
Research Article

Optimization of mixed fare structures: Theory and applications

Received (in revised form): 7th April 2009

Thomas Fiig

is Chief Scientist in the Revenue Management Development department at Scandinavian Airlines System (SAS). He is responsible for developing methods and strategy for revenue management systems at SAS, including the overall design and methodologies of the O&D forecasting and optimization systems. His recent work has focused on methodologies for O&D optimization in semi-restricted fare structures and estimating price elasticities. Dr Fiig holds a PhD in Mathematics and theoretical Physics and a BA in Finance from the University of Copenhagen.

Karl Isler

is Head of Operations Research and Strategy in the Revenue Management, Pricing and Distribution department of Swiss International Airlines. He holds a PhD in Theoretical Physics from ETH Zurich. He developed the concepts for the integrated O&D pricing and inventory control strategy used by Swiss.

Craig Hopperstad

is currently president of Hopperstad Consulting. Previously, as Project Director in the Boeing Commercial Airplane Group, he was a principal in the development of passenger preference, fleet planning, scheduling and revenue management models. He is the author of numerous papers and presentations, many of which deal with the application of the Passenger Origin/Destination Simulator (PODS), which he developed at Boeing and for which Hopperstad Consulting now holds a license.

Peter Belobaba

is Principal Research Scientist at the Massachusetts Institute of Technology (MIT), where he teaches graduate courses on The Airline Industry and Airline Management. He is Program Manager of MIT's Global Airline Industry Program and Director of the MIT PODS Revenue Management Research Consortium. Dr Belobaba holds a Master of Science and a PhD in Flight Transportation Systems from MIT. He has worked as a consultant on revenue management systems at over 40 airlines and other companies worldwide.

Correspondence: Peter Belobaba, MIT International Center for Air Transportation, 77 Mass. Ave., Room 33-215, Cambridge, MA 02139 USA
E-mail: belobaba@mit.edu

ABSTRACT This paper develops a theory for optimizing revenue through seat inventory control that can be applied in a variety of airline fare structures, including those with less restricted and fully undifferentiated fare products that have become more common in the recent past. We describe an approach to transform the fares and the demand of a general discrete choice model to an equivalent independent demand model. The transformation and resulting fare adjustment approach is valid for both static and dynamic optimization and extends to network revenue management applications. This transformation allows the continued use of the optimization algorithms and seat inventory control mechanisms of traditional revenue management systems, developed more than two decades ago under the assumption of independent demands for fare classes.

Journal of Revenue and Pricing Management (2010) 9, 152–170. doi:10.1057/rpm.2009.18;
published online 19 June 2009

Keywords: seat inventory control; airline fare structures; network revenue maximization; O-D control; marginal revenue transformation; DAVN-MR, PODS



INTRODUCTION AND MOTIVATION

Traditional airline revenue management (RM) systems, developed in the late 1980s, make use of forecasting and optimization models that assume independent demands for each fare class on a flight leg, and/or for each passenger itinerary (path) and fare class in an airline network. This assumption of demand independence was facilitated by airline fare structures characterized by multiple fare products, each with different restrictions such as minimum stay requirements, advance booking and ticketing rules, cancellation and change penalties, as well as non-refundability conditions. That is, passengers that purchase a given fare type are assumed to be willing to purchase only that particular fare type, with no possibility of choosing a lower fare, another itinerary (path) or another airline. Although this assumption of demand independence was never completely accurate, virtually all airline RM forecasting and optimization models were developed on this basis.

With changes to airline pricing practices in recent years, what was a questionable yet tolerable assumption for RM models became almost completely invalid. The emergence and rapid growth of low cost carriers (LCC) with less restricted fare structures has led to the introduction of a variety of 'simplified' fare structures with fewer restrictions and in some cases no restrictions at all. Without modifications to traditional RM forecasting and optimization models, along with their associated seat inventory control mechanisms, less restricted fare structures result in more passengers (higher load factors) purchasing lower fares (lower yields), but lower total revenues for the airline.

This paper develops a theory for optimizing revenue through seat inventory control in a completely arbitrary airline fare structure. Examples that will be considered here include independent product-oriented demand (also known as 'yieldable' demand) in traditional differentiated fare structures; price-oriented (that is, 'priceable') demand with undifferen-

tiated fares; and hybrid demand in fare structures with a mixture of differentiated and undifferentiated fare products.

The applicability of this approach to any fare structure (and hence the associated passenger choice demand) might suggest that the optimization theory is complex. *On the contrary* – we will derive a marginal revenue transformation of the fares and the choice demand to an equivalent independent demand model. This transformation allows the continued use of existing optimization algorithms and control mechanisms in traditional RM systems. The marginal revenue transformation can be applied to all existing leg-based and O-D control RM systems, including EMSRb leg-based control, Displacement Adjusted Virtual Nesting (DAVN) and Dynamic Programming (DP), as will be illustrated.

This paper restricts its discussion to the optimization theory and thus assumes that we are able through some other means to forecast the demand of the choice of fare products in each fare structure, as affected by both passenger willingness to pay and the RM system's inventory controls. Demand forecasts that reflect fare product choice require the estimation of demand elasticity and sell-up parameters using separate estimation models, as described in Hopperstad (2007) and Guo (2008), for example.

The contribution of this paper is therefore in the presentation of a new theory for revenue management optimization that can be applied to any airline fare structure. The practical importance of this contribution is substantial – use of this new optimization theory can allow airlines to continue using traditional RM systems and seat inventory control mechanisms originally developed under the assumption of independent fare product demands.

After a brief review of relevant revenue management literature in the next section, we define the current optimization problem in the subsequent section and present a general formulation in the section after that. We first assume a number of simplifications: static optimization (no time dependence), determi-

nistic demand (no stochasticity), single airline leg (no network) and a single market (one specific price–demand relation). The theory is by no means limited by any of these simplifications, which are only introduced to present the theory more clearly. We then discuss specific applications of the optimization theory to different fare structures. In the section ‘Extensions to stochastic models’ we show that the theory extends beyond the simplifications made above to stochastic models, including dynamic programming. Simulated revenue results from the Passenger Origin Destination Simulator (PODS) are presented in the section ‘Simulation results using PODS’.

LITERATURE REVIEW

Most revenue management techniques were developed under a critical assumption that demand for a given fare class is independent of the demand for other fare classes. The effective use of these revenue management (RM) systems by airlines worldwide has been estimated to generate on the order of 4–6 per cent in incremental revenues (Smith *et al*, 1992). Barnhart *et al* (2003) describe the evolution of revenue management systems since the early 1980s. McGill and van Ryzin (1999) provide a comprehensive survey of revenue management literature, focusing on the evolution of forecasting and optimization models.

The optimization of airline seat inventory to maximize revenue given multiple fare products on a single flight leg can be traced back to Littlewood (1972), who solved the fare class mix problem for two nested fare classes. Belobaba (1987, 1989) extended the nested seat allocation problem to multiple fare classes with the development of the Expected Marginal Seat Revenue (EMSR) heuristic. Curry (1990), Wollmer (1992) and Brumelle and McGill (1993) described optimal solutions for multiple nested fare classes. Belobaba (1992) then modified the EMSR heuristic to better approximate the optimal solution to the multiple nested fare class problem. The modified

version became known as ‘EMSRb’, and is also described in detail in Belobaba and Weatherford (1996). Dynamic programming models have also been applied to the single-leg seat inventory control problem, by Lee and Hiersch (1993) and Lautenbacher and Stidham (1999), among others.

Network RM systems (or origin–destination controls) use network optimization mathematics to determine fare class seat availability based on paths (that is passenger itineraries), not by individual leg, making them particularly valuable to airlines that operate networks of connecting flights. Smith and Penn (1988) implemented one of the first O–D control methods, virtual bucketing, at American Airlines. Many network RM systems evolved from this initial concept, leading to the development of Displacement Adjusted Virtual Nesting (DAVN). DAVN adjusts the value of each path’s O–D fare for the estimated revenue displacement of a passenger on the connecting leg (Saranathan *et al*, 1999). Williamson (1992) and Vinod (1995) offer more detailed descriptions of O–D control and specifically the DAVN approach.

An alternative mechanism for O–D control that makes use of similar network optimization logic is ‘Bid-Price Control’. Developed by Simpson (1989) and elaborated by Williamson (1992), this method requires that the airline store a bid-price value (an estimate of the marginal network revenue value of each incremental seat) for each leg in the network. For each itinerary request, seat availability is determined by comparing the relevant O–D fare to the sum of the bid prices on the legs that the path traverses. Algorithms for calculating bid prices in an airline network include mathematical programming models by Talluri and Van Ryzin (1998), heuristic approaches by Belobaba (1998), and a probabilistic convergence algorithm by Bratu (1998). Dynamic programming can also be used to determine optimal network bid prices, in theory at least, as formulated by Gallego and Van Ryzin (1997). In practice, the ‘curse of



dimensionality' makes the size of the network DP problem infeasible to solve for realistic airline networks. Instead, some network RM systems apply DP methods to single flight legs to determine a vector of bid prices, in conjunction with a mathematical programming approach for network optimization.

The implementation of network RM systems by airlines has been estimated to generate an additional 1–2 per cent in total revenues, above the 4–6 per cent realized from the traditional leg-based RM systems (Belobaba, 1998). Yet, virtually all of the optimization models employed by these systems were based on the assumption of independent demands, first by flight leg and fare class, and then by passenger itinerary (path) and fare class. None of the traditional optimization models described in the literature explicitly addresses the complications caused by the move to less restricted fare structures since 2000. As noted by Tretheway (2004), the introduction of these simplified fare structures 'has undermined the price discrimination ability of the full service network carrier, and is the most important pricing development in the industry in the past 25 years.'

Without adequate demand segmentation restrictions, simplified fare structures allow passengers that previously would buy higher fares to purchase lower fare classes. The RM system's historical booking database then records fewer bookings in the higher fare classes which, in turn, leads to lower forecasts of future demand for higher fare types. The seat allocation optimizer then protects fewer seats for higher fare classes and makes more seats available to lower booking classes, which encourages even more high-fare demand to buy down. A mathematical model that describes this 'spiral down' effect was developed by Cooper *et al* (2006).

Several works describing methods for reversing this spiral down effect have focused on modifications to the forecasting models used in traditional RM systems. Boyd and Kalllesen (2004) proposed a segmentation between

price-oriented (or 'priceable') demand that will always purchase the lowest available fare and product-oriented (or 'yieldable') demand that is willing to buy a higher fare due to specific attributes or lack of restrictions. They suggested that identification of price- vs product-oriented demands in an RM database could enable the airline to generate separate forecasts of each type of demand, and to combine these forecasts into a single 'hybrid' forecast. While traditional time series forecasting models can be applied to the product-oriented demand data, a forecasting approach that includes estimates of willingness to pay (or sell-up potential) is required to forecast price-oriented demand for each fare class. Belobaba and Hopperstad (2004) describe such a model, called 'Q-forecasting', which can provide the estimates of price-oriented demand required for hybrid forecasts.

Talluri and van Ryzin (2004) formulated a DP model for the case when consumer behaviour is described by a general discrete choice model, and showed that only strategies on the efficient frontier are relevant in the optimization. In this paper, we show that any optimization algorithm for an airline network can be extended to a completely general fare structure, as developed and first presented by Fiig *et al* (2005) and Isler *et al* (2005). The approach, also referred to as 'fare adjustment', makes use of a marginal revenue transformation that adjusts the fare and demand inputs to traditional RM optimizers, effectively transforming the choice model representing a general fare structure into independent demand of a fully differentiated fare structure.

THE OPTIMIZATION PROBLEM

In this section we describe the general optimization problem and define the notation to be used throughout the paper. We consider the simplest possible example to develop the optimization theory, that is, a single leg with capacity *cap*. We begin with the solution of the optimization problem when fares are fully




Figure 1: (a) Incremental revenue as function of demand (quantity). Total revenue shown as shaded area below curve. (b) Total revenue versus total demand (quantity).

Table 1: Data for example in Figures 1 and 2

Fare product	Fare	Demand		Fully differentiated		Fully undifferentiated	
	f_i	d_i	Q_i	TR_i	MR_i	TR_i	MR_i
1	\$1200	31.2	31.2	\$37 486	\$1200	\$37 486	\$1200
2	\$1000	10.9	42.1	\$48 415	\$1000	\$42 167	\$ 428
3	\$800	14.8	56.9	\$60 217	\$800	\$45 536	\$228
4	\$600	19.9	76.8	\$72 165	\$600	\$46 100	\$28
5	\$400	26.9	103.7	\$82 918	\$400	\$41 486	\$(172)
6	\$200	36.3	140.0	\$90 175	\$200	\$28 000	\$(372)

differentiated (that is independent or ‘product-oriented’ demand). The airline has published a set of fares f_i , in decreasing fare order: $f_i > f_{i+1}$, $i = 1, \dots, n-1$. We assume that demand is deterministic, with demand d_i for each fare product.

In a fully differentiated fare structure, traditional RM models assume that the demand for each fare product is independent of whether other products are available. As shown in Figure 1(a), the revenue maximizing solution is to allocate seats for all demands in decreasing fare order until capacity is reached. The seats available to the lowest fare product must be limited by capacity. The step-wise decreasing curve is the incremental revenue as function of quantity (Q). The capacity constraint is shown as the vertical line at $cap = 100$. Note that the data behind Figures 1 and 2 are summarized in Table 1.

Figure 1(b) illustrates the quantities TR_k and Q_k , denoting the total revenue and the total quantity sold given k is the lowest fare level available, as

$$Q_k = \sum_{j=1}^k d_j \text{ and } TR_k = \sum_{j=1}^k f_j d_j.$$

The total revenue increases monotonically with quantity, calculated as the area below the curve in Figure 1(a).

Now assume that the restrictions on the fare products are removed, making them fully undifferentiated (that is, ‘priceable’ demand). The only distinction between the fare products is their price levels and therefore customers will always buy the lowest fare available. Even customers willing to buy the higher fares will now buy the lowest available fare. The total revenue and the total quantity sold given k is the

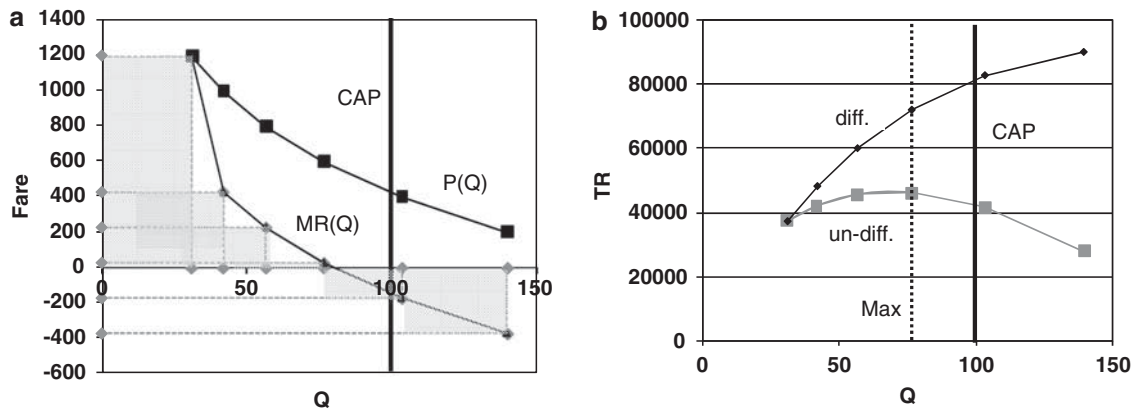


Figure 2: (a) The demand curve $P(Q)$ and marginal revenue $MR(Q)$ for the undifferentiated case. (b) $TR(Q)$ for both differentiated and undifferentiated case.

lowest fare available can now be expressed as

$$Q_k = \sum_{j=1}^k d_j \text{ and } TR_k = f_k Q_k.$$

The total revenue is shown for both the differentiated and undifferentiated case in Figure 2. With all passengers buying the lowest available fare, the total revenue for the undifferentiated case falls below that of the differentiated fare structure.

The incremental demand accommodated by opening up fare class k is d_k . However the incremental revenue is not $d_k f_k$ but a reduced amount, which we can calculate by accounting for the revenue lost due to buy down. The loss from buy-down is given by $Q_{k-1}(f_{k-1} - f_k)$, since the aggregated demand that previously bought class $k-1$ now will buy class k . The incremental revenue contribution by opening up class k can therefore be written as $d_k \times f_k - Q_{k-1}(f_{k-1} - f_k)$. Another way to state this is that, instead of receiving the fare f_k by opening up class k , we obtain an adjusted fare $f'_k = [d_k \times f_k - Q_{k-1}(f_{k-1} - f_k)] / d_k$. The adjusted fare is also known as the marginal revenue MR_k and may be calculated as the increment in total revenue per capacity unit:

$$MR_k = \frac{TR_k - TR_{k-1}}{Q_k - Q_{k-1}} = f'_k$$

The marginal revenue is illustrated by the lower curve, $MR(Q)$, on Figure 2(a). Note that the area below the curve corresponds to the total revenue. As more seats are made available to the lower fare, the marginal revenue becomes negative, meaning the total revenue decreases. Therefore the fare classes below \$600 in this example should always remain closed.

We can therefore summarize the solution principle of the revenue optimization problem as follows: *Order the fares in decreasing marginal revenue and open fares until capacity is reached or the marginal revenue becomes negative.* This general principle is valid for both differentiated and undifferentiated fare structures.

GENERAL FORMULATION

In this section we develop the theory for a completely general fare structure, which may have any set of restrictions attributed to each fare product. The fare and the restrictions of a fare product determine the probability that a booking occurs for that product, which will generally depend on what other products are available as well.

An optimization strategy consists of opening up a given set of fare products, Z . Examples include opening all products $N = \{1, 2, \dots, n\}$; or keeping all products closed $\{\}$. In the previous example, we considered the subsets of opening fare products in decreasing fare

order $\{\}; \{1\}; \{1, 2\}, \dots, \{1, 2, \dots, n\}$, which are ‘nested’ strategies. If we now allow a strategy to be any set of fare classes open, there will be 2^n different strategies, $Z \subseteq N$.

The demand for a fare product depends in general on the particular optimization strategy Z . The demand for fare product j given strategy Z will be denoted $d_j(Z)$. The total demand for a given strategy is calculated as $Q(Z) = \sum_{j \in Z} d_j(Z)$. The corresponding total revenue for a given strategy (assuming no capacity constraint) is given by $TR(Z) = \sum_{j \in Z} d_j(Z) f_j$.

The general optimization problem can then be stated as: $\max TR(Z)$ subject to $Q(Z) \leq \text{cap}$. The solution to the problem can most easily be solved by analyzing $TR(Z)$ versus $Q(Z)$ in a scatter plot where each of the 2^n different strategies are plotted. These correspond to what we can call ‘pure’ strategies. In addition to the pure strategies, we may construct ‘mixed’ strategies by combining the pure strategies such that each of the pure strategies are only used for a fraction of the time.

The mixed strategies trace out the convex hull of the pure strategies. This is marked in Figure 3 as the shaded area, meaning that any point within the shaded area is feasible. The solution to the optimization problem $\max TR(Z)$ subject to $Q(Z) \leq \text{cap}$ is identified as the upper boundary on the convex hull. This portion marked with the solid lines is termed the ‘efficient frontier’. Let the subset of strategies that lies on the efficient frontier (efficient sets) be denoted S_0, S_1, \dots, S_m . Here $S_0 = \{\}$ is the empty set with $TR(S_0) = 0$ and $Q(S_0) = 0$. The total demand and total revenue of strategy S_k is denoted by Q_k and TR_k respectively. Note that Figure 3 illustrates the general case where the optimal solution is obtained by a mixed strategy.

These sets are ordered according to quantity $Q_{k-1} \leq Q_k$ and $TR_{k-1} \leq TR_k$ for $k = 1, \dots, m$. The end-point for the efficient frontier is strategy S_m , representing the maximum total revenue.

The first non-empty set is in general $S_1 = \{1\}$, the highest fare in the fare structure, since it has the highest value slope, $TR(Z)/$




Figure 3: Construction of convex hull.

$Q(Z)$, which is a weighted average over the fares in Z .

The marginal revenue transformation maps the original fare structure into the transformed fare structure. The demands for the transformed products are given by the marginal demand d'_k while the fares of the transformed fare products are given by the marginal revenue.

$$d'_k = Q_k - Q_{k-1}, \quad k = 1, \dots, m$$

$$f'_k = (TR_k - TR_{k-1}) / (Q_k - Q_{k-1}), \quad k = 1, \dots, m$$

The marginal demand associated with each strategy k is also called the *partitioned* or *incremental* demand. The marginal revenues are called *adjusted fares*. The transformation is identical to the previous MR expression. The usefulness of the marginal revenue transformation follows from the following theorem, which we have proven by direct construction of the convex hull (since the independent demand model with demand d'_k and fares f'_k would trace out exactly the same convex hull and are therefore equivalent).

Theorem: The marginal revenue transformation transforms the original fare structure into an equivalent model with independent fare products.

Implementation of control mechanism using fare adjustment

We have shown in the last section that the marginal revenue transformation transforms the optimization problem for an arbitrary fare



structure to an equivalent independent demand model. But how is the booking control mechanism transformed? A booking control mechanism in traditional RM systems calculates numerical seat availability for each booking class based on the solution of the optimization problem and the current seat inventory. However, in the transformed model the strategies on the efficient frontier can involve correspond to a subset of the original booking classes, but not necessarily all classes in decreasing fare order.

Consider for example an efficient frontier with the efficient strategies $\{1, 2\}$ at capacity 1, $\{1, 3\}$ at capacity 2, while $\{1, 2, 3\}$ is inefficient. If the capacity is 2, a nested booking limit control would protect 1 seat for $\{1, 2\}$ and allow 1 booking for $\{1, 3\}$, thus the seat availability for $\{1, 2\}$ would be 2 and for $\{1, 3\}$ 1 correspondingly. It is not possible to map this to seat availability for the original classes without opening the inefficient strategy $\{1, 2, 3\}$. New distribution mechanisms would be required, where the exact number of requested seats has to be known in advance and the corresponding strategy can be shown accordingly.

Many choice models have the desirable property that the strategies on the efficient frontier are nested, $S_k \subset S_l$, $k < l$, for example $\{1\}$, $\{1, 3\}$, $\{1, 2, 3\}$. Note that the nesting does not necessarily have to follow the fare order. On such nested efficient frontiers, a class once contained in an efficient strategy will also be contained in any of the next efficient strategies as capacity increases. It is therefore possible to map the original booking classes to the strategies in which they occur first with increasing capacity (or going from left to right on the efficient frontier).

When the efficient frontier is nested, the marginal revenue transformation can therefore be interpreted as the assignment of demands and fares to the original booking classes which, when fed into the optimization algorithm and the control mechanism for independent demand, produces the correct availability control. This ease of implementing the control

mechanism suggests that we will in practice apply a nested approximation, even in cases where the nesting property is not exactly satisfied.

This simplification also applies to bid price control. Here, the marginal revenue going from one strategy to the next can be assigned to the newly added classes. The bid price control is therefore modified by using the marginal revenue associated with the booking class rather than the fare itself.

APPLICATIONS TO DIFFERENT FARE STRUCTURES

In this section, we describe the application of the marginal revenue transformation and fare adjustment theory to different airline fare structures – from traditional differentiated fares to completely undifferentiated fares, along with ‘hybrid’ fare structures that include both types of fare products.

Fully differentiated fare products

Under the traditional RM assumption of fare class independence, the fare products are assumed to be adequately differentiated, such that demand for a particular fare product will only purchase that fare product. Therefore, for any set of fare products with fares f_j , and demands d_j , $j = 1, \dots, n$, the marginal revenue transformation yields

$$\begin{aligned} d'_k &= Q_k - Q_{k-1} \\ &= \sum_{j=1, \dots, k} d_j - \sum_{j=1, \dots, k-1} d_j = d_k \end{aligned}$$

and

$$\begin{aligned} f'_k &= \frac{TR_k - TR_{k-1}}{Q_k - Q_{k-1}} \\ &= \frac{\sum_{j=1, \dots, k} d_j f_j - \sum_{j=1, \dots, k-1} d_j f_j}{d_k} = f_k \end{aligned}$$

$k = 1, \dots, n$. Thus for the differentiated fare product case, the marginal revenue transformation acts as an identity and both demands and fares are unaffected. We can refer to this

demand as being ‘product-oriented’. Let us denote the unadjusted fare by f_k^{prod} .

Completely undifferentiated fare products

We now apply the marginal revenue transformation to a single flight leg with a fully undifferentiated fare structure, again assuming deterministic demand. The definitions of fares f_j , and demands d_j , $j = 1, \dots, n$ are unchanged. Due to the fully undifferentiated fare structure, passengers will only buy the lowest available fare. Let k be the lowest open fare product, then the demand for all other fare products becomes zero, $d_j = 0$ for $j \neq k$. Hence the cumulative demand equals the demand for the lowest open fare product $d_k = Q_k$, $k = 1, \dots, n$.

The demand $Q_k = Q_n psup_k$ (assuming f_k is the lowest open) can be expressed as a sell-up probability $psup_k$ from $n \rightarrow k$, times a base demand Q_n for the lowest class n .

The marginal revenue transformation determines the adjusted fares. The demand is calculated as $d'_k = Q_k - Q_{k-1} = d_k - d_{k-1} = Q_n(psup_k - psup_{k-1})$.

The adjusted fares are given by

$$f'_k = \frac{TR_k - TR_{k-1}}{Q_k - Q_{k-1}} = \frac{f_k Q_k - f_{k-1} Q_{k-1}}{Q_k - Q_{k-1}} = \frac{f_k psup_k - f_{k-1} psup_{k-1}}{psup_k - psup_{k-1}}$$

Let us denote this by f_k^{price} .

Note that the highest fare value is unadjusted:

$$f'_1 = \frac{TR_1 - TR_0}{Q_1 - Q_0} = \frac{f_1 d_1 - 0}{d_1 - 0} = f_1$$

Consider the special case of exponential sell-up $Q_k = Q_n psup_k = Q_n \exp(-\beta(f_k - f_n))$ and assume for simplicity an equally spaced fare-grid with $f_{k-1} - f_k = \Delta$. The marginal revenue transformation yields:

$$d'_k = Q_n \exp(\beta f_n) (\exp(-\beta f_k) - \exp(-\beta f_{k-1}))$$

$$f'_k = f_k - f_M, \text{ where } f_M = \Delta \cdot \frac{\exp(-\beta \cdot \Delta)}{1 - \exp(-\beta \cdot \Delta)}$$

The adjusted fare can therefore be expressed as the original fare minus a positive constant called the *fare modifier*, f_M . For a general fare structure, the fare modifier will differ by class. Because the fare modifier is the cost associated with buy-down, it is also termed the *price elasticity cost*.

In Table 1, the constant fare modifier can be seen $f_M = \$572$. Since seats are only allocated to fare classes with positive adjusted fares, we can interpret f_M as the lowest open fare. In the limit $\Delta \rightarrow 0$, using l’Hopital’s rule we obtain $f_M \rightarrow 1/\beta$ for $\Delta \rightarrow 0$, which depends only on the beta parameter in the exponential demand model.

Hybrid demand in mixed fare structures

In this section we show how to optimize a single leg with a fare structure consisting of a mix of differentiated and undifferentiated fare products assuming deterministic demand, using the marginal revenue transformation. This is the so-called hybrid demand case. The fare class demands are decomposed into contributions from both differentiated (product-oriented) and un-differentiated (price-oriented) demand:

$$d_j = d_j^{prod} + d_j^{price}, j = 1, \dots, n$$

The marginal revenue transformation determines the adjusted fares. The demand for each fare class is obtained as

$$d'_k = Q_k - Q_{k-1} = \left(\sum_{j=1, \dots, k} d_j^{prod} + d_k^{price} \right) - \left(\sum_{j=1, \dots, k-1} d_j^{prod} + d_{k-1}^{price} \right) = d_k^{prod} + Q_n(psup_k - psup_{k-1})$$

Define $r_k = d_k^{prod} / (d_k^{prod} + Q_n(psup_k - psup_{k-1}))$ as the ratio of product demand to total demand.



The adjusted fare is given by

$$\begin{aligned}
 f'_k &= \frac{TR_k - TR_{k-1}}{Q_k - Q_{k-1}} \\
 &= \frac{\left(\sum_{j=1,\dots,k} d_j^{prod} f_j + d_k^{price} f_k\right) - \left(\sum_{j=1,\dots,k-1} d_j^{prod} f_j + d_{k-1}^{price} f_{k-1}\right)}{d_k^{prod} + Q_n(psup_k - psup_{k-1})} \\
 &= \frac{d_k^{prod} f_k + Q_n(f_k psup_k - f_{k-1} psup_{k-1})}{d_k^{prod} + Q_n(psup_k - psup_{k-1})} \\
 &= r_k f_k^{prod} + (1 - r_k) f_k^{price}.
 \end{aligned}$$

Thus the adjusted fare in the hybrid case is calculated by a demand-weighted average of the unadjusted fare used for the product-oriented demand and the adjusted fare for the price-oriented demand.

EXTENSIONS TO STOCHASTIC MODELS

In the previous sections, we considered deterministic models on a single flight leg, in order to present the theory more clearly. In this section, we show that the marginal revenue transformation also holds for stochastic models. Furthermore, it extends to network optimization provided the choice probabilities for the different O-D paths are independent.

Talluri and van Ryzin (1998) provided a consistent formulation of a stochastic network dynamic programming model, but it cannot be implemented in practice because of the curse of dimensionality. In practice, stochastic network optimization methods for airline revenue management typically require some heuristic aspects to generate approximate solutions. Therefore, if we show that the marginal revenue transformation holds exactly for the DP approach, we can argue that it can also be applied to any other solution method. We start with the one leg model first, followed by a straightforward generalization to the network problem with independent O-D paths.

DP formulation for a discrete choice model of demand

The dynamic program for one leg with an arbitrary fare structure, such that demand behaviour is modelled by a discrete choice model, was developed by Talluri and van Ryzin (2004). Dynamic programming models assume that the bookings can be modelled as a Markov process with small time steps such that the probability of more than one booking is negligible. The arrival probability of a booking is denoted by $\lambda_t \geq 0$ and the remaining capacity x_t for each t serves as state variable. At the beginning of each time interval, the airline can choose the set of products available for sale. A strategy therefore consists of the set of open classes $Z_t(x) \subseteq N$, where N is the set of all classes.

Given an arrival and strategy Z_t , the probability that product $j \in Z_t$ is chosen is denoted by $p_j(Z_t)$. The expected demand for product j in time slice t is therefore $d_j(Z_t) = \lambda \cdot p_j(Z_t)$, which is the probabilistic analogy of the deterministic models of the sections ‘General formulation’ and ‘Applications to different fare structures’. The DP model can handle time dependent choice probabilities and arrival rates as well, but for the sake of simplicity, we did not introduce t indices for the choice probabilities and the arrival rate.

The probability that a booking occurs for any product, given an arrival and choice set Z , is therefore

$$Q(Z) = \sum_{j \in Z} p_j(Z)$$

and the corresponding expected total revenue is

$$TR(Z) = \sum_{j \in Z} p_j(Z) \cdot f_j$$

Bellman equation

The dynamic program proceeds by considering the concept of expected optimal future revenue, or revenue to go, $J_i(x)$. The goal is to find

the strategy which maximizes $J_0(cap)$, the optimal revenue at the beginning of the booking process. The Bellman equation for our problem now states that the optimal future revenue satisfies the recursion relation:

$$J_{t-1}(x) = \max_{Z \subseteq N} \{ \lambda \cdot (TR(Z) + Q(Z) \cdot J_t(x - 1)) + (1 - \lambda \cdot Q(Z)) \cdot J_t(x) \},$$

with boundary conditions $J_T(x) = J_t(0) = 0$. The first term in the above corresponds to the case when a sell occurs; the airline receives one booking from the set of open classes in Z , and continues with one seat less. The second term corresponds to the no-sell situation, where the airline continues with the same amount of seats. The above optimization problem has to be solved for each x and t in order to find the optimal strategy Z^* . Introducing the bid-price vector

$$BP_t(x) = J_t(x) - J_t(x - 1)$$

one can write the optimization problem as:

$$J_{t-1}(x) = J_t(x) + \lambda \cdot \max_{Z \subseteq N} (TR(Z) - BP_t(x) \cdot Q(Z)) \quad (1)$$

Observe now that for all x and t , this problem is of the same form, namely

$$\max_{Z \subseteq N} (TR(Z) - BP \cdot Q(Z)) \quad (2)$$

where only the value of the bid-price may change.

Efficient frontier

The problem of optimizing $\max_{Z \subseteq N} (TR(Z) - BP \cdot Q(Z))$ is most easily analyzed by drawing $TR(Z)$ versus $Q(Z)$ in a scatter plot for all possible strategies of open classes, together with the straight line $BP \cdot Q$ for specific values of the bid-price. For comparison, the plot in Figure 4 is identical to Figure 3 describing the efficient frontier for the deterministic case.

The optimum for the specific BP in the chart is attained where the parallel shifted line $BP \cdot Q$ is tangent to the convex hull of the




Figure 4: Construction of convex hull.

points $(Q(Z), TR(Z))$. This strategy is marked as S_i in the chart.

Since the bid-price is always positive, only strategies which lie on the efficient frontier from S_0 to S_m the set with maximal total revenue can be solutions to the problem. We refer to the discussion regarding the efficient strategies from the deterministic case, which translates identically to the DP case.

Marginal revenue transformation

The optimal strategy as function of the bid-price can only consist of one of the vertices on the efficient frontier. As long as the bid-price is higher than the highest fare, we should close all classes.¹ We denote this strategy by S_0 . If the bid-price falls below f_1 the optimal strategy is initially S_1 , until the bid-price falls below the slope of the convex hull between S_1 to S_2 . The slopes of the efficient frontier correspond to the critical values of the bid-price, where the optimal strategy changes. The optimal strategy for a given bid price BP is found by walking along the efficient frontier from one strategy to the next, increasing sales volume, as long as

$$MR_k = \frac{TR(S_k) - TR(S_{k-1})}{Q(S_k) - Q(S_{k-1})} \geq BP.$$

The last strategy obtained this way is the optimal strategy of open booking classes for bid price BP .

Once the efficient frontier is known, the bid-prices can be calculated from the DP formulation recursively. Because the efficient frontier contains only a few of the possible



strategies, there are potentially many fare structures with corresponding choice models (identified by a set of $p_j(Z)$) which will lead to the same efficient frontier. All these models will have the same bid-price. In particular consider the independent demand model defined by:

$$p'_k = Q(S_k) - Q(S_{k-1}), \quad k = 1, \dots, m.$$

$$f'_k = MR_k = \frac{TR(S_k) - TR(S_{k-1})}{Q(S_k) - Q(S_{k-1})}, \quad k = 1, \dots, m$$

This is identical to the marginal revenue transformation introduced in the section ‘General formulation’. Using the transformed choice model (primed demand and fares) in an independent demand DP instead of the original choice model DP, the Bellman equation will produce the same bid-prices. Furthermore, as explained in the section ‘General formulation’, we can associate the marginal revenue or adjusted fares f'_k to the newly added classes in $S_k \setminus S_{k-1}$ provided the efficient frontier is nested.

Network DP formulation for a discrete choice model of demand

Let us also touch briefly on the network DP formulation of Talluri and van Ryzin (1998), generalized to independent discrete choice models for each itinerary. In each time slice there is at most one request for the whole network and the rate λ is the probability that the request is for a given itinerary. The state becomes a vector (x_1, \dots, x_L) of the remaining capacities for each leg. The second term on the right hand side of equation (1) above becomes a sum over all itineraries, where the bid-price for an itinerary is the difference in expected future revenue if on each leg one seat is removed. An optimal strategy for each itinerary has to be chosen. The optimization problem therefore decomposes into a set of

independent optimizations of the form (2) for each itinerary.

The DP network model has no practical importance, since the number of possible states explodes with cap^L . However the marginal revenue transformation holds as well. This leads us to conclude that any heuristic network algorithm which assumes independent demand by O-D itinerary can be easily extended to deal with arbitrary fare structures by applying the marginal revenue transformation. This will be illustrated in the example with DAVN-MR below.

Recover pricing DP model as a special case

Let us consider as an example the case of a fully undifferentiated fare structure. This model is nested by fare order and we can therefore identify the efficient strategy index with the original booking class index. A strategy S_k therefore has class k and all classes with higher fares open. If we insert the transformed demand and transformed fares into the Bellman recursion formula for a DP with independent demand, we obtain the following problem.

$$\begin{aligned} & \max_S \sum_{i \in S} p'_i (f'_i - BP) \\ & = \max_k \sum_{i=1}^k (Q(S_i) - Q(S_{i-1})) \\ & \quad \cdot \left(\frac{TR(S_i) - TR(S_{i-1})}{Q(S_i) - Q(S_{i-1})} - BP \right) \\ & = \max_k [Q(S_k)(f_k - BP)] \end{aligned}$$

In the last equation we have used that we have a telescopic sum and that all demand resides in the lowest open class and thus $TR(S_k) = f_k Q(S_k)$. The last expression is identical to the pricing DP model of Gallego and Van Ryzin (1997). The first optimization problem above for independent demand is

easily solved since only classes where $f'_k \geq BP$ should be included in the strategy S . If we further insert the exponential sell-up with equally spaced fares we obtain

$$\begin{aligned} & \max_S \sum_{i \in S} p'_i(f'_i - BP) \\ & = \max_k [\exp(-\beta \cdot (f_k - f_n)) \cdot (f_k - BP)] \end{aligned}$$

We have shown in the section ‘Applications to different fare structures’ that the fare modifier for exponential sell-up and equally spaced fare grid is a constant. Thus the open classes are exactly those which have $f_i - f_M \geq BP$, a result which would not be immediately clear from optimization problem on the right hand side.

Leg RM: EMSRb-MR

Here we show how EMSRb can be extended to optimize a general fare structure. First we will briefly consider the results for standard unadjusted EMSRb, in order to be able to compare with the transformed EMSRb-MR. In standard EMSRb we consider independent fare products with fares f_i , with normally distributed demands $d_i \sim N(\mu_i, \sigma_i^2)$, $i = 1, \dots, n$. The aggregated demand for fare product $i = 1, \dots, k$ is also normally distributed with average $\mu_{1,k} = \sum_{i=1}^k \mu_i$, variance $\sigma_{1,k}^2 = \sum_{i=1}^k \sigma_i^2$ and average fare $\bar{f}_{1,k} = \mu_{1,k}^{-1} \sum_{i=1}^k \mu_i f_i$.

The joint protection level π_k is then given by $\pi_k = \mu_{1,k} + \sigma_{1,k} \Phi^{-1}(1 - f_{k+1}/\bar{f}_{1,k})$, and the booking limit by $BL_k = cap - \pi_k$.

For a general fare structure we need to apply the formalism presented in the previous sections, which consists of the following steps:

- (1) Determine the strategies $k = 1, \dots, m$ on the efficient frontier.
- (2) Apply the marginal revenue transformation to both demands and fares.
- (3) Apply EMSRb in the normal fashion using the transformed demands and fares.

This method, called EMSRb-MR, thus considers independent normally distributed demands $d'_k \sim N(\mu_k, \sigma_k^2)$, with fares f'_k , $k = 1, \dots, m$. Note that the demands are related to the strategies not the individual fare products. The aggregated demand, variance, and average fares are calculated analogously using the transformed demands and fares and hence the formulas for the joint protection level and the booking limit become exactly the same: $\pi'_k = \mu_{1,k} + \sigma_{1,k} \Phi^{-1}(1 - f'_{k+1}/\bar{f}'_{1,k})$, and $BL'_k = cap - \pi'_k$.

Consider again the case of fully undifferentiated fare products with equal spaced fares, as presented in Table 1 and in section ‘Applications to different fare structures’. In this case the strategies become nested and we may associate an adjusted fare with each strategy. For simplicity (and ease of comparison) we will assume that the demands and standard deviations are identical to those of the fully differentiated case.

Table 2 compares the standard EMSRb results in the case of differentiated fares with

Table 2: Comparison of Standard EMSRb and EMSRb-MR Booking Limits

Fare product	Fare value	Mean Demand	Standard Deviation	EMSRb Limits	Adjusted Fares (MR)	EMSRb-MR Limits
1	\$1200	31.2	11.2	100	\$1200	100
2	\$1000	10.9	6.6	80	\$428	65
3	\$800	14.8	7.7	65	\$228	48
4	\$600	19.9	8.9	46	\$28	16
5	\$400	26.9	10.4	20	\$(172)	0
6	\$200	36.3	12.0	0	\$(372)	0

The bold columns contain the booking limits being compared.



the modified booking limits coming from the EMSRb-MR approach in the case of undifferentiated fares. Note that the adjusted fares used as input to EMSRb after the marginal revenue transformation lead to closure of the lowest two classes with negative marginal revenues. Classes 2 through 4 also get reduced seat availability, given that the estimated potential for sell-up has been incorporated into the marginal revenue values of each class.

The adjusted EMSRb-MR overcomes spiral down by modifying the resulting booking limits to account for the estimated sell-up potential, and thereby the dependent demand forecasts, which results in closing availability to the inefficient fare products (that is with negative adjusted fares).

Network RM: DAVN-MR with differentiated and undifferentiated fare structures

For a large network airline, it is common to offer a variety of fare structures in different O-D markets, each with different restrictions, as required to respond to competitors in each market. However, individual flight legs to or from a connecting hub provide seats to multiple O-D markets, each of which can have different fare structures. In this section we show how to optimize a network with a mix of differentiated fare products and undifferentiated fare products. As an example we will use DAVN (Displacement Adjusted Virtual Nesting).

The marginal revenue transformation applied to DAVN (termed DAVN-MR) results in transforming both demands and fares as described in the previous sections. A particularly simple expression for the adjusted fare is obtained by assuming exponential sell-up and equally spaced fares, as shown previously. The adjusted fare becomes

$$adjf = f' - \sum DC = f - f_M - \sum DC.$$

Here $\sum DC$ refers to the standard DAVN approach of calculating the marginal network

revenue contribution of accepting an O&D passenger by subtracting the displacement costs of carrying a passenger on connecting legs. Further, we have used the previous result that the adjusted fare can be expressed as the fare minus a fare modifier. Thus this formula extends the standard DAVN formula by just an additional term $-f_M$. Note that $f_M = 0$ for differentiated fare products. The undifferentiated fare products are mapped to lower buckets since $f_M > 0$ (except the highest undifferentiated fare class). Also note that all undifferentiated fare products with fares $f < f_M$ will be closed (regardless of remaining capacity) since they map to negative adjusted fares. DAVN-MR is then applied by running DAVN in the standard fashion on the transformed fare structure to produce booking limits by bucket (in marginal revenue scale) for each leg. Finally, note that the formula for the adjusted fare is valid for an arbitrary fare structure with nested efficient frontier, but there the fare modifier will vary by class.

SIMULATION RESULTS USING PODS

In this section, we present simulation results from PODS (Passenger Origin Destination Simulator) to illustrate the revenue impacts of applying the marginal revenue transformation in DAVN-MR, as described above. PODS and its various components have been widely

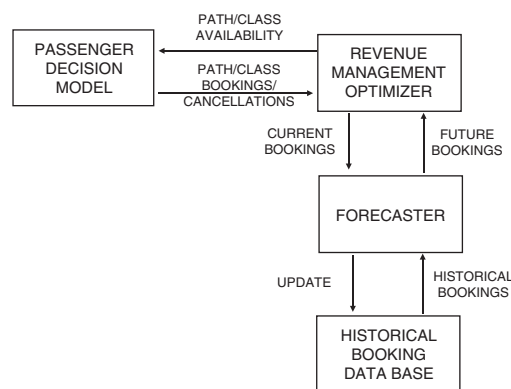


Figure 5: PODS simulation architecture.

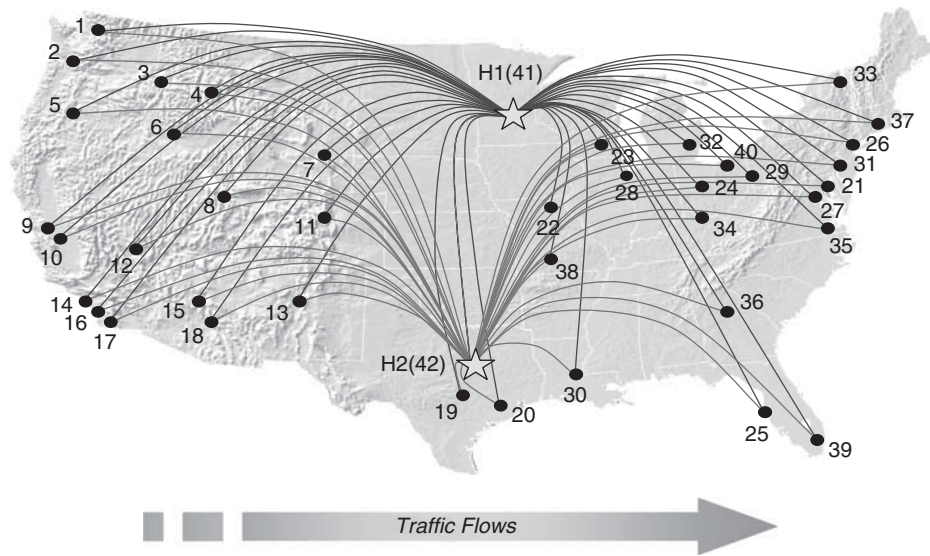


Figure 6: PODS network D.

UNDIFFERENTIATED					SEMI-DIFFERENTIATED				
FARE CLASS	AP	Min Stay	Cancel Fee	Non Refund	FARE CLASS	AP	Min Stay	Cancel Fee	Non Refund
1	0	NO	NO	NO	1	0	NO	NO	NO
2	0	NO	NO	NO	2	0	NO	YES	NO
3	0	NO	NO	NO	3	0	NO	YES	YES
4	0	NO	NO	NO	4	0	NO	YES	YES
5	0	NO	NO	NO	5	0	NO	YES	YES
6	0	NO	NO	NO	6	0	NO	YES	YES

Figure 7: Undifferentiated and semi-differentiated fare structures.

described in the literature – see for example, Hopperstad (1997), Belobaba and Wilson (1997), Lee (1998), and Carrier (2003). The central concept in PODS is that the models of the revenue management systems and the model of passenger choice are contiguous but independent. That is, the revenue management system(s) define the availability of path/classes in a market and the generated passengers pick from the available path/classes and book. The only information available to the revenue management system(s) of the simulated airlines for setting path/class availability is the history of (previously simulated) passenger bookings. Figure 5 illustrates this basic architecture.

In PODS, each simulated passenger chooses among available path/class alternatives (that is, itinerary and fare product combinations) with fares lower than that passenger’s randomly drawn maximum willingness-to-pay (WTP). Business passengers have a higher input mean WTP than leisure travellers. If there exist more than one available path/class meeting this WTP criterion, the passenger then chooses the alternative with the lowest generalized cost, defined as the sum of the actual fare and the disutility costs with the restrictions (if any) associated with each fare class. A passenger will ‘sell up’ from a preferred lower fare class (with a lower generalized cost) to a more expensive

one, therefore, if the lower class is closed due to RM controls and if the higher fare is still lower than the passenger's WTP. The simulated sell-up behaviour is thus imbedded in the PODS passenger decision model, rather than specified as an input sell-up parameter.

Figure 6 shows the geometry of network D, one of the standard PODS networks. In network D there are two airlines, each with their own hubs. The traffic flow is from the 20 western cities to the hubs and to the 20 eastern cities. Both airlines employ three connecting banks per day, resulting in each airline operating 126 legs, which produce 1446 paths in 482 markets.

Six fare classes are assumed, and simulation results are provided here for both an undifferentiated and a semi-differentiated fare structure, as shown in Figure 7. In the undifferentiated case there are no fare restrictions or advance purchase rules. In the semi-differentiated case, two restrictions (no refunds in case

of cancellation and a fee for changing reservations) are applied to lower fare classes (Figure 7).

The focus of the simulation study is on Airline 1 and its use of hybrid forecasting and DAVN-MR. In the simulations, the effects of each of the two components of the transformation were studied separately. Three scenarios were considered:

- Standard: Airline 1 uses standard path/fare class forecasting with DAVN (baseline).
- Hybrid: Airline 1 changes to Q/hybrid forecasting, *without* fare adjustment.
- DAVN-MR: Airline 1 adds fare adjustment to its Q/hybrid forecasting.

Note that hybrid forecasting for the fully undifferentiated fare structure corresponds to only using the price-oriented forecasting model (so-called Q-forecasting) [Belobaba and Hopperstad (2004)].

The three scenarios are examined, for the undifferentiated and the semi-differentiated

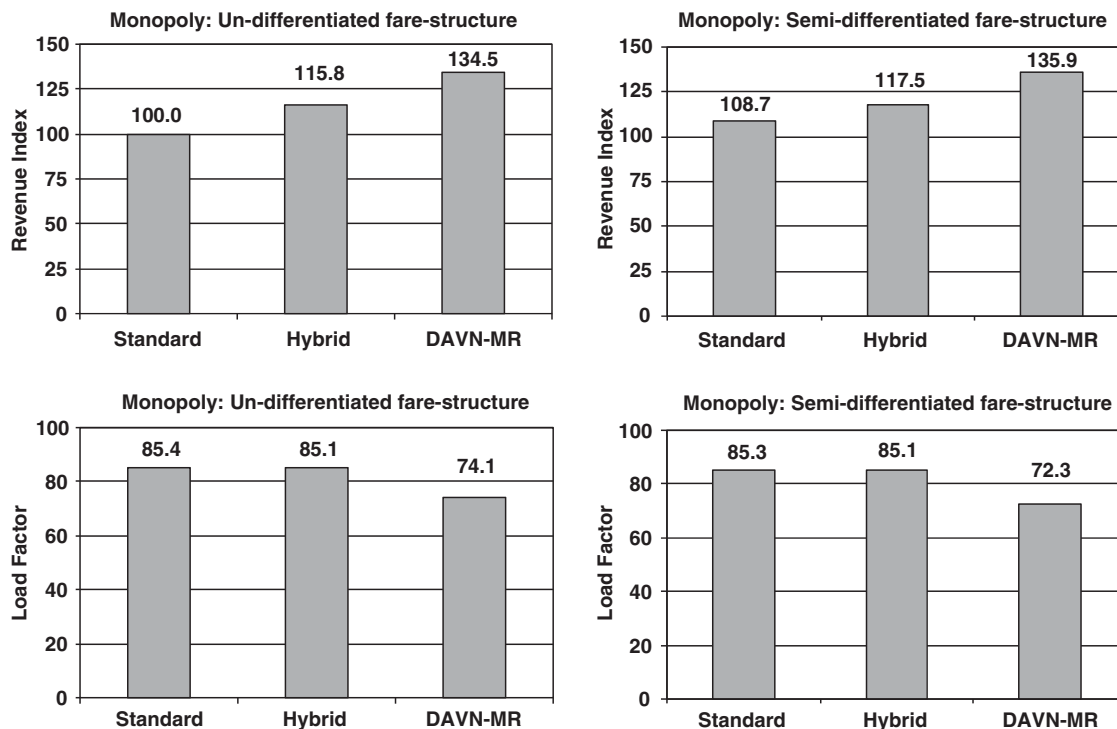


Figure 8: Simulation results for the monopoly case with undifferentiated and semi-differentiated fare-structures in network D.

fare structures, first in a monopoly environment where only Airline 1 is present and then in a competitive environment where the second airline uses standard path/fare class forecasting and traditional DAVN network RM optimization.

Monopoly case

The simulation results are presented in Figure 8. In the panels the revenue index and load factor (LF) are shown both for the undifferentiated and the semi-differentiated fare structures. The revenue index is reported relative to the baseline monopoly case in which the airline employs standard forecasting and optimization.

The results show that applying hybrid forecasting only, for both fare structures, leads to significant revenue gains on the order of 10 per cent–16 per cent. The reason is that hybrid forecasting considers the buy-down

potential of the price oriented demand and reduces spiral down. Note that load factors for standard and hybrid forecasting are almost identical. Hybrid forecasting can be thought of as an attempt to ‘re-map’ passengers that bought down back to their original fare class.

Applying fare adjustment by moving to full DAVN-MR leads to a further revenue increase of about 16 per cent in both cases. The revenue benefit of applying fare adjustment is significant since its effect is to close lower fare classes, which further restricts buy-down. This effect can clearly be seen from the load factors. Applying fare-adjustment leads to a significant drop in load factor from 85 per cent to 72 per cent–74 per cent.

Competitive case

The simulation results for the two-airline competitive case are presented in Figure 9,

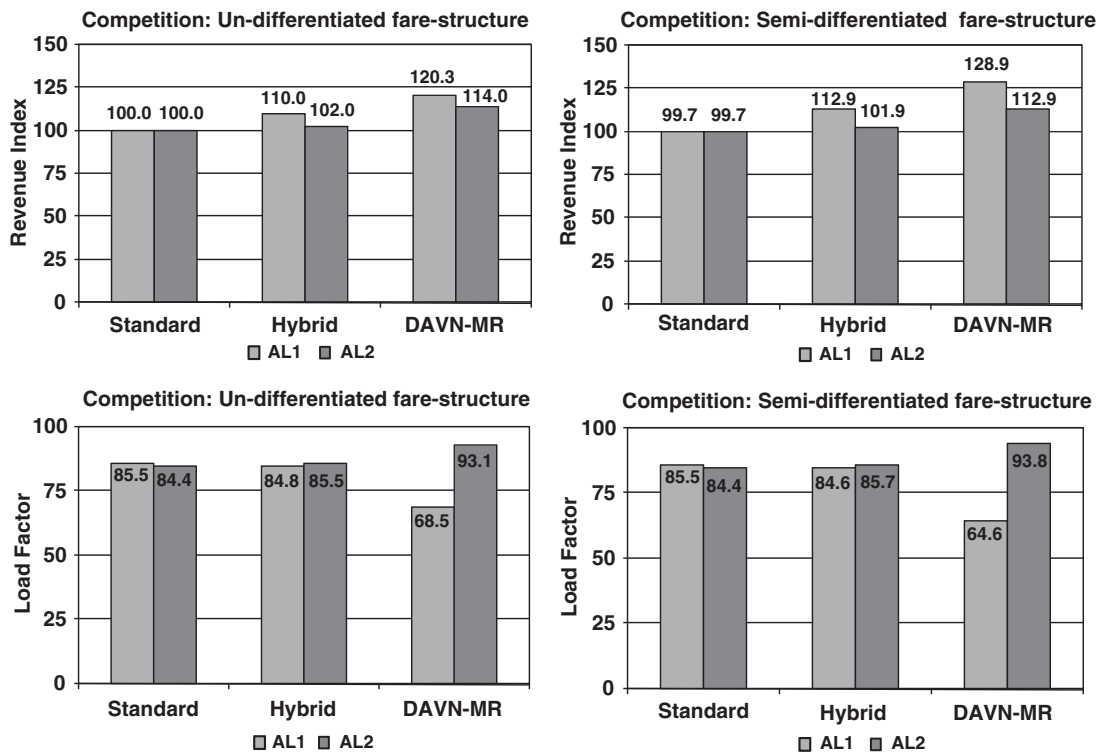


Figure 9: Simulation results in the competitive case with undifferentiated and semi-differentiated fare-structures in network D.



with the panels organized as in Figure 8. The revenue index is reported relative to a baseline scenario in which both airlines employ standard forecasting and DAVN network RM controls.

In a competitive simulation environment, the results show that applying hybrid forecasting again leads to significant revenue gains on the order of 10 per cent–13 per cent. In contrast to the monopoly case, the revenue index for standard forecasting is nearly the same for the unrestricted and semi-restricted cases. The reason for this is that in a competitive environment spiral-down is magnified by competitive feedback, as passengers have a greater set of path/class alternatives from which to choose. With more alternatives, the lowest class 6 is more frequently available on one of the airlines. The consequence of this availability is that passengers who might have bought-up to class 1 or 2 (from, say, classes 3 or 4) do not, because the fare difference between these classes and class 6 is mostly greater than the disutility the passengers attribute to the restrictions. Quite simply, the sell-up potential in a competitive market is lower than in the monopoly case, even with the same passenger choice characteristics.

Applying fare adjustment and thus full DAVN-MR gives a further 9 per cent–14 per cent revenue increase. The use of DAVN-MR by Airline 1 causes load factors to drop from 85 per cent to 65–68 per cent. The competitor picks up the spill-in of low fare passengers rejected by Airline 1 and obtains load factors of 93 per cent–94 per cent. Note that applying DAVN-MR is not a zero-sum game. When Airline 1 moves to DAVN-MR, Airline 2 also benefits from Airline 1's change in RM control. Airline 1's gain comes from extracting higher yields from passengers willing to buy up to higher fare classes. The result is a drop in Airline 1's load factor (from 86 per cent to 65 per cent) with a consequent 10–20 per cent increase in the apparent demand for Airline 2, which gives it a revenue increase due to higher load factors, albeit at much lower yields.

CONCLUSION

We have described an approach to transform the fares and the demand of a general discrete choice model to an equivalent independent demand model. This transformation is of great practical importance to airlines, as it allows the continued use of the optimization algorithms and control mechanisms of traditional revenue management systems, developed more than two decades ago under the assumption of independent demands for fare classes. The transformation and resulting fare adjustment approach is valid for both static and dynamic optimization and extends unchanged to the case of network problems (provided the network problem is separable into independent path choice probability). Assuming the efficient frontier is nested (or approximately nested), the implementation of the control mechanism can be implemented very elegantly by associating to each class the corresponding adjusted fare obtained from the marginal revenue transformation and then applying the control mechanism in the standard way.

NOTE

- 1 While this cannot be possible in our simple stationary one-leg world, it could happen for the network problem, or when the fare structure would depend on the time step.

REFERENCES

- Barnhart, C., Belobaba, P.P. and Odoni, A.R. (2003) Applications of operations research in the air transport industry. *Transportation Science* 37(4): 368–391.
- Belobaba, P.P. (1987) Air travel demand and airline seat inventory management. Unpublished Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Belobaba, P.P. (1989) Application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37: 183–197.
- Belobaba, P.P. (1992) *Optimal vs. Heuristic Methods for Nested Seat Allocation*, Proceedings of the AGIFORS Reservations and Yield Management Study Group, Brussels.
- Belobaba, P.P. (1998) The evolution of airline yield management: Fare class to origin–destination seat inventory control. In: D. Jenkins (ed.) *The Handbook of Airline Marketing*. New York,

- NY: The Aviation Weekly Group of the McGraw-Hill Companies, pp. 285–302.
- Belobaba, P. and Hopperstad, C. (2004) Algorithms for revenue management in unrestricted fare markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology, Cambridge, MA.
- Belobaba, P.P. and Weatherford, L.R. (1996) Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences* 27(2): 343–363.
- Belobaba, P.P. and Wilson, J.L. (1997) Impacts of yield management in competitive airline markets. *Journal of Air Transport Management* 3(1): 3–10.
- Boyd, E.A. and Kallesen, R. (2004) The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management* 3(2): 171–177.
- Bratu, S.J.-C. (1998) Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Brumelle, S.L. and McGill, J.I. (1993) Airline seat allocation with multiple nested fare classes. *Operations Research* 41: 127–137.
- Carrier, E. J. (2003) Modeling airline passenger choice: Passenger preference for schedule in the passenger origin-destination simulator. Unpublished Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Cooper, W.L., Homem-de-Mello, T. and Kleywegt, A.J. (2006) Models of the spiral-down effect in revenue management. *Operations Research* 54(5): 968–987.
- Curry, R.E. (1990) Optimal airline seat allocation with fare classes nested by origin and destinations. *Transportation Science* 24(3): 193–204.
- Fiig, T., Isler, K., Hopperstad, C. and Cleaz-Savoyen, R. (2005) DAVN-MR: A Unified Theory of O&D Optimization in a Mixed Network with Restricted and Unrestricted Fare Products. *AGIFORS Reservations and Yield Management Study Group Meeting*, Cape Town, South Africa, May.
- Gallego, G. and van Ryzin, G. (1997) A multi-product, multi-resource pricing problem and its application to network yield management. *Operations Research* 45: 24–41.
- Guo, J.C. (2008) Estimation of sell-up potential in airline revenue management systems. Unpublished Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Hopperstad, C.H. (1997) PODS Modeling Update. *AGIFORS Yield Management Study Group Meeting*, Montreal, Canada.
- Hopperstad, C.H. (2007) Methods for Estimating Sell-Up. *AGIFORS Yield Management Study Group Meeting*, Jeju Island, Korea.
- Isler, K., Imhof, H. and Reifenberg, M. (2005) From Seamless Availability to Seamless Quote. *AGIFORS Reservations and Yield Management Study Group Meeting*, Cape Town, South Africa, May.
- Lautenbacher, C.J. and Stidham, S.J. (1999) The underlying markov decision process in the single-leg airline yield management problem. *Transportation Science* 33: 136–146.
- Lee, A.Y. (1998) Investigation of competitive impacts of origin-destination control using PODS. Unpublished Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Lee, T.C. and Hirsch, M. (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27: 252–265.
- Littlewood, K. (1972) *Forecasting and Control of Passenger Bookings*, 12th AGIFORS Annual Symposium Proceedings, Nathanya, Israel, pp. 95–117.
- McGill, J.I. and van Ryzin, G.J. (1999) Revenue management: Research overviews and prospects. *Transportation Science* 33(2): 233–256.
- Saranathan, K., Peters, K. and Towns, M. (1999) Revenue Management at United Airlines. *AGIFORS Reservations and Yield Management Study Group*, April 28, London.
- Simpson, R.W. (1989) Using network flow techniques to find shadow prices for market and seat inventory control, memorandum M89-1, MIT Flight Transportation Laboratory, Cambridge, MA: Massachusetts Institute of Technology.
- Smith, B.C., Leimkuhler, J.F. and Darrow, R.M. (1992) Yield management at american airlines. *Interfaces* 22(1): 8–31.
- Smith, B.C. and Penn, C.W. (1988) Analysis of alternative origin-destination control strategies. *AGIFORS Symposium Proceedings* 28: 123–144.
- Talluri, K.T. and van Ryzin, G. (1998) An analysis of bid-price controls for network revenue management. *Management Science* 44: 1577–1593.
- Talluri, K.T. and van Ryzin, G. (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.
- Tretheway, M.W. (2004) Distortions of airline revenues: Why the network airline business model is broken. *Journal of Air Transport Management* 10(1): 3–14.
- Vinod, B. (1995) Origin and destination yield management. In: D. Jenkins (ed.) *The Handbook of Airline Economics*. New York, NY: The Aviation Weekly Group of the McGraw-Hill Companies, pp. 459–468.
- Williamson, E.L. (1992) Airline network seat inventory control: methodologies and revenue impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Wollmer, R.D. (1992) An airline seat management model for a single leg route when lower fare classes book first. *Operations Research* 40(1): 26–37.