PRACTICE PAPERS

The science of revenue management when passengers purchase the lowest available fare

E. Andrew Boyd and Royce Kallesen
Received: 6th April, 2004

PROS Revenue Management, 3100 Main Street, Suite 900, Houston, TX 77002, USA
Tel: +1 713 335 5151; E-mail: aboyd@prosrm.com

E. Andrew Boyd *is Chief Scientist and Senior Vice President of Science and Research at PROS Revenue Management, www.prosrm.com. He received his PhD in operations research from MIT in 1987.*

Royce Kallesen *is the Senior Scientist in charge of airline product science at PROS Revenue Management, and can be reached at rkallesen@prosrm.com.*

Abstract
KEYWORDS: revenue management, yield management, low-fare carriers, low-cost carriers, yieldable demand, priceable demand, airline demand models

The changing business environment is making ticket price an ever more important consideration in passenger purchasing behaviour. This paper discusses how traditional revenue management models change as the ability to segment passengers deteriorates and passengers focus on price. The concepts are applicable to most industries where revenue management is practised.

THE CHANGING BUSINESS ENVIRONMENT

Historical methods of segregating customers are slowly eroding, necessitating changes in the underlying science of revenue management. While it is not the purpose of this paper to sort through the many reasons why this erosion is occurring, it is worth mentioning a few of the major reasons in order to place the science discussion in perspective.

For years, airlines could count on travel agents to help keep customers in their proper segments — business travellers on non-refundable tickets, leisure travellers staying Saturday nights — with incentives and sales agreements to keep the system working. In the late 1990s, a number of factors came together which continue to break down this system. Rising ticket prices in a booming global economy helped entice new airlines into the market, most of which touted a simplified fare structure with generally lower ticket prices. At the same time, the rise of the internet provided a new distribution channel which is proving to be a mixed blessing. Overall

Journal of Revenue and Pricing Management, Vol. 3, No. 2, 2004, pp. 171–177

1477–657X

distribution costs have decreased because there are fewer people involved; but these are the same people who helped keep the old way of doing business alive. The internet has torn down natural barriers to the availability of competitive ticket price information. In well under a minute, a passenger can access the fares of dozens of airlines with route and time of departure for any city pair.

Many traditional carriers have responded by creating new fare structures in which some or all classes have no significant differentiating features other than price. Often, this change is in response to the actions of low-fare competitors in certain markets. Among those carriers and markets where price-differentiated fare classes have not been created, passengers are more frequently tending to ignore ticketing restrictions and purchase on price. In both cases, passenger buying behaviour is very different from the way it is presently modelled in most of the revenue management literature.

MODELLING PROBLEM

Modelling passenger demand is a challenging task (Boyd and Bilegan, 2003; McGill and van Ryzin, 1999; Talluri and van Ryzin, 2004). A model is presented here which captures an important characteristic missing from traditional models, while satisfying essential practical considerations, namely, that the model is computationally tractable, comprehensible to users and generates better forecasts.

The basic modelling issue is straightforward to understand. For descriptive purposes, consider an airline with only one flight leg and two fare classes, Y and Q, where Y has the higher fare. The traditional demand model used in the vast majority of published research leads to the results shown in Table 1 when Y and Q class passengers arrive and try to book.

Using a demand model that assumes passengers will buy the lowest available fare — either because fare class restrictions are minimal or because price is more important than restrictions — leads to the results shown in Table 2.

The key difference between the two models is the behaviour of Y passengers when Q class is open. Under a *yieldable* model of demand, the Y passenger is specifically interested in the Y class product and will purchase that product even when a less expensive Q class product is available. Under a *priceable* model of demand, the Y class passenger is primarily concerned with price and will purchase a Q class ticket because it is the lowest fare available. Note that, under the priceable demand model, there is a shift in what it means to be a Y class passenger. Rather than wanting to purchase a Y class product, a Y class passenger is someone who is willing to pay the Y class fare but will take a lower fare if it is available.

This example brings to the surface a very important practical question: are air-

Table 1: Booking behaviour with a yieldable demand model

Passenger arrival	*Class availability*	*Resultant booking*
Y	Y Open, Q Open	Y
Q	Y Open, Q Open	Q
Y	Y Open, Q Closed	Y
Q	Y Open, Q Closed	None

Table 2: Booking behaviour with a priceable demand model

Passenger arrival	*Class availability*	*Resultant booking*
Y	Y Open, Q Open	Q
Q	Y Open, Q Open	Q
Y	Y Open, Q Closed	Y
Q	Y Open, Q Closed	None

line fare classes *different products* (yieldable) or *different prices for the same product* (priceable)? When fare class restrictions are negligible and the distribution environment presents fares transparently, the priceable demand model is appropriate. When fare class restrictions are significant or the distribution environment promotes segmentation by allowing passengers to see only those fares an airline wants them to see, the yieldable demand model is appropriate.

The reality is that demand lies somewhere between the yieldable and priceable extremes, and may vary by market, distribution environment and a host of other factors. There is, however, an overall shift toward priceable demand.

TWO TYPES OF DEMAND

Yieldable and priceable demand behave differently and give rise to the need for alternative forecasting and optimisation methods.

Understanding forecasting differences

Consider once again an airline with only one flight leg and two fare classes, Y and Q, where Y has the more expensive fare. Suppose the carrier leaves both classes open and that, during the booking period, ten passengers book into Q class.

If the carrier is assuming a yieldable demand model, a reasonable estimate of the demand on future flights would be ten Q class passengers and no Y class passengers. If, however, passengers are actually exhibiting priceable demand behaviour, in the worst case the ten observed Q class passengers could be made up of ten Y class passengers who purchased Q class tickets because they were available.

In general, when forecasting models for yieldable demand are applied when demand is actually priceable, forecasts will overestimate low-fare demand at the expense of high-fare demand and lead to *revenue dilution*. If the process continues from one forecasting cycle to the next, a carrier will experience a spiralling down of revenues as more and more tickets are made available in the lower fare classes.[1]

Forecasting models for priceable demand seek to account for the buy-down behaviour of passengers, recognising that an observed Q class passenger may have been willing to pay a Y class fare. To see how this might be done, consider the information for two different departures of the same flight, as shown in Table 3. Given this information, a reasonable estimate of Y class passengers is four and of Q class passengers is six ($=10-4$).

When forecasting models for priceable demand are applied when demand is actually yieldable, forecasts will overestimate high-fare demand at the expense of low-fare demand and lead to revenue loss from spoilage as a result of setting prices too high. The situation with respect to forecasting is summarised in Table 4.

Table 3: Y and Q class bookings for two different departures of the same flight

Departure date	*Class availability during booking period*	*Total bookings*
1st May, 2001	Y Open, Q Open	Y: – Q: 10
8th May, 2001	Y Open, Q Closed	Y: 4 Q: –

Table 4: Effect of using the wrong forecasting model

		Forecasting model	
		Yieldable	*Priceable*
Actual demand	Yieldable	Good	Overestimates high fare demand at the expense of low fare demand
	Priceable	Overestimates low fare demand at the expense of high fare demand	Good

Understanding optimisation differences

To understand how the different demand models affect optimisation, consider the four fare classes depicted in Table 5. The nested number of seats authorised for sale (AU levels) for the different fare classes are shown for the cases of yieldable and priceable demand. The yieldable AU levels were determined by EMSRb, and the priceable AU levels from an algorithm developed specifically for priceable demand. Note that the priceable AU levels are lower than those of EMSRb, reflecting the fact that there is a potential dilutionary cost in having too many low-class seats available.

To get a better idea of why the AU levels for priceable demand generate higher revenues, it is useful to step through a specific arrival stream and see what bookings are made. Table 6 shows accepted bookings for the case where actual arrivals are the same as mean demand and arrive in low-to-high order.

In simulations where the different algorithms are invoked at 23 points in the booking cycle, the revenue differences can be quite dramatic. Table 7 shows the average revenues using 100 different demand streams from the example in Table 5.

In developing algorithms for priceable demand, it is invaluable to examine the rich history of published research for yieldable demand. The literature has almost exclusively made one of two assumptions about arrivals: either demand arrives in order of low-to-high value (the order Q, B, M, Y in the example) or demand arrives

Table 5: Example of different AU levels for priceable demand (flight capacity: 80)

Fare class	*Price ($)*	*Demand mean*	*Demand std. dev.*	*EMSRb AU levels*	*Priceable AU levels*
Y	800	10	4	80	80
M	680	20	8	75	71
B	540	30	12	56	51
Q	400	40	16	25	25

Table 6: Bookings and revenues for a specific passenger arrival stream

Fare class	*Price ($)*	*Arrivals*	*EMSRb AU levels*	*Booked*	*Priceable AU levels*	*Booked*	*Revenue difference ($)*
Y	800	10	80	5Y	80	9Y	3200
M	680	20	75	19M	71	20M	680
B	540	30	56	30B, 1M	51	26B	–2700
Q	400	40	25	25	25	25Q	0
			Revenue: $43,660		*Revenue*: $44,840		1180

Table 7: Revenue improvement from priceable algorithm

	EMSRb	*PROS priceable*	*Difference*
Revenue ($)	35,919	40,445	4526

Revenue improvement: 12.6%

according to stochastic processes (normally independent Poisson processes) with time varying arrival rates. These two demand models are addressed at length in Boyd and Bilegan (2003).

Priceable demand arriving in order of low-to-high value can be analysed in a manner similar in spirit to that used by Brumelle and McGill (1993), Curry (1990), Robinson (1995) and Wollmer (1992), but starting with an assumption of priceable as opposed to yieldable demand behaviour. Priceable demand leads to a much more complicated analysis.

When appropriate stochastic processes are assumed to govern arrivals in each fare class, dynamic programming yields optimal AU levels. Yieldable demand leads to one form of recursive equation, and priceable demand to another.

A MIX OF TWO TYPES OF DEMAND

Thus far, yieldable and priceable demand have been discussed as distinct in an effort to clarify the differences in forecasting and optimisation methods necessitated by the two types of demand. In reality, both types of demand exist and must be dealt with.[2]

Attempting to segregate the markets is one approach. In local markets where a low-fare competitor exists, a carrier might introduce fare classes for non-connecting traffic and offer these classes only through the carrier's home website. These fare classes could be treated as priceable, while the existing fare classes are treated as yieldable. This approach makes sense when yieldable and priceable demand can truly be segregated in this way.

In most instances, however, it is safer to assume that yieldable and priceable demand book against the same inventory availability. Some important questions present themselves.

Given that I only observe actual bookings in fare classes, how do I determine which bookings were yieldable and which were priceable? If priceable, what is the actual fare class they would have been willing to pay for? How do I use this information to forecast future yieldable and priceable demand?

Given that I have forecasts of future priceable and yieldable demand but only one set of inventory controls, what are the optimal AU levels to set in order to maximise revenue?

At first glance, separating yieldable from priceable demand may seem daunting. Two observations, however, conceptually illustrate how separation can be achieved.

First, when a fare class is open and there is also another open class with a lower fare, then by definition any demand observed for the class is yieldable. Secondly, when a fare class is the lowest open fare class, then the observed demand is a mix of yieldable demand for the fare class and priceable demand for all the fare classes of the same or higher value. Taken together, these two observations provide the basis for estimating yieldable and priceable demand.

Of course, actual forecasting addresses many more details, but it is quite possible to generate forecasts of yieldable and priceable demand by class, with forecasting methods determining the mix of demand. The forecasting methods are a generalisation of both yieldable and priceable methods, and will automatically determine whether demand is purely yieldable, purely priceable or any mixture.

Using estimates of both yieldable and priceable demand by fare class, extensions of the yieldable and priceable optimisation algorithms can be developed to determine optimal AU levels for mixed demand. The extension of dynamic programming to handle arrivals governed by stochastic processes is especially natural, since the recursive equation consists of a probabilistic mix of the equations for yieldable and priceable demand.

A WORD ABOUT O&D CONTROL

The modelling techniques for yieldable and priceable demand can be applied in an O&D control environment. In fact, consumer behaviour is most appropriately modelled at the O&D level. The fundamental analysis does not change, though the implementation details require some modification.

SUMMARY

Understanding yieldable and priceable demand models is extremely important in an era when traditional methods of market segmentation are eroding. Priceable demand models recognise that passengers are price sensitive and not perfectly segmented, and therefore may purchase a fare product which costs less than they are actually willing to pay. Without accounting for this buy-down behaviour, passenger willingness to pay will be underestimated, and carriers will charge less for their tickets than the market is willing to bear. As a general rule, *when priceable demand is accounted for, the net effect is to keep revenues from spiralling downwards by better limiting the availability of low-fare seats.*

It is important to realise that priceable demand is not a function of how a carrier does business, but of how passengers behave; and the rise of low-fare carriers and the growth of internet bookings are both working to encourage priceable behaviour. While carriers cannot control passenger behaviour, they can recognise changes in behaviour and adapt their response accordingly. Revenue management models that account for priceable behaviour are an important part of this response.

Notes

1. The potential for spiral down has long been recognised and discussed. See Boyd *et al.* (2001), Kleywegt *et al.* (2003) and Cooper *et al.* (2004) for recent treatments of the topic.
2. Of course, even more sophisticated models of demand can be postulated, though a host of practical difficulties must be dealt with in their application.

References

Boyd, E. A. and Bilegan, I. C. (2003) 'Revenue management and e-commerce', *Management Science*, **49**, 1363–1386.

Brumelle, S. L. and McGill, J. I. (1988) 'Airline seat allocation with multiple nested fare classes', paper presented at the Fall ORSA/TIMS Conference, Denver, CO. Also presented at the University of British Columbia, 1987.

Cooper, W. L., Homem de Mello, T. and Kleywegt, A. J. (2004) 'Models of the spiral down effect in revenue management', Working Paper, University of Minnesota, Minneapolis, MN.

Curry, R. E. (1990) 'Optimal airline seat allocation with fare classes nested by origin and destinations', *Transportation Science*, **24**, 193–204.

Kleywegt, A. J., Homem de Mello, T. and Cooper, W. L. (2003) 'Models of the spiral down phenomenon', paper presented at the Meeting of the INFORMS Section on Revenue Management, Columbia University, New York, NY.

McGill, J. I. and van Ryzin, G. J. (1999) 'Revenue management: research overview and prospects', *Transportation Science,* **33**, 233–256.

Robinson, L. W. (1995) 'Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes', *Operations Research*, **43**, 252–263.

Talluri, K. T. and van Ryzin, G. J. (2004) *The Theory and Practice of Revenue Management*, Kluwer Academic, Norwell, MA.

Wollmer, R. D. (1992) 'An Airline Seat Management Model for a Single Leg Route When Lower Fare Classes Book First', *Operations Research*, **40**, 26–37.