Case 2 uses the same fares as Case 1 and has the same level of mean total demand, but has more of the total demand concentrated in the lowest fare class.

Table 4: Fares and means and standard deviations of demand by class for the two example cases described in the text

Class	*Fare*	*Case 1*		*Case 2*	
		μ_i	δ_i	μ_i	δ_i
1	$1050	17.3	5.8	15	5.8
2	950	45.1	15.0	22	15.0
3	699	39.6	13.2	29	13.2
4	520	34.0	11.3	70	35.0
	Total	136	23.7	136	40.7

For both cases, we calculate the efficient frontier between revenue and load factor. To do this, we first normalize by dividing all of the fares by the lowest fare ($520). The resulting normalized fares are (1, 1.34, 1.83, 2.02). Table 5 shows the values of $v_i(\alpha)$, the protection levels for different values of α, and the corresponding expected revenues and expected load factors in Case 1. (The protection level for class 4 is always 0 and is not shown). Note that, as α increases from 0, the values of $v_i(\alpha)$ converge toward 1 and the protection levels converge toward 0. Under the policy of load-factor maximization ($\alpha = 1$), 'First Come, First Served' is the optimal policy and there is no incentive to protect seats for late booking customers. However, whenever revenue is positively weighted, there is more of an incentive to protect seats for late booking customers.

Table 6 shows the results for Case 2 and Figure 4 illustrates the efficient frontier between expected load factor and expected revenue for both cases. For both cases, the efficient frontier is highly concave – that is, initial increases in load factor can be achieved relatively cheaply in terms of forgone revenue but, at some point in each case, marginal increases in load factor become increasingly expensive in terms of revenue foregone. Thus, in Case 2, it costs less than $1000 to increase expected load factor

Table 5: Weighted-sum objective coefficients, protection levels, expected load factor and expected revenue for Case 1 as described in the text

α	*$v_i(\alpha)$ by Class*				*Protection levels by class*			*LF*	*Revenue*
	1	*2*	*3*	*4*	*1*	*2*	*3*		
0	2.02	1.83	1.34	1.00	9.7	53.3	96.8	0.93	$79 649
0.2	1.82	1.66	1.28	1.00	9.3	51.6	94.0	0.95	$79 520
0.4	1.61	1.50	1.21	1.00	8.8	49.5	90.3	0.96	$79 181
0.6	1.41	1.33	1.14	1.00	8.0	46.3	85.4	0.97	$78 428
0.8	1.20	1.17	1.07	1.00	6.3	41.0	77.4	0.99	$74 874
0.95	1.05	1.04	1.02	1.00	3.6	31.0	63.6	0.99	$71 327
1.0	1.00	1.00	1.00	1.00	0	0	0	0.99	$70 133

Table 6: Weighted-sum objective coefficients, protection levels, expected load factor and expected revenue for Case 2 as described in the text

α	$v_i(\alpha)$ *by class*				*Protection levels by class*			*LF*	*Revenue*
	1	*2*	*3*	*4*	*1*	*2*	*3*		
0	2.02	1.83	1.34	1.00	7.4	28.3	60.6	0.92	\$66 477
0.2	1.82	1.66	1.28	1.00	7.0	26.7	57.7	0.93	\$66 367
0.4	1.61	1.50	1.21	1.00	6.5	24.5	54.1	0.94	\$66 096
0.6	1.41	1.33	1.14	1.00	5.7	21.3	49.1	0.95	\$65 507
0.8	1.20	1.17	1.07	1.00	4.3	15.9	41.1	0.96	\$62 666
0.95	1.05	1.04	1.02	1.00	1.3	6.0	27.4	0.96	\$60 371
1.0	1.00	1.00	1.00	1.00	0	0	0	0.96	\$58 595

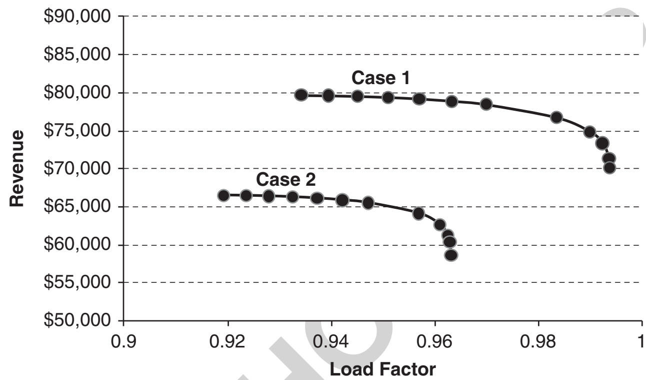

Figure 4: Efficient frontiers between load factor and revenue for Case 1 and Case 2 described in the text.

from 92 to 95 per cent. But increasing from 95 to 96.3 per cent costs almost \$7000 in lost revenue.

CONCLUSIONS AND DISCUSSION

This article discussed the computation of the efficient frontier for pricing and revenue management decisions. It provided straightforward conditions under which the efficient frontier for a single product price will be a downward-sloping, continuous, concave function defined on an interval. The article also discussed the use of the weighted-sum approach for the dynamic single-resource problem and showed that the approach would generate the stochastically undominated points on the efficient frontier. However, solving the weighted-sum dynamic program to determine a point on the efficient frontier is of the same order of difficulty as solving the revenue maximization (or profit maximization) dynamic program. A more tractable approach is to use the hybrid objective function within the standard EMSR setting to determine the corresponding protection levels after which simulation can be used to determine the points on the efficient frontier. To the extent that an airline uses an EMSR heuristic to calculate protection limits, then the EMSR-based approach to calculating the efficient frontier should give a better view of the tradeoffs that it faces than using a dynamic approach.

Numerical experiments show that individual flights can demonstrate very different revenue/load factor tradeoffs depending upon their fare structures, demand distribution among booking classes and capacity levels. An airline that wished to maximize profitability but maintain a certain level of total load or load factor would do well to search for those flights on which additional load can be 'bought' with only slight reductions in profitability.

The approach has been proposed and developed in a single-leg context with independent booking classes. In principle, the weighted sum approach could be applied to the standard Bellman equation for a network so that equation (2) would be replaced with:

$$V_t^{AB}(x) = \max_{u_{it}} \sum_{i=1}^{n} \{ p_{it} u_{it} [w_i^{AB}(\alpha), + V_{t+1}^{AB}(x - A_i)] + (1 - p_{it} u_{it}) V_{t+1}^{AB}(x) \}$$

where x is a vector of remaining leg capacities and A_i is the column of the incidence matrix specifying which legs are used by product i. Proposition 3 implies that solving this dynamic program with different values of α would generate the stochastically undominated outcomes on the efficient frontier. However, although this formulation has been widely studied (Talluri and van Ryzin, 2005), it remains a formidable challenge to solve for realistic networks.

A second challenge is determining the efficient frontier when booking class demands are not independent. The use of a consumer choice model does not change the Bellman equation (2), but it makes the problem much harder to solve because the booking classes are not necessarily nested in fare order. However, many airlines use various modifications of the EMSR heuristic to incorporate consumer 'buy-up' and 'dilution' behaviors as described in Belobaba and Weatherford (2007) and Gallego *et al* (2009). The method of the Section 'An EMSR-based heuristic' would generally be applicable to these modified EMSR heuristics.

ACKNOWLEDGEMENT

I am grateful to Robin Raffard of Nomis Solutions for his very detailed and useful comments on an earlier version of this article. I am also grateful to two anonymous reviewers for their very useful comments and suggestions.

NOTES

1 The terms 'efficient frontier', 'Pareto frontier' and 'Pareto set' are often used interchangeably in the literature.

2 It should be noted that the term 'efficient frontier' has been used in the revenue management literature to denote the undominated combinations of total expected demand and total expected revenue that could be achieved from different combinations of choice sets. This use of the term 'efficient frontier' – which is different from ours – was introduced by Talluri and van Ryzin (2004) and used in this sense in some other papers dealing with choice modeling such as Maglaras (2006).

3 *Load factor* is defined as the number of passengers carried on a flight departure (the *load*) divided by the capacity of the flight. It is a commonly used metric in the passenger airline industry.

REFERENCES

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

Belobaba, P. and Weatherford, L. (2007) Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences* 27(2): 343–363.

Boyd, S. and Vandenberghe, L. (2004) *Convex Optimization*. Cambridge, UK: Cambridge University Press.

Britting, B. (2006) Profit-Based Pricing: Time to Stop Leaving Money on the Table. Tower Group Report V48:13C. Tower Group. Needham, MA.

Charnes, A. and Cooper, W.W. (1961) Management Models and Industrial Applications of Linear Programming, New York: Wiley.

Das, I. and Dennis, J.E. (1997) A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural Optimization* 14(1): 63–69.

Gallego, G., Li, L. and Ratliff, R. (2009) Choice-based EMSR methods for single-leg revenue management with

demand dependencies. *Journal of Revenue and Pricing Management* 8(2–3): 207–240.
Geoffrion, A.M. (1968) Proper efficiency and the theory of vector maximization. *Journal of Mathematical Analysis and Applications* 22(3): 618–630.
Koski, J. (1988) Multicriteria truss optimization. In: W. Stadler (ed.) *Multicriteria Optimization in Engineering and the Sciences*. New York: Plenum Press, pp. 263–307.
Kumbhakar, S.C. and Knox Lovell, C.A. (2000) *Stochastic Frontier Analysis*. Cambridge, UK: Cambridge University Press.
Lariviere, M.A. (2006) A note on probability distributions with increasing generalized failure rates. *Operations Research* 54(3): 602–604.
Lee, H.L. and Billington, C. (1995) The evolution of supply-chain-management models and practice at Hewlett-Packard. *Interfaces* 25(9): 42–63.
Lee, S.M. (1972) Goal Programming for Decision Analysis. Philadelphia, PA: Auerback.
Lee, T.C. and Hersh, M. (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27(3): 252–265.
Levin, Y., McGill, J. and Nediak, M. (2008) Risk in revenue management and dynamic pricing. *Operations Research* 56(2): 326–343.
Luenberger, D. (1994) *Microeconomic Theory.* New York: McGraw-Hill.
Maglaras, C. (2006) Dynamic pricing strategies for multi-product revenue management problems. *Manufacturing and Service Operations Management* 8(2): 136–148.
Markowitz, H. (1959) *Portfolio Selection*. New Haven, CT: Cowles Foundation, Yale University.
Mitra, D. and Wang, Q. (2005) Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management. *IEEE/ACM Transactions on Networking* 13(2): 221–238.
Phillips, R. (2005) *Pricing and Revenue Optimization*. Stanford, CA: Stanford University Press.
Smith, B.C., Leimkuhler, J.F. and Darrow, R.M. (1992) Yield management at American Airlines. *Interfaces* 22(1): 8–31.
Talluri, K. and van Ryzin, G. (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.
Talluri, K. and van Ryzin, G. (2005) *The Theory and Practice of Revenue Management*. New York, NY: Springer.
Thomas, L.C., Oliver, R.W. and Hand, D.J. (2005) A survey of the issues in consumer credit modelling research. *Journal of the Operations Research Society* 56(9): 1006–1015.
Tóth, S.F., McDill, M.E. and Rebain, S. (2006) Finding the efficient frontier of a bi-criteria, spatially explicit, harvest scheduling problem. *Forest Science* 52(1): 93–113.
Walczak, D. (2010) Balancing expected profit contribution and resource utilization. Presented at 2010 INFORMS Revenue Management and Pricing Conference; 18 June, Ithaca, NY.
Zadeh, L.A. (1963) Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control, AC* 8(1): 59–60.
Ziya, S., Aydan, H. and Foley, R.D. (2004) Relationships among three assumptions in revenue management. *Operations Research* 52(5): 804–809.