

Modelling the Performance of Revenue Management Systems in Different Competitive Environments

by

Maital Dar

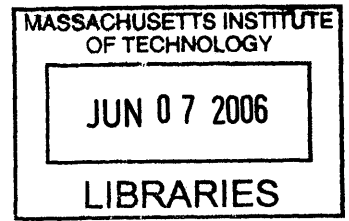
B.E.(Aeronautical Engineering), The University of Sydney, Australia, 2004

B.A.(Archaeology), The University of Sydney, Australia, 2004

Submitted to the Department of Civil and Environmental Engineering
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Transportation
at the
Massachusetts Institute of Technology

June 2006



© 2006 Massachusetts Institute of Technology. All Rights Reserved

Signature of Author :
Department of Civil and Environmental Engineering
May 12, 2006

Certified by :
Peter P. Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Accepted by :
Nigel H. M. Wilson
Professor of Civil and Environmental Engineering
Director, Master of Science in Transportation Program

Accepted by :
Andrew J. Whittle
Chairman, Departmental Committee for Graduate Students

BARKER

Modelling the Performance of Revenue Management Systems in Different Competitive Environments

by
Maital Dar

Submitted to the Department of Civil and Environmental Engineering
on May 12, 2006, in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Transportation

ABSTRACT

In the wake of contemporary widespread fare simplification in many major airline markets, this thesis is concerned with the possibilities and the potential for airline revenue management in less-differentiated fare environments. Traditional revenue management has relied upon the assumption that independent demands exist for different fare class products, and can be forecast as such. However, in less-differentiated fare environments this assumption has been shown to lead to “spiral-down” in revenues.

Hence, in this thesis, seat inventory control methods are simulated in less-differentiated fare environments and their relative performances are compared. The methods tested are: EMSRb-based Fare Class Yield Management (FCYM); Heuristic Bid Price (HBP); Displacement Adjusted Virtual Nesting (DAVN); and Probabilistic Bid Price (ProBP). Each of the methods is tested in conjunction with two different demand forecasting philosophies: the traditional pick-up (or moving average) forecaster which is based on the assumption of independent demands; and a hybrid forecasting method based on the notion that there is one demand for flexible products and another demand for the cheapest product. The methods are simulated in two different competitive airline network environments: a symmetric network with simplified fares; and a more complex non-symmetric network with mixed fare structures.

Simulation shows that the performance of all four revenue management methods suffers in less-differentiated fare environments if they continue to use traditional forecasting. Methods that forecast demand at the path level see inflated forecasts for more expensive products, leading them to reject too much lower-class demand; methods that forecast demand at the leg level see diminished forecasts for the more expensive products, leading them to accept too much lower-class demand.

The efficacy of FCYM improves in less-differentiated fare environments, providing a gain of about 19% over “First Come First Served” revenues (as compared to the 6% gains seen previously), nevertheless, fare product simplification still results in overall network revenue losses of around 16%. Incremental gains from O-D control when using traditional forecasting range from 0.44% to 1.93% over FCYM.

In contrast, when the new hybrid forecaster is used, revenue management performance improves significantly, and all methods provide larger revenue gains in all competitive network environments. Revenues under FCYM are now 1.7–2.6% higher than when traditional forecasting is used. When using hybrid forecasting, the incremental gains from O-D control now range from 2.6% to 4% over FCYM.

Thesis Supervisor: Peter P. Belobaba

Title: Principal Research Scientist, Department of Aeronautics and Astronautics

Acknowledgements

First and foremost, loving and grateful thanks to my mother and to my late grandmother. It has been their lifelong support that has enabled me to reach this point academically, professionally, and personally. I will never be able to repay them, but I hope that I can live my life with the wit, dedication, and integrity that they have exemplified.

Acknowledgements to my research sponsor, the PODS Consortium, which has made this research possible. I thank the various airline representatives as well as our programmer, Craig Hopperstad, for their feedback and guidance about my work.

My two best friends back home, Miriam Sved and Naomi Stead, earn my eternal gratitude for their friendship, their advice on writing, and for encouraging me from afar.

Thanks to my colleagues here at MIT—in the Master of Science in Transportation program, in the International Center for Air Transportation, and in the PODS research group—for their help in getting through the long years and the longer days. Special thanks to Christos Tsonis, Célia Geslin, Gregory Zerbib, and to Michael Reyes for allowing me access to his results for use in my discussions.

Thanks to my fabulous deskmate Richard Cléaz-Savoyen, who showed me the ropes and was a great travelling companion, and who taught me everything that there is to know about PowerPoint and about PODS and about Revenue Management.

And many thanks of course to my advisor, Peter Belobaba, who taught me—and the rest of the world—everything else that there is to know about Revenue Management. I appreciate the opportunities he has given me, his guidance and mentorship, and for being a manager who was so easy to manage.

And finally, to Anne – my heart and my eyes – for everything. You are a good woman, and I love you.

Contents

Contents	7
List of Figures	9
List of Tables	11
1 Introduction	13
1.1 Airline Revenue Management	14
1.1.1 Airline Pricing.....	14
1.1.2 Seat Inventory Control	17
1.2 Recent Developments.....	19
1.2.1 Revenue Crisis	19
1.2.2 How to Segment Demand?	22
1.2.3 Non-ticket Revenues	25
2 Literature Review	31
2.1 Airline Industry Background	31
2.2 Differential Pricing and Passenger Choice	32
2.3 Evolution of Revenue Management and Forecasting Methods.....	33
2.4 Advent of Less-Differentiated Fare Environments	36
2.5 Revenue Management System Simulation—PODS	39
3 Revenue Management Methods and Forecasters	43
3.1 Seat Allocation Optimizers.....	45
3.1.1 EMSRb—Expected Marginal Seat Revenue.....	46
3.1.2 HBP—EMSR Heuristic Bid Price.....	53
3.1.3 DAVN—Displacement Adjusted Virtual Nesting.....	56
3.1.4 ProBP—Probabilistic Bid Price	58
3.2 Forecasting Demand.....	58
3.2.1 Pick-up Forecasting.....	59
3.2.2 Hybrid Forecasting.....	59
3.3 Summary	62
4 Simulation Environment	63
4.1 PODS	63
4.1.1 Willingness-To-Pay Considerations	65
4.2 Network D	68
4.2.1 Characterizing Network D	68

4.2.2	Developing a Simplified Network Environment	74
4.2.3	Re-building Network D—"Network D6"	78
4.2.4	Price Elasticity Multipliers	81
4.2.5	Summary—Network D6	86
4.3	Network R.	88
4.4	Conclusion.	93
5	Results	95
5.1	Network D6—Pick-up Forecasting.	95
5.1.1	DAVN	96
5.1.2	ProBP	99
5.1.3	HBP	102
5.1.4	Summary.	103
5.2	Network R—Pick-up Forecasting	104
5.2.1	DAVN	105
5.2.2	ProBP	106
5.2.3	HBP	107
5.2.4	Summary.	108
5.3	Network D6—Hybrid Forecasting	109
5.3.1	FCYM	110
5.3.2	DAVN	110
5.3.3	ProBP	111
5.3.4	HBP	112
5.3.5	Summary.	112
5.4	Network R—Hybrid Forecasting.	113
5.4.1	FCYM	114
5.4.2	DAVN	114
5.4.3	HBP	116
5.4.4	Summary.	117
5.5	Summary of Results	117
6	Conclusion.	121
7	Bibliography	124

List of Figures

Chapter 1 Introduction	13
Figure 1-1 : How much is that seat in the cabin?	13
Figure 1-2 : Economics of Differential Pricing.	16
Figure 1-3 : Differential Pricing in Action.	18
Figure 1-4 : Sacrificing Consumer Willingness-To-Pay	23
Figure 1-5 : Men in Suits Would Like to Pay for Legroom	27
Chapter 3 Revenue Management Methods and Forecasters	43
Figure 3-1 : Nested Inventory	47
Figure 3-2 : EMSRb Process	48
Figure 3-3 : EMSRb with Fare Compression	52
Chapter 4 Simulation Environment	63
Figure 4-1 : PODS Simulation Flow	64
Figure 4-2 : Willingness-To-Pay.	65
Figure 4-3 : Example PODS Generalized Costs.	67
Figure 4-4 : Network D Layout	68
Figure 4-5 : Connecting Paths	69
Figure 4-6 : Airline 1 Local Markets/Paths.	69
Figure 4-7 : Airline 2 Paths	69
Figure 4-8 : Market Fares by Market Distance—Classic Network D	71
Figure 4-9 : Benefits of DAVN—Classic Network D	72
Figure 4-10 : Fare Class Mix—Classic Network D—ALF=84	73
Figure 4-11 : Generalized Costs—Classic Network D	75
Figure 4-12 : Generalized Costs—Network D6 Flat	75
Figure 4-13 : Fare Class Mix—Network D6 Perturbed—ALF=90.8	77
Figure 4-14 : Market Fares by Market Distance—Network D6 Perturbed.	78
Figure 4-15 : Market Fares by Market Distance—Network D6	78
Figure 4-16 : Fare Class Mix—Network D6 Perturbed—ALF=88.	80
Figure 4-17 : Business emult and Network D6 Fare Ratios.	82
Figure 4-18 : Leisure emult and Network D6 Fare Ratios	83
Figure 4-19 : Fare Class Mix—Network D6 (new mults.)—ALF=84.5	84
Figure 4-20 : Generalized Costs—Network D6	85
Figure 4-21 : Fare Class Mix—Network D6—ALF=84	86
Figure 4-22 : Airline 1—MSP Local Markets/Paths.	89
Figure 4-23 : Airline 2—ORD Local Markets/Paths	90
Figure 4-24 : Airline 3—MCI Local Markets/Paths.	90

Figure 4-25 : Airline 4—DFW Local Markets/Paths	90
Figure 4-26 : Generalized Costs—Network R Non-LCC Markets	92
Figure 4-27 : Generalized Costs—Network R LCC Markets.	92

Chapter 5 Results95

Figure 5-1 : Incremental Benefits of O-D Control—Classic Network D.	96
Figure 5-2 : Incremental Benefits of DAVN—Classic Network D.	97
Figure 5-3 : Incremental Benefits of DAVN—Network D6.	97
Figure 5-4 : Incremental Benefits of ProBP—Classic D & Network D6	100
Figure 5-5 : DAVN and ProBP Performance.	101
Figure 5-6 : Incremental Benefits of HBP—Classic D & Network D6	102
Figure 5-7 : Incremental Benefits of O-D Control—Network D6	103
Figure 5-8 : Incremental Benefit of DAVN—Network R.	105
Figure 5-9 : Incremental Benefit of ProBP—Network R.	106
Figure 5-10 : Incremental Benefit of HBP—Network R.	107
Figure 5-11 : Incremental Benefits of O-D Control—Network R	109
Figure 5-12 : Incremental Benefits of Hybrid Forecasting—Network D6	112
Figure 5-13 : Incremental Benefit of FCYM+HF—Network R	114
Figure 5-14 : Incremental Benefit of DAVN+HF—Network R.	115
Figure 5-15 : Incremental Benefit of HBP+HF—Network R.	116
Figure 5-16 : Incremental Benefits of Hybrid Forecasting—Network R	117
Figure 5-17 : Incremental Benefits of O-D Control.	118
Figure 5-18 : Incremental Benefits of O-D Control—Hybrid Forecasting.	119

List of Tables

Chapter 1 Introduction	13
Table 1-1 : Differentiated Products	15
Table 1-2 : Differentiated SimpliFares Products	22
Chapter 3 Revenue Management Methods and Forecasters	43
Table 3-1 : Summary of Revenue Management System Elements	45
Table 3-2 : Fare Compression and EMSRb Protection Levels	51
Table 3-3 : Example “Greedy” Value Buckets	53
Table 3-4 : Example Displacement Adjusted Value Buckets	57
Chapter 4 Simulation Environment	63
Table 4-1 : Example PODS Fare Structure	66
Table 4-2 : Example Generalized Costs	66
Table 4-3 : Network D Legs and Paths	70
Table 4-4 : Fare Structure—Classic Network D	70
Table 4-5 : Fare Levels—Classic Network D	71
Table 4-6(a) : Performance Metrics—Classic Network D	72
Table 4-6(b) : Revenue Performance—Classic Network D	72
Table 4-7 : Fare Structure—Network D6 Flat	74
Table 4-8 : Fare Levels—Network D6 Flat	75
Table 4-9(a) : Performance Metrics—Network D6 Flat	76
Table 4-9(b) : Revenue Performance—Network D6 Flat	76
Table 4-10 : Revenue Performance—Network D6 Perturbed	77
Table 4-11 : Fare Levels—Network D6	79
Table 4-12(a) : Performance Metrics—Network D6	79
Table 4-12(b) : Revenue Performance—Network D6	80
Table 4-13 : Revenue Performance—Network D6 (new elasticity parameters)	84
Table 4-14 : Fare Structure—New Network D6	85
Table 4-15 : Revenue Performance—New Network D6	86
Table 4-16 : Fare Levels & Structure—New Network D6	87
Table 4-17 : Network R Legs and Paths	88
Table 4-18(a) : Fare Levels—Network R Non-LCC Markets	91
Table 4-18(b) : Fare Levels—Network R LCC Markets	91
Table 4-19(a) : Performance Metrics—Network R	93
Table 4-19(b) : Revenue Performance—Network R	93

Chapter 5 Results	95
Table 5-1 : DAVN Changes wrt Baseline—Network D6	97
Table 5-2 : ProBP Changes wrt Baseline—Network D6	100
Table 5-3 : Fare Structure—Network D6/Network D6 “Restricted”	101
Table 5-4 : HBP Changes wrt Baseline—Network D6	102
Table 5-5 : O-D Control Changes wrt Baseline—Network D6	103
Table 5-6 : Benefits of Revenue Management—Network R	104
Table 5-7 : DAVN Changes wrt Baseline—Network R	105
Table 5-8 : ProBP Changes wrt Baseline—Network R	106
Table 5-9 : HBP Changes wrt Baseline—Network R	107
Table 5-10 : O-D Control Changes wrt Baseline—Network R	108
Table 5-11 : FCYM+HF Changes wrt Baseline—Network D6	110
Table 5-12 : DAVN+HF Changes wrt Baseline—Network D6	111
Table 5-13 : ProBP+HF Changes wrt Baseline—Network D6	111
Table 5-14 : HBP+HF Changes wrt Baseline—Network D6	112
Table 5-15 : Hybrid Forecasting Changes wrt Baseline—Network D6	113
Table 5-16 : FCYM+HF Changes wrt Baseline—Network R	114
Table 5-17 : DAVN+HF Changes wrt Baseline—Network R	115
Table 5-18 : HBP+HF Changes wrt Baseline—Network R	116

1 Introduction

The market for passenger air transportation services is quite unique, and is certainly unlike the markets for most other goods and commodities.

The specific output that airlines produce are seats – or quantities of seat-miles. Each seat takes off at a certain time from a certain airport in the service of a specific—and ephemeral—flight leg, and then travels some amount of miles and lands at a later time at another airport. At landing, the leg is complete, and those seat-miles ceases to exist.

In contrast, the products that airlines actually sell to consumers are trips from origins to destinations. They sell more than this too, of course – they sell an image, an experience, a “journey” – but setting aside these intangibles, the item with a price attached to it is travel from an ‘Origin’ to a ‘Destination’. Even though clearly the trip is comprised of a certain amount of seat-miles, and has various other attributes attached to it such as times of departure and arrival or level of on board service, nevertheless the destination is key: consumers are not primarily interested in a quantity of seat-miles, they are interested in an outcome.

The challenge in airline marketing, therefore, is to maximize revenue while allocating these specific perishable commodities—this individual seat on that distinct and transient flight leg—amongst the many possible journeys which are available for purchase and which could possibly include the seat in question. If the goal was simply to achieve 100% load factors then there would be no problem selling every seat in the sky, given sufficiently low prices. But if we are looking to maximize revenue, then we need a way to decide: How much is each seat “worth”? How much revenue can each seat earn for the airline? Which seats will be matched to which journeys? What combination of seats and journeys can provide the best revenue for the airline?



Figure 1-1: How much is that seat in the cabin?

(Braniff International Airways, 1977, Magazine Advertisement, Sourced online [eBay.com])

In this thesis I will test a series of quantitative methods for addressing this problem of controlling seat inventory. I will compare the relative performance of four seat inventory optimizers in conjunction with two different demand forecasting methodologies, and across different competitive environments.

In the remainder of this chapter, I will describe some recent developments in airline pricing and revenue performance, particularly the growing emergence of fare product simplification—i.e., lower fares and fewer restrictions. In Chapter 2 there will be a review of the literature on revenue management, and in particular on revenue management for use in fare environments that are seeing fare product simplification.

I will then, in Chapter 3, describe in detail the four seat inventory control methods and the two forecasting methodologies that will be tested. In Chapter 4, two network simulation environments will be developed and characterized. These will be designed to represent current airline competitive environments. I will also outline a series of assumptions regarding current passenger choice behaviour.

Then in Chapter 5, using as a basis the passenger choice and revenue management tools provided by the PODS simulator, I will test the efficacy of the series of inventory control methodologies in these different environments. Finally, I will present some conclusions regarding the most effective and appropriate methods to use in less-differentiated fare environments, as well as some directions for future research. But first, I will present some background material on pricing and revenue management in airlines.

1.1 Airline Revenue Management

Airline revenue management is a combination of pricing and seat inventory control. That is to say, a combination of product differentiation and the ability to vary the numbers of seats allocated to each product. Revenue management has been much studied over the past 25 years and it has consistently been found that it benefits airline revenues in many different competitive environments.

1.1.1 Airline Pricing

The modern fare environment is very complex, given contemporary conditions of open

competition (deregulation) and relatively full information (internet booking facilities). Airline pricers try to estimate demand for travel and willingness-to-pay for travel amongst market segments, and then develop special fares to try to fill empty seats. Revenue management analysts try to estimate how many bookings can be expected in each fare class and then they need to decide which bookings to accept.

This routine assumes an operating environment in which differential pricing is extant. Differential pricing is the practice of charging different prices for products with different attributes, in an attempt to segment a market. Essentially, airlines sell the “same” thing – a journey between a particular origin and destination at a given time of the day – but associate different attributes with this thing to create distinct products with different prices.

Table 1-1 : Differentiated Products

Fare	Class	Advance Purchase	Minimum Stay	Change Fee	Comments
\$458	N	21 days	Saturday Night	Yes	Travel on Tue/Wed/Sat
\$707	M	21 days	Saturday Night	Yes	Travel on Tue/Wed
\$760	M	21 days	Saturday Night	Yes	Travel Thu-Mon
\$927	H	14 days	Saturday Night	Yes	Travel on Tue/Wed
\$1001	H	14 days	Saturday Night	Yes	Travel Thu-Mon
\$2083	B	3 days	None	No	–
\$2262	Y	No	None	No	–
\$2783	F	No	None	No	First Class

*(Source: Peter Belobaba
 “PRICING AND REVENUE MANAGEMENT RESEARCH”
 Presentations to MIT Industry Advisory Board Meeting November 4, 2005)*

At one level, airlines differentiate between on board classes by laying out different cabins and associating very different levels of service with the First Class, Business Class and Economy Class cabins. At another level, airlines create differentiated products within these on board classes by associating different conditions and restrictions with different ticket prices. These conditions can be restrictions on possible travel dates, restrictions on allowable length of stay for return tickets, penalties for changes or cancellations, and advance purchase restrictions. An example of this can be seen in Table 1-1. This is a fare structure that was used by American Airlines in October of 2001. The first seven rows are Economy Class products, and the last row is a First Class product.

Why do airlines practice product differentiation? As noted above, airlines would have no problem selling all of the seats that they fly at the same price, if that price were low enough. Obviously, however, this price would lead to very low total revenues. Differen-

tial pricing allows airlines to increase their revenues by more effectively exploiting consumer surplus. In fact, there is an argument to be made that airlines wouldn't even be able to cover their costs if they didn't charge different prices (see e.g. Williamson 1992; Cusano 2003).

In order to break even, airlines need to charge at least a certain minimum critical fare per seat (see Figure 1-2(a)). But then, given fluctuations in demand by time of day or day of week, it is unlikely that flights will have constant loads, so the fare charged would need to be slightly higher than this. However, it is then very possible – and quite likely – that not enough passengers would be willing to pay this break-even fare, which would necessitate raising the price further, with the result that even fewer passengers would be willing to pay it. Hence this is an unstable strategy to pursue. However, if airlines can charge two or more different prices, then the weighted average fare received from all passengers can be higher than the minimum critical fare, and the amount of revenue earned can approach total consumer willingness-to-pay (see Figure 1-2(b)).

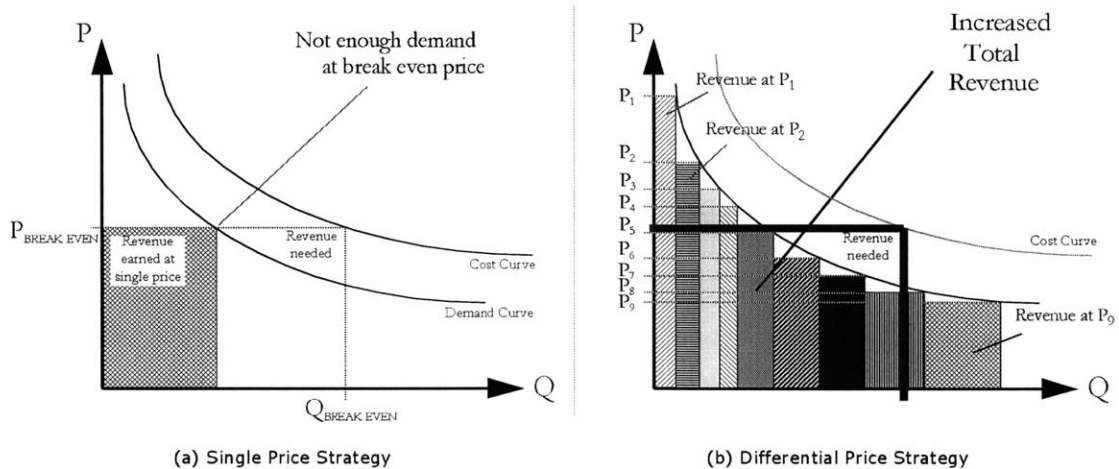


Figure 1-2: Economics of Differential Pricing

This all depends on the airlines' ability to segment demand. In order for differential pricing to work, airline customers must be seen to display heterogeneity in preferences, such that different products are more or less attractive to different segments of the market. Given this, airlines can impose increasing levels of restrictions on their cheaper products in an attempt to prevent diversion – that is, to discourage customers from buying down from the higher priced products to lower priced products. Otherwise, all customers would purchase the cheapest available ticket.

Differential pricing allows airlines to capture greater levels of consumer surplus by making the cheapest products generally unattractive, and then providing higher priced products for consumers with high willingness-to-pay for air travel; and it also stimulates demand by making some very cheap products available to consumers with low willingness-to-pay for air travel. It increases revenues for airlines, and it also increases total consumer welfare. Thus, all participants benefit from differential pricing, and airlines would likely not be able to stay in business without it. Nevertheless, it is a practice much misunderstood by the public, and particularly resented by those customers who do end up paying the higher fares.

1.1.2 Seat Inventory Control

In the hierarchy of airline decision-making processes, pricing and seat inventory control are short-term strategic decisions. Prices, i.e. published fares, can be extremely complex, and they can be changed weekly or even daily. Airlines employ pricing as an agent in numerous strategic and competitive contexts. Nevertheless, prices are transparent and publicly available, and the published fare between an origin and a destination is usually consistent across day of the week, time of day, and path.

On the other hand, seat inventory control is less transparent and seat allocations for a given departure will be different every single day, such that available prices will not seem consistent across day of the week, time of day, or path. Seat inventory control is thus a more dynamic and more tactical level of decision making, and can be used to respond to competitive situations quickly without its implications being immediately obvious to customers or to competitors.

Seat inventory control is the practice of trading off different booking requests for the same inventory of seats in the pursuit of maximizing revenue. There are many ways of carrying this out, ranging from manual *ad hoc* methods to extremely intricate mathematical programs. Particulars of various algorithms will be described in more detail in Chapter 3. At this point, however, a brief discussion of some historical and practical considerations will suffice.

In the early stages of commercial air transportation in the US, all seats were sold at one price, and all on-board service was First Class. From the 1940s onwards, some cheaper seats became available at a discount. As time went by, the size of the discount became larger, and the numbers of “discount” seats sold increased to the extent that this lower

“coach” fare had in essence become the standard fare, and the First Class fare had become a premium product. During the 1970s, as airline competition intensified, this process of product differentiation continued as airlines began to offer discounts on their coach fares, in an effort to compete with charter airlines, sell empty marginal seats, and boost revenues.

Since these seats were not in a different cabin from the full fare paying Economy Class passengers, and since the on-board service on offer was also the same, the issue for airline marketers was in deciding how many of these discounted seats to sell, and how many to reserve for passengers who would pay the full fare. Some intuitive rules of thumb could be used – for example, demand was known to vary by day of the week and by time of day, this was even reflected in the pricing structures. But aside from this, there is still a certain amount of demand stochasticity from day to day, week to week, and departure to departure, and this variability is not so straightforward to predict.



Figure 1-3 : Differential Pricing in Action

(Scandinavian Airlines System, 1964, Magazine Advertisement, Sourced online [eBay.com])

The earliest attempts to predict the amount of demand for the different products on offer were based on observations of bookings, and histories recording which ticket products were booked, purchased, and flown at particular times, and at what prices. The process of deciding upon the limits for the cheaper fare products was largely manual. As the power and magnitude of available computing increased, conditions became amenable to the implementation of more systematic inventory control frameworks. At the late 1980s, human experience and judgement were still the primary techniques used to set and adjust booking limits, until the development of expected marginal seat revenue (EMSR) methods (Belobaba 1987).

EMSR methods are based on two major principles. The first is the recognition of the fact that full priced unrestricted products were no longer the “primary” product on offer to consumers of air travel. Previously, low-fare seats were thought of as by-products – left-over seats that resulted from the variability of demand for what was in fact a full fare product. These surplus seats needed to be sold at a discount to earn some marginal revenue. As such, only demand for the full-fare product was estimated.

By the late 1980s, however, it was clear that the lower priced, more restricted products were also part of the airline’s primary output. Hence it was now necessary to try to estimate demand for each of these multiple fare products and then to allocate inventory in a revenue maximizing way, trading off the relative revenue contributions that could be expected from each product and the likelihood that the seats would go empty (Belobaba 1987: 30).

Consequently, the second principle that is critical to EMSR—and to all methods that are based on it—is the assumption that there actually does exist separate and identifiable demand for these different products. This is not, however, a straightforward supposition. The notion that there is distinct demand for each class, and that this demand is independent of demands for other fare classes, might be accurate if differentiation between products is extremely sharp: that is to say, if the price differences between products are large and if the conditions and restrictions imposed are sufficiently distinct and severe.

Notwithstanding this, if the assumption of distinct demands was ever or has ever been true, recent developments in airline pricing – and their observed effects on airline financial performance – would suggest that it is certainly not the case at present. The problem of effectively segmenting demand, or alternatively of modifying revenue management systems so that they can work with unsegmented demand is an urgent challenge facing airlines today.

1.2 Recent Developments

1.2.1 Revenue Crisis

As described above, airlines segment the demand for air travel by making cheaper products less and less attractive to consumers. They try to achieve this by creating levels of disutility. Demand for travel service is directly derived from demand for the outcome that

the trip will produce, and so travel itself can be thought of as a disutility. Air travel is an inconvenience, and the market can be segmented according to consumers' willingness to trade off between price levels and levels of inconvenience.

Examples of traditional fare restrictions include advance purchase requirements, penalties for changes or cancellation, round-trip travel requirements with minimum stay conditions, restrictions on routings and allowable day-of-week or time-of-day for travel. In the past, restricted fare products would tend to include at least an advance purchase requirement and a change fee of some sort, next a minimum stay requirement, and then possibly other types of restrictions on top of these. Leverage came from the ability to alter the degrees of restriction— 1-, 3-, 7-, or 14- day minimum stay conditions, for example—making the products less flexible with decreasing price.

However it can be argued that in recent years the very nature of the products on offer has changed. The general structure of the differentiation is similar—the most expensive product is still unrestricted, disutilities increase with decreasing price—but the mechanisms of differentiation and the attributes being emphasized are now very different. There has been a trend towards simplified fare products, meaning that round-trip and minimum stay requirements have largely been abolished, and the range of prices charged has been compressed. Why have the pricing structures changed? This can be read as a symptom of the extensive revenue problems facing the industry.

Airline revenues and yields have been decreasing steadily over the past five or six years, resulting from the combined effects of the global slow down in economic growth and the growing supply of air transportation capacity. This increasing over-capacity has been enabled by progressive airline deregulation and consequent increased competition—particularly the proliferation of new airlines offering “low cost” travel at lower fares. The larger older airlines have been undercut by and have lost market share to these newer smaller airlines, which have lower operating costs and can offer cheap and frequent service on popular routes. Their prices on average are lower and their fare products are far less complex—but also far less differentiated—than those of the major airlines.

In addition, market power has been shifting to the side of consumer with the advent of internet distribution, such that consumers almost have “full information” about which products are on offer. Thus, the major airlines have lost their pricing power. Generally, they have matched fares with the low-cost carriers (LCCs) in markets where they are in direct competition. Initially, in an attempt to maintain revenues, they resisted this and responded by offering lower fares for market segments with perceived elastic demand

(such as holiday travellers) and raising the fares for segments with inelastic demand (such as business travellers). But in markets where conditions are appropriate, the LCCs have been able to capture high-willingness-to-pay passengers. For example, in short-haul markets, time-sensitive passengers have demonstrated that they are not prepared to pay premiums for “luxury” service and are more than happy to purchase budget travel, provided that the capacity and frequency are available. In response, the older airlines have reduced their levels of on-board service on short-haul routes.

In response to these developments, a few of the older airlines also began to change their fare structures, making their array of products less complex and more like those offered by the LCCs. In the United States, it was initially Alaska Airlines and America West Airlines who simplified their products. These airlines’ networks were highly exposed to competition from Southwest Airlines – a LCC offering simpler fare products – but they were both relatively small and concentrated to specific limited regions of the country, so the impact of their actions was limited.

But the position of the established airlines worsened, and in January 2005 Delta Air Lines, the third largest airline in the US, implemented a new pricing scheme called SimpliFares¹. SimpliFares was not really very simple, but it promised that there would be only six fare levels in Economy Class, capped at \$499, and only two fare levels in First Class, capped at \$599. There would also be only four levels of Advance Purchase restrictions (3, 7, 14, and 21 days), no ‘Saturday Night Stay’ requirement, and a maximum change fee of \$50 (see delta.com). SimpliFares was implemented throughout Delta’s network and so naturally the other major US airlines matched it.

The result of this was that the established airlines saw increased load factors, and they did regain market share; however revenues fell as their yields continued to decrease through 2005 (Doganis 2005). The reasons for this are twofold: airline revenue management is a combination of pricing and seat inventory control. If prices are lowered and product differentiation is removed or reduced, then there is nothing to prevent diversion, and consumers with high willingness-to-pay have no inducement not to buy the cheaper products. Furthermore, if the pricing structures no longer function to segment demand, then it is very difficult for seat inventory control algorithms to work. How can the demands for different products be traded-off against each other if the different products – and, indeed, the distinct demands – no longer exist?

1. “Delta Cuts, Simplifies Ticket Prices” USAToday, January 5, 2005

1.2.2 How to Segment Demand?

As mentioned above, the fare products on offer today are still differentiated, but the emphases are now different. Many tickets are now sold on a one-way basis, such that the ‘Saturday Night Stay’ (SNS) and most other minimum stay/round-trip conditions are being phased out altogether. The ‘Saturday Night Stay’ (SNS) requirement was extremely effective in preventing business and other time-sensitive travellers from purchasing cheaper products, but was widely reviled by consumers. The various round-trip requirements were very good at segmenting demand, and their removal has greatly weakened the segmentation power of pricing.

Advance purchase restrictions are also being curtailed by some airlines, especially low-cost carriers. Advance purchase restrictions have the effect of reducing product options as departure day approaches, that is, they remove the cheaper products from the choice sets of time-sensitive customers. Of course, they also eliminate purchasing opportunities for late-booking customers with lower willingness-to-pay. Removing or reducing advance purchase restrictions now gives these late-booking lower willingness-to-pay customers the opportunity to book, however it also enlarges the choice set for time-sensitive customers and allows them to buy products that are cheaper than their maximum willingness-to-pay, that is, it allows them to “buy down”. Hence, adjusting advance purchase restriction is a trade-off between these two revenue streams.

Table 1-2 : Differentiated SimpliFares Products

Fare	Class	Advance Purchase	Minimum Stay	Change Fee	Comments
\$124	T	21 days	None	\$50	Non-refundable
\$139	U	14 days	None	\$50	Non-refundable
\$184	L	7 days	None	\$50	Non-refundable
\$209	K	3 days	None	\$50	Non-refundable
\$354	B	3 days	None	\$50	Non-refundable
\$404	Y	No	None	No	Full Fare
\$254	A	No	None	No	First Class
\$499	F	No	None	No	First Class

(Source: Belobaba and Dar 2005)

On the other hand, change and cancellation fees do appear to be increasing, particularly with the growing popularity of on-line instant purchases. However, these do tend to be uniform across an airline’s products. Hence, with reference to Table 1-2, we can see that fare product differentiation is now reflected mainly in different price levels and different

advance purchase restrictions, and the presence or absence of ticket flexibility. In this Delta Air Lines example from April of 2005, we can see that there is one no-advance purchase, fully flexible Economy Class product, and five other Economy Class products whose only differentiable features are price and advance purchase restriction.

We have seen that it is the practice of differential pricing that enables traditional airlines to cover their costs. But with these barely differentiated products, segmentation power has been all but removed from pricing structures. This means that the burden of segmentation now rests with seat inventory control, i.e., the optimizers that decide how many seats to sell at each fare level. So, how is seat inventory control affected by simplification?

The effectiveness of EMSR based methods is reduced by product simplification, since these methods are based on the assumptions of distinct demands. Fare compression has also been found to harm the performance of more sophisticated inventory control methods (Cusano 2003). The difficulty comes from the fact that without inducements to pay for more expensive products, customers with high willingness-to-pay will buy cheaper tickets, and so the potential full fare demand remains concealed. Therefore the expected revenue contributions from the expensive products are under-predicted, and too much preference is given to cheaper products by the seat allocation optimizer. As a result, booking limits on lower classes are too high and there is a great deal of opportunity for buy-down.

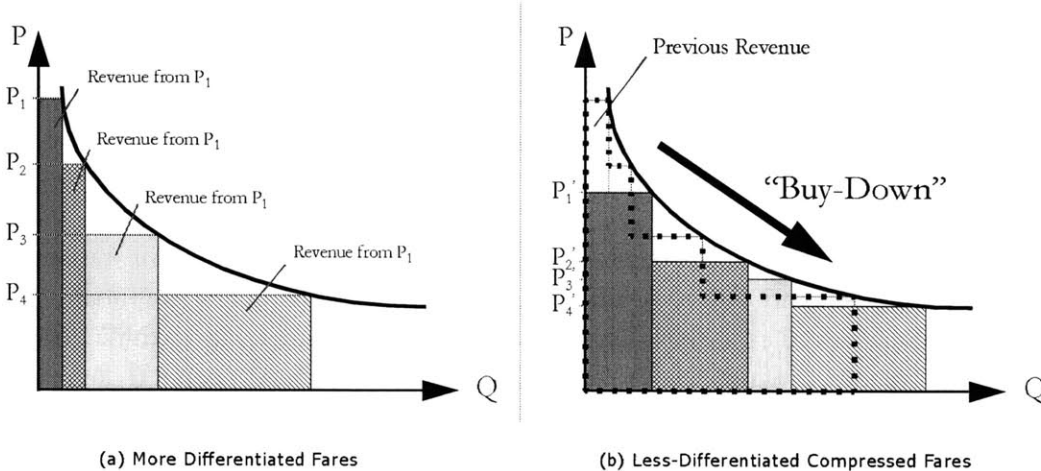


Figure 1-4 : Sacrificing Consumer Willingness-To-Pay

And thus — the crux of the problem. In an environment without differentiated fares, and using tools which depend on differentiation, how can airlines forecast demand for high fare products and segment the market, hence maintaining revenues by exploiting consumer willingness-to-pay. The key question here relates to the problem of forecasting willingness-to-pay – or more precisely, the elasticities of demand for air transportation services.

Traditionally, demand segments have been divided into categories such as “business” and “leisure”, with the assumption that these segments have very different demand elasticities, and cross-elasticities of zero. When EMSR methods were developed starting in 1987, it was considered that demand is a function of price and level-of-service and that these were the strongest agents of segmentation (Belobaba 1987). The term “service” was used to refer to on-board service as well as to the convenience associated with ticket flexibility (e.g., minimum stay requirements, routing restrictions).

However, given that on-board service has declined in recent years, and that differentiation in terms of other convenience levels has almost been removed, it is hard to argue that the products on offer to consumers of air travel are still truly complements with zero cross-elasticities. Indeed, we might really only be able to distinguish two different Economy Class products—unrestricted fully flexible products at full price, and partially restricted products at a range of decreasing prices—but even here the qualitative differences are slight, it’s usually just a matter of price and the presence of some advance purchase restriction.

The domain of seat inventory control is no longer merely concerned with selling marginal seats or protecting seats for predicted high-yield demand. Circumstances have evolved such that its function is to force market segmentation so that airlines can remain operational, even profitable. But the principal inputs required by inventory control systems are demand forecasts. Thus, inventory control must forecast segmentation by estimating willingness-to-pay.

One way to theorize willingness-to-pay is indeed to think of the demand for air transportation as one single demand, and to estimate its price-elasticity and hence potential demand at higher prices (Gorin 2004; Cleaz-Savoyen 2005). Another way to think about it is in terms of urgency. An individual’s willingness-to-pay might be expected to increase with their level of urgency, and this is exploited to a certain extent by the continuing imposition of advance purchase requirements.

Researchers are studying new ways to forecast demand, particularly in less differentiated fare environments. It is a pressing and challenging topic of inquiry. This problem will be returned to again, and the implications will become more apparent in later chapters. But first, a short discussion of some possible consequences of this compressed, less differentiated fare environment for fundamental passenger choice behaviour.

1.2.3 Non-ticket Revenues

A result of the recent changes to airline fare structures has been a great deal of diversion away from expensive products. This could reasonably be expected when differentiation is lessened, and even more so since fare levels have been lowered. We might also expect a certain amount of low-fare stimulation, and this has certainly occurred in many previously un(der)served markets. Classical economic theory suggests that lowering prices will stimulate new demand. However, load factors have been high for a number of years. So it's not that people weren't flying and now they are. It's that everyone is paying less to fly, because of overcapacity and competition. Has this opportunity to pay less for tickets changed consumers' fundamental appraisal of their willingness-to-pay?

The market price leader – i.e., likely the airline with the lowest operating costs – sets the prices. As the lower fares on offer have progressively become cheaper, it has been harder for traditional airlines to maintain very large fare ratios, resulting in fare compression. Variation in fares is less than it was, and fare ratios are as low as 4:1. This is problematic for airline revenues, and for the functioning of airline revenue management systems.

Nevertheless, self-described “business” travellers have long been tired of what they consider to be “irrational” pricing structures—“the traditional business model of business travellers subsidizing cheap leisure travel” according to USA Today.com—and the business and travel sections of major newspapers regularly carry features on pricing matters of interest to this segment. Needless to say, they have been quite happy with the recent evolutions in airline pricing.

As noted above, market power now lies with consumers, many of whom are increasingly highly informed as to what products are on offer, and they are demanding more transparency. They are not choosing their products based on level of “service”, they are mainly interested in price. The differentiation between products in different cabins—First Class, Business Class, Economy Class—is physical and material and intuitive. However, the difference between products within Economy Class is not, and consumers who pay high

prices are offended by this. They are not interested in arguments about consumer welfare or system optimality – they just don't want to feel like they're being gouged.

Reportedly, business travellers are adapting their patterns of choice behaviour. A survey of carried out by Accenture (quoted in the *New York Times* of October 25, 2005) suggests that “business” travellers are becoming more price sensitive, and are learning to behave more like “leisure” travellers and are planning their travel further in advance. In another *New York Times* article¹, airline economist Steven Morrison speculated that traditional carriers have lost their market power over business travel, and declared that “it's a new world” – at least in terms of sustained limits on high fares.

Airlines are clearly also doing PR aimed at the business segment, expressing their remorse for the complex and “irrational” pricing heretofore. In the same article an executive vice president from Delta Air Lines apologized for airline pricing and said that he hopes that customers “feel better” about the new, more “rational” pricing structure. The meaning of the word “rational” in these kinds of contexts is supposed to evoke notions of transparency and correlation with perceived costs, rather than any economic market-based definition of the word, which would refer to the interactions of supply and demand.

Many airlines also try to show consumers that their fares are transparent by laying out the fare structure on their websites—clearly outlining the attributes of the different fare products, and showing their prices and availability. This gives customers the impression that nothing nefarious is taking place, even though seat inventory control mechanisms are still clearly being used to limit seat availability behind the scenes.

Nevertheless, it has been very interesting to watch in recent months as airlines have been trying to re-capture some consumer surplus through non-ticket products. With the compression and conflation of fare products, airlines are moving towards a more “transparent” style of pricing. In other words, if it is too difficult to estimate willingness-to-pay and segment the market accordingly, instead help consumers to reveal their willingness-to-pay explicitly by allowing them to pay for things that they value.

Business travellers don't want to feel as if they are being gouged by large fare differentials that are ostensibly reflecting opaque differences in “service”, but they are still interested in convenience and on-board “service”. This slippage provides an opportunity for airlines to make back some revenue.

1. “*Businesses Are Getting Price Break On Fares*”, *New York Times*, January 14, 2006

"It feels like the airlines are nickel and diming their best customers by offering less and less for the cost of the ticket," says Joseph Gordon, a 27-year-old consultant from Iowa who flies mostly on American. "I'm just resigned to the fact that in economy class you have to pay for almost everything now."

"Booze no longer free -- not even on international flights"
Today In The Sky, <http://blogs.usatoday.com/sky/>

In response to steep falls in revenues and yields, the older airlines drastically scaled back their Economy Class short-haul and domestic on-board service: they cut out meals, newspapers, blankets, pillows; they began to charge fees for the use of non-internet reservation channels. They have focussed on trying to better understand the needs of their different market segments. It had become clear that customers with low willingness-to-pay were happy to fly with no service amenities provided that the price was low enough. At the same time, high willingness-to-pay customers resented having to pay high fares with no obvious better on-board service. Hence, since fares were now lower, the airlines re-introduced meals, but this time for a fee. Many airlines now charge for alcohol, some even charge for water on very short routes. Some are experimenting with selling luxury products, like business lounge access and legroom, on a fee-basis.

The quote above is in reference to United Airlines' charging for alcohol in international Economy Class. The speaker complains, but he still pays for the drink. Indeed, many other travellers are not even offended at being asked to pay for something that was once free, especially given that the price is lower than the airport bar.



Figure 1-5 : Men in Suits Would Like to Pay for Legroom

Western Airlines, n.d., Magazine Advertisement, Sourced online [eBay.com]

Another example where airlines are helping customers to pay for things that they value is the recent trend toward charging for elite Economy privileges which were previously reserved for customers with elite loyalty status. United Airlines is now selling pre-paid packages which offer access to their airport lounges, and first preference for better seats on board:

With the introduction of these Economy Plus AccessSM packages, you can now reserve, at the time of booking, Economy Plus seats for yourself and a guest traveling on the same reservation.

<http://united.com/epaccess>

According to United a spokesperson on USAToday.com, sign ups for the Economy Plus Access program exceeded their expectations. United also allows passengers checking-in for heavily booked flights the option of securing an Economy Plus seat, for a fee of \$35. Similarly, Northwest Airlines allows customers the option of paying \$15 to be allocated to an aisle seat or seat in an emergency exit row.

It is worth noting that not all services are amenable to this kind of development—airline telephone has been available for many years and has never been popular, and wireless internet service is now becoming widely available but has similarly not been popular. These services are different from catering or comfort types of services: they are not particularly critical in-flight, and the fees charged for them have tended to be excessive when compared to rates available on the ground.

It will be interesting to observe this trend, and to see if airline pricing does in fact live in a whole new *a la carte* world, or if the environment will once again resemble its previous structures. This is after all a cyclical process. If we look at historical patterns in airline pricing over a longer time frame, we can see that the current situation might not be entirely unique. Twenty years ago, Belobaba (1987) reported that Business Class fares were only 10% higher than those in Economy Class, such that full Economy ‘Y’ class was being displaced by Business—in other words, the airlines were providing “services” instead of just convenience and flexibility. The problem of distinguishing full ‘Y’ when seating inventory is shared is not a new one, and there are airline executives who believe that traditional airlines can always expect some level of revenue premium.

Another element to take note of is the fact that the most powerful LCCs are also having revenue and yield problems and are now themselves trying to raise fares. It has been reported that Southwest Airlines is adding \$10 to each of its one-way fares, thus allowing

the rest of the industry to do likewise¹. In addition, that JetBlue Airways has been trying to improve its yield and raise revenues by \$5-\$10 per ticket by focussing on less competitive routes². JetBlue's CEO was also quoted in USA Today³ saying that he was going to be spending a lot more time "upstairs" with their revenue management department. In this context, we can see that continued research into appropriate applications of revenue management seat inventory control algorithms and improved demand forecasting methods is still essential.

To this end, in this thesis I will be looking at the effects that the new simplified fare environments have had on revenues and on the effectiveness of traditional revenue management systems. The purpose of the thesis is to investigate some methods that can aid revenue management systems in these environments. The data that this study is based on will come from simulation tests performed using the Passenger Origin-Destination Simulator (PODS). This application will be described in more detail in Chapters 2 and 4. PODS is a computer simulation tool that originated at Boeing, and has subsequently been developed by researchers at MIT in conjunction with Craig Hopperstad of Hopperstad Consulting, and a consortium of major airlines from around the world. It is used to test revenue management methods, by simulating a competitive airline network and by generating passengers who make decisions based on a sophisticated passenger choice model.

Using this simulator I will test a number of revenue management methods (which will be described in Chapter 3) in two different competitive network environments (which will be described in Chapter 4), with the aim of evaluating the potential of these methods for simplified fare environments. The results of these studies will be presented in Chapter 5. But first, in the next chapter, I will outline the literature on airline pricing, airline revenue management, airline revenue management in less-differentiated fare environments, and simulation using PODS.

1. "Southwest Raises Fares, Others Follow", Reuters, March 13, 2006

2. "Jet Blue Says New Routes Will Boost Profitability", Reuters, March 15, 2006

3. "Loss shifts JetBlue's focus to climbing back into black", USA Today, February 21, 2006

2 Literature Review

This thesis is concerned with the question of appropriate revenue management methods for less-differentiated fare environments, and with modelling the performance of these methods in different competitive environments.

In order to investigate revenue management in less-differentiated fare environments we need to consider four underlying areas of inquiry: the airline industry and recent developments in it; the economics of pricing and passenger choice; the evolution of methods of revenue management and forecasting; and simulation techniques.

In this thesis, I will build on the latter three of these areas. I will be considering changes to passenger preferences, as manifested in the willingness-to-pay behaviour of business and leisure travellers. I will be investigating the performance of several leading revenue management optimizers and an emerging method for demand forecasting. Finally, I will be testing these, utilizing some new capabilities in the Passenger Origin-Destination Simulator (PODS) simulation package.

2.1 Airline Industry Background

To understand why it is necessary to reconsider revenue management methods for less-differentiated fare environments, some background in recent trends in the airline industry is helpful. Doganis (2005), for example, provides a wide-ranging review of issues in airline management and the airline business, with an up-to-date introductory chapter on changes in the industry between 2000-2005. Relevant aspects he identifies include the increase in competition with progressive market deregulation and with the emergence of low-cost airlines, and the subsequent global trend of falling yield.

The less differentiated fare environments that form the context of this work have become dominant with the growth of the low-cost airline sector. Characterizing the low-cost airline sector and its business models is beyond the scope of this thesis, and a general introduction to the low-cost sector can be found in Lawton (2002). Focusing primarily on Europe, it includes sections on costs, levels of service, and the impacts of low-cost competition on the industry.

Airline economics and marketing theory provide the theoretical bases of all research into revenue management. A general background to these fields can be found in Holloway (2003) and Shaw (2004). Shaw discusses concepts of marketing, market segmentation and the "customer" for different air travel segments, and includes a basic introduction to airline pricing and revenue management. Holloway contains a more thorough background to airline economics and strategy, with sections on air transportation demand drivers, pricing and costs, capacity management, and revenue issues.

2.2 Differential Pricing and Passenger Choice

In the context of this industry background, the first important principle to understand is the economics of differential pricing. An introduction to the practice of differential pricing in airline markets can be found in Belobaba's PhD dissertation (1987). The relationship between airline differential pricing and yield management is explained further and in lucid detail by Belobaba's article in the *Handbook of Airline Marketing* (1998a), in which he explains how revenue management is necessitated by having differentiated fare products. However, he also notes that the efficacy of traditional revenue management systems depends on the successful market segmentation that these differentiated fare products can achieve. Indeed, Mak (1992) has shown that having more different products on offer, i.e., a greater number of fare classes, leads to higher revenues and greater revenue management effectiveness.

Differential pricing of airline products in combination with seat allocation has been shown to result in greater allocative efficiency than a first-come first-served situation. This was demonstrated by Botimer (1994) who performed a review of contemporary airline fare structures and an economic analysis of consumer welfare. Botimer also elaborates the classic price discrimination formulation by developing a generalized cost model of airline differential pricing. In this model, differently priced products do not simply have different prices: since differently priced products also have different restrictions associated with them, he proposes that they should be modelled with different demand curves.

The next extensions of this work explored the effective or 'perceived' costs of air travel inconvenience and restrictions. Lee (2000), for example, investigated the effects of passenger disutilities in greater detail (i.e., advance purchase requirements, product restrictions, path quality characteristics) and developed a sophisticated model of cost

equivalents for travel disutilities, which was then used for passenger choice simulation. Then, as the airline fare environment changed in the early 2000s, Carrier (2003) re-examined passenger choice modelling in the context of the resulting possible structural changes in passenger behaviour. In this thesis, I will also be modelling structural changes in passenger choice—in particular, the willingness-to-pay behaviour of business and leisure travellers.

2.3 Evolution of Revenue Management and Forecasting Methods

Current revenue management systems have evolved from the databases that airlines used to keep track of their bookings, into automated tools for decision-making which make use of past and current bookings to forecast future demand for seats on flight legs (Williamson 1992).

Reviews of the current practices and future research prospects in revenue management include Ratliff and Vinod (2005), Barnhart et al. (2003), and McGill and Van Ryzin (1999). McGill and Van Ryzin in particular provide a very good review of the key components of revenue management—overbooking, pricing, seat inventory control, and forecasting of demand. In this thesis, I will only be looking at these latter two for use in less-differentiated fare environments, i.e., seat allocation algorithms and forecasting methods.

Since the basis of this thesis is the evaluation of different seat inventory control methods in less-differentiated fare environments, I will first focus on some of the more pertinent work that has been done in this area. The other key aspects of revenue management have been subjected to more specific study elsewhere. For example, a review of the airline overbooking problem can be found in Rothstein (1985). As for pricing, as described in the previous chapter, in recent years we have seen increasing fare compression and fare product simplification (Belobaba and Dar 2005). A review of US DOT ticket price statistics over the past five years has shown that prices in the top one thousand markets in the US have fallen by an average 16% (Geslin 2006).

The root of the seat inventory control problem lies in managing the fixed and shared inventory of seats on a leg, such that a sufficient amount seats are saved at full price for passengers who are willing to pay higher fares, while seats which are otherwise not

expected to be sold can be made available at discounted prices to passengers with lower willingness-to-pay. There are two basic types of approaches for addressing the seat inventory control problem: methods that are fundamentally leg-based, and methods that utilize full network optimization methods. Since all of the revenue management systems that I am focusing on in this thesis have essentially developed out of leg-based Expected Marginal Seat Revenue (EMSR) techniques, the history of these methods is particularly relevant.

The first application of EMSR methods to the seat inventory control problem was at the level of the single-leg two-fare class environment (Littlewood 1972; expanded upon by Bhatia and Parekh 1973; Richter 1982). This kind of approach was subsequently extended to more complex environments, such as two-leg flights (Buhr 1982) and later even to flights of up to four legs with six fare types (Wang 1983). A more generally applicable solution framework based on EMSR methods was developed by Belobaba (1987), who implemented and tested it in a real airline environment.

EMSR techniques involve weighing expected revenues – that is, the product of expected demand multiplied by expected fares - for each fare class and trading them off. This technique has been shown to be optimal only in the simplest of environments (single-leg two-fare class environment). An optimal solution, via convolution integrals, for environments with more than two nested fare classes on a single flight leg was developed by Curry (1990 [as described in Williamson 1992]). However, this method is complicated to implement. When EMSR methods are applied to multiple leg networks with many fare classes, some simplification through heuristics is needed. These methods are easier to implement and intuitive to understand, and are now widespread in industry.

Different heuristics regarding the precise modes of trading off the expected revenues can be incorporated into Belobaba's EMSR framework (Belobaba 1992). For example the methods known as EMSRa and EMSRb are described and investigated in Mak (1992); Ratliff (2005) has also developed a variant. The version commonly in use is called EMSRb. Early simulation testing (Wilson 1995; Belobaba & Wilson 1995) showed that leg-based fare class yield management in a symmetric single market environment, and utilizing EMSRb techniques, provided revenue benefits to the airlines which implemented it and also to the industry as a whole.

The basic EMSR techniques are frameworks for solving the inventory control problem for networks in which all itineraries have single legs. This is known as leg-based inventory control. However, most airlines sell itineraries which consist of multiple leg journeys, and

so inventory is shared not just between different fare classes, but also between “local” (single-leg journey) passengers and “connecting” (multiple-leg journey) passengers. Thus the problem of controlling inventory is complicated further, because the fixed and shared inventory is now being shared along more axes. This step from control by fare class alone to control by origin-destination and fare class concurrently is referred to as Origin-Destination Fare Class (ODF) control.

Although the ODF seat inventory control problem is obviously suited to formulation at the network level, there are a number of difficulties – both practical and theoretical – in solving the problem using network optimization methods. An early example of the use of these network optimization methods can be found in the work of Glover et al. (1982 [as described in Williamson 1992]), who frame the problem as a large network flow problem. Williamson (1992) and McGill and Van Ryzin (1999) provide thorough reviews of these approaches.

However, there are significant problems with applying these methods in real airline environments. Firstly, the size and complexity of their formulations are problematic in terms of data and computation requirements. A second, and more systemic problem relates to nesting. Network optimization methods provide solutions for formulations that assume distinct inventories rather than nested inventories. This means that higher-fare passengers are not allowed access to the entire inventory of seats, but rather are restricted to the seats that the solution specifically allocates to those fare classes. In contrast, in a nested system higher-fare passengers are allowed access to any seats that might be allocated to lower priced fare classes. That is, their access extends even to seats for which lower-fare passengers are expected to have demand.

Hence, research emphasis has been on what might be termed leg-based approaches to Origin-Destination control—that is, leg-based optimizers that do incorporate information about network passenger flows and network revenues. Williamson (1988, 1992) investigates this by first outlining some network linear programming techniques that solve the problem “optimally”, and then by developing more practical and implementable leg-based ODF solutions, which perform comparably and which use leg-based optimizations and hence can incorporate nesting.

The issue here is accounting for displacement costs. This means finding a way to decide which itinerary booking requests to accept on a given leg at the risk of displacing potential passengers on connecting legs of other itineraries. Developing on work presented by Smith and Penn (1988) and Wysong (1988), Williamson accounts for network flow

effects by calculating displacement costs for connecting itineraries, and then adjusting the leg-based revenue associated with each ODF accordingly—hence communicating network effects to the EMSR-based seat allocation optimizer. She suggested a variety of ways to calculate these costs (1988, 1992). These displacement cost techniques were further refined by Tan (1994), and further developments can be seen in Wei (1997) and Lee (1998).

Today there are two major types of Origin-Destination control approaches in common use: bid-price methods and virtual nesting methods. Bid-price methods were developed and are described in Wei (1997) and Belobaba (1998b). The main bid-price method discussed in this thesis – and simulated in PODS – is known as Heuristic Bid Price (HBP). Another version of the bid-price method that is here examined is a network optimization variant called Probabilistic Bid Price (ProBP), which was developed and tested by Bratu (1998). The virtual nesting method that will be investigated is called Displacement Adjusted Virtual Nesting (DAVN), as described in Belobaba (2002), and this will be the third Origin-Destination control approach that I will be testing.

In this thesis I will look at these three approaches to Origin-Destination control (HBP, ProBP, and DAVN), comparing their performances against each other and also with respect to leg-based control (as exemplified by EMSRb). The evolution and refinement of these methods, and previous comparisons of their performances can be seen in studies such as Wei (1997), Lee (1998), Zickus (1998), Bratu (1998), Gorin (2000), and Cusano (2003). In these prior studies, testing and evaluation was performed under symmetric head-to-head competitive environments. I will be extending the evaluation in two directions: firstly I will be considering the performance of these methods under less-differentiated fare environments; secondly I will be comparing their effectiveness in a more competitive and asymmetric network.

2.4 Advent of Less-Differentiated Fare Environments

The successful (re-)emergence and growth of low-cost airlines in recent years has had a significant impact on the world airline industry. We have seen the advent of less-differentiated fare environments, with negative effects on average fares and total airline revenues. Estimates claim that revenues network-wide revenues have fallen on the order of 15-20% and some interesting questions emerge for revenue management research (e.g. Belobaba and Dar 2005; Cakulev et al. 2005; Fiig et al. 2005; Saranathan and Zhao 2005).

Namely – is this revenue decrease evidence of substantively different passenger choice behaviour that is evolving in the new environment? or is it merely a result of the new fare environment's specific conditions? (e.g. see Franke 2004; Mason 2005)

The response of established airlines to Low-Cost Carrier competition has generally been to match their fares (eg Windle and Dresner 1999; Forsyth 2003; Donnelly et al. 2004; cf. Morrell 2005). This on its own could be expected to result in revenue decreases. But moreover, the deleterious relationship between fare simplification, revenue management and revenue degradation has led to a phenomenon termed "spiral-down": that is, as there is less inducement to buy high-priced products, fewer of these products are bought; as fewer of these products are bought, airlines' booking histories contain fewer records of these products being bought; as airlines' booking histories contain fewer records of these products being bought, revenue management system demand forecasters forecast less demand for these products; and finally, as revenue management system demand forecasters forecast less demand for these products, fewer seats are protected for high-priced products and hence there is even less inducement to buy high-priced products. In other words, forecasting inadequacies are at the root of underestimation of passenger willingness-to-pay for air transportation and this leads to a systematic decrease in protection levels for high fare classes and hence sales and revenues (Cooper et al. 2004).

Spiral-down can be restated as a problem with the forecasting methodologies of traditional revenue management systems. Therefore, the next component of revenue management that I will be investigating relates to the problems of forecasting demand. EMSR-based systems depend on the assumption that demand for the different fare products is independent (Belobaba 1987) which is no longer a very sustainable assumption, and is seen as a key factor in spiral-down. Hence, the question is how to adapt revenue management systems so that they don't depend on this assumption. Some ideas about how to do this follow below, but first a brief introduction to traditional forecasting.

Forecasting is about predicting future demand based on previous demand and on assumed passenger sell-up characteristics. Some forecasters types are outlined in Wickham (1995), who described time-series forecasting, forecasting based on regression models, and a hybrid of these termed "pick-up" forecasting. He found that pick-up forecasting – a function of the bookings already on hand and projected bookings to come – provided the best revenue outcomes. Since these forecasters depend on historical data, they need to be aided by a detruncation process for the cases where the historical data becomes constrained for one reason or another. Detruncation methods are described and tested in Skwarek (1997). Finally, these methods were also surveyed by Zickus (1998)

who performed simulation tests in PODS.

The other factor to consider is the potential for passenger sell-up (Bohutinsky 1990; Gorin 2000), which is the probability that a passenger will purchase a higher priced ticket. Bohutinsky (1990) performed some early empirical investigations of sell-up behaviour, in which she found that sell-up could only be observed on specific flights, which supports the practice of predicting future passenger behaviour from observed behaviour. She also found that passengers' willingness to "sell-up" was related to their fare class, implying that price elasticity decreased with increased willingness-to-pay. This is consistent with the idea that people who purchase cheaper products largely make their decisions on the basis of price, whereas people who purchase more expensive products may be privileging other attributes.

So, returning to the problem of spiral-down—why is the assumption of distinct demands so problematic? As outlined above, compressing fare ratios and removing restrictions allows passengers to buy lower-priced products and prevents a history of demand for higher priced products from being established. Therefore, it may be necessary to develop forecasting methods that do not rely so heavily on histories of observed bookings.

In anticipation of increased fare simplification in the industry, Cusano (2003) looked at the performance of traditional EMSRb-based revenue management systems (EMSRb, DAVN, HBP, ProBP) doing some initial experiments with fare reductions and fare product simplification. He noted at the time of writing that fare structures had become significantly different, with less product differentiation. Cusano found that these new fare structures resulted in lower industry revenues, but showed that simple leg-based revenue management was still better than not using revenue management at all, and that ODF control was still better than leg-based revenue management. In my work, I will be looking at revenue management performance in the wake of much more widespread fare simplification and endemic spiral-down.

The next step was to look at the influence of explicit Low-Cost Carrier (LCC) competition. Gorin (2004) describes the changes in the US airline industry in detail. He documents the recent history and influence of low-cost competition, and noting that low-fare market share had increased from 5% to 25% between 1990 and 2003, he investigated the revenue management implications when low-fare airlines enter a small subset of a large hub network. He analysed empirical data from two real-world case studies and also performed simulation testing. He found, unsurprisingly, that incumbents see decreased revenues and traffic, although total industry revenue and traffic does increase. At the

network level, he saw that revenue management allows the incumbent airlines to focus on connecting traffic. He also found that the performance of the revenue management methods was robust across different network environments. I will also be modelling the effects of a LCC competitor, but I will be extending the analysis to a LCC with its own connecting hub.

Having demonstrated the effects of using traditional forecasting in less-restricted environments, the next innovation was to try to alter the forecasting process. Cléaz-Savoyen (2005; also Fiig et al. 2005) used the same network, with a LCC entering a small part of a larger hub network, and tested two new techniques. The first is called fare adjustment, and its aim is to handle competitive market situations where “traditional” differentiated products and less-differentiated products are being sold on the same flight legs. Its scope is to adjust local fares in order to account for potential sell-up. The second is a new forecasting method which no longer assumes independent demands for fare products. Q-Forecasting, developed by Belobaba and Hopperstad (2004) tries to fracture or segment the observed demand—which is unsegmented—and then estimate willingness-to-pay or price elasticity of demand in order to produce demand forecasts that can be used with a traditional EMSRb-based optimizer. Q-forecasting tests were also carried out by Vanhaverbeke (2006) in a dynamic programming framework.

An extension of Q-forecasting that will be tested in this thesis (see also Reyes 2006) is based on the notions of price/product demand, also known as priceable and yieldable demand (Boyd and Kallesen 2004). This is the theory alluded to above that there is on the one hand demand for specific products, and on the other hand also demand for the “cheapest” product, and that forecasts for these demands can be made separately and then combined. These ideas will be explained in more detail in the next chapter.

2.5 Revenue Management System Simulation—PODS

Finally, it is worth noting that it is very hard to determine the effects of particular revenue management techniques or forecasters explicitly in real airline environments. Those environments are far too complex with far too many variables. Competitive environments change very quickly along multiple axes, making it hard to discern the effects of any particular action. This is why simulation studies, in which it is possible to change one variable at a time and hold all others constant, are useful for research in this field. Hence the utility of the Passenger Origin-Destination Simulator (PODS).

Early simulation of revenue management algorithms was done using monte carlo simulation, or spreadsheet applications. For example, Williamson (1992) developed an elaborate Monte Carlo based booking process/seat inventory control optimization simulation: Bratu (1998) also used this model. In contrast, Wickham (1995) performed his forecasting simulations with the help of Matlab.

The Passenger Origin-Destination Simulator is much more sophisticated than these earlier applications, and its architecture is well described in great detail by Wilson (1995) and Cléaz-Savoyen (2005). It contains independent simulations of seat optimization processes and the customer booking process, based on a highly elaborate passenger choice model (described in Carrier 2003). It is subtle and robust, and well suited to performing all-else-equal sensitivity tests of different scenarios.

PODS was developed at Boeing by Hopperstad, Berge and Filipowski, and it is used as a tool by a research consortium consisting of MIT researchers and a group of large airlines. It has been extended and enriched over the years as the functionalities of more and more revenue management techniques and forecasters are added to it.

The first academic testing was by performed in Wilson (1995). Wilson built a single-market environment, with two airlines competing head-to-head in a single non-stop market. He tested basic leg-based EMSR methods in this environment. Subsequently, Skwarek (1997) and Zickus (1998) used PODS to test various permutations of forecasting models.

The testing environment was enlarged by Lee (1998) and Zickus (1998) who created small hub network environments in order to evaluate the performance of leg-based OD methods. These networks had two airlines competing in one of three networks:

- Network 1 had one hub and four spokes - the two airlines shared the hub and hence competed for the same local traffic
- Network 2 had two hubs and four spokes - the two airlines were based in different hubs and only competed for connecting traffic
- Network 3 had two hubs and now six spokes

Gorin (2000) used this last network (Network 3) to look at the performance of estimation techniques - different forecasters, sell-up estimators, and detruncaion methods.

Passenger choice and disutilities were investigated thoroughly by Lee (2000). She used a much larger network—“Network D”—consisting of two airlines of equal size, two hubs,

and forty spoke cities, and incorporated the effective costs of disutilities into the passenger choice model. She tested both leg-based and O-D control methods and found that this more realistic modelling of passenger choice behaviour lead to higher gains from OD control. Carrier (2003) also looked at passenger choice models in Network D and also in another network (“Network E”) which simulated trans-atlantic alliance traffic.

As noted above, in anticipation of fare simplification Cusano (2003) used PODS to model the effects of lower fares and fewer restrictions. He did this in Network D and changed some of the default fares and restriction settings. Building on this, a new-entrant LCC was added to Network D to compete directly with one of the airlines in ten of its non-stop markets. Gorin (2004) used this network and the single market case to model various scenarios of low-fare entrant competition. He found that the main determinants of the resulting observed revenues and loads were the LCC's capacity levels, whether the incumbent matched the LCCs prices and fare structures, and the different revenue management optimizers used.

Cléaz-Savoyen (2005) continued to look at Network D with an LCC in ten non-stop markets. His focus was on helping revenue management methods to deal with spiral down in two ways, using fare adjustment and Q-forecasting. The fare adjustment methodology was designed to deal with mixed network situations such as this, i.e., where there is more than one different fare structure on some of the legs, and Q-forecasting was developed for forecasting in less-restricted environments.

I will build on this work in two ways. Firstly, I adjust the symmetric environment -- Network D - so that all of its markets have simplified fare products. Then I also look at a a more complex competitive environment - Network R - which is a mixed network: there are two “traditional” airlines which are similar geographically and in scope, one “traditional” airline which is dissimilar geographically and smaller in scope; and an enlarged Low-Cost Carrier that has a connecting hub and competes with all of the “traditional” airlines. The characteristics of these networks will be described in Chapter 4.

The PODS simulator has also been utilized in research into such other topics as the competitive and revenue management issues surrounding alliances and code-sharing (Darot 2001); the implications of changes to airline distribution channels (Dorinson 2004); and the traffic and revenue impacts of LCCs offering hub bypass competition in traditional networks (Zerbib 2006).

3 Revenue Management Methods and Forecasters

Traditional leg-based revenue management systems collect historical booking data by flight leg and by fare class and use these to forecast future demand by departure date and fare class. Based on these demand forecasts, they calculate booking limits on low-value fare classes in order to protect high-fare class seats, hence maximizing revenues. The goal is to save enough seats for high-fare customers by limiting bookings and protecting seats, but also to not overprotect by leaving too many seats empty. Forecasting and optimization are intertwined in revenue management, and forecasts and booking limits are revised and re-optimized periodically in response to materializing bookings.

In the first section of this chapter I will provide some details about the mechanics of the four revenue management systems that will be studied in this thesis. The decisions that each system makes will be described; as well, reference will be made to each system's relative benefits and disadvantages, the kind of data that each requires, and some previously reported performance results. Finally some notes regarding the implications of less-differentiated fare environments on their operation and effectiveness will be included. In the second section of this chapter, I will describe the two kinds of forecasting approaches that will be tested with these revenue management systems.

A number of types of revenue management systems have been developed over the years: these different systems are based on different methodologies and can have different mechanisms and decision rules, and they each are constrained by various specific practical considerations. Thus, the system that any one airline chooses to employ depends on its particular needs and the conditions in which it operates.

Perhaps the most important categorization to be made currently is between those inventory control systems that focus on maximizing leg revenues, and those that focus on maximizing network revenues. This distinction has implications for what each system can be expected to achieve, which methodological tools it employs, and the amount, quality, precision and complexity of the data and data management systems that are required to support it. Put simply, leg-based inventory control systems attempt to maximize revenue by optimizing seat allocations on each flight leg, whereas network-based inventory control systems attempt to maximize revenue by optimizing seat allocations over the

entire network. Maximising revenue on each leg is not the same as maximising revenue over the network, and, needless to say, the latter is a far more complex task, however it is also preferred as a goal.

It is easy to appreciate that on any individual flight leg that starts or ends at a network hub, while there will certainly be passengers on single-leg journeys, there will also be passengers who have purchased travel on any one of hundreds of possible Origin-Destination Itineraries (ODIs). On top of this, each of the possible ODIs served by this leg can also be booked in any one of 5 or 10 or 30 different booking classes. Comparing the relative revenue contributions of these itineraries to this specific flight leg is no trivial matter.

In implementing a leg-based system, traffic on each leg is forecast by fare class, and then protection levels and booking limits are optimized based on expected revenue for each fare class. The disadvantages of using a system based on this type of methodology mainly stem from the fact that many of the passengers on any given leg are not travelling on single-leg itineraries. First, there is the problem of deciding how to allocate connecting revenues to the leg in the first place — is the entire fare used? or is it prorated somehow? — and then there is the issue of the wide range of prices that can be assigned the same fare class designation. This makes forecasting loads and revenues on each leg by fare class an imprecise process. Finally, leg-based methods can't "see" connecting traffic, and so it is not possible for them to perform any trade-offs between connecting itineraries and local itineraries.

However, the advantages of this class of methods is that they are computationally relatively simple to implement and to operate. They are intuitive to interpret and to monitor. The data which they use are richly and easily available. By contrast, in implementing a network-based system, network traffic is forecast by ODI and by fare class (ODIF), and then protection levels and booking limits are optimized as before based on expected revenue for each ODIF. This obviates the need to allocate revenues to each leg because the question of the value of each leg becomes irrelevant: the network solution implicitly determines the revenue contribution of each leg to the network. In addition, this type of system can "see" connecting traffic, and can judge whether to accept or reject connecting itineraries.

However, a significant drawback with network-based methods is that they can be very difficult to implement. The mathematical programs involved can be extremely large and can take a long time to run, and while this is not an insurmountable problem these days,

it can still be very expensive. A more serious problem is that these systems require high quality and high precision booking histories as data inputs. There can be hundreds of thousands of possible ODIF's, many of which will have very limited booking histories. Furthermore, demand forecasting is based upon observed booking means and variances. These rarely-purchased ODIF's will tend to have low means and relatively large variances which will result in overestimation of future demand, and thus poor quality forecasts.

Another refinement that can be made to this categorization between leg-based and network-based systems is to further divide them according to whether the forecasting process is performed on a leg basis or on a path (network) basis; and whether the seat allocation optimizer is leg or network based. In the following sections the specific characteristics of some revenue management system methodologies and some forecasting approaches will be described, with consideration given to the issues raised above.

3.1 Seat Allocation Optimizers

The four systems to be considered in this thesis are summarized in Table 3-1 and below:

Table 3-1 : Summary of Revenue Management System Elements

System	Level of Forecasting	Level of Optimization	Level of Control	Mechanism for Control
Leg EMSRb	Leg Forecasts	Leg Optimization	Leg-based Control	Booking Limits
DAVN	Path Forecasts	Network and Leg	Leg-based Control	Booking Limits
HBP	Leg Forecasts	Leg Optimization	Path Based Control	Bid Prices
ProBP	Path Forecasts	Network Optimization	Path Based Control	Bid Prices

- Leg-based EMSRb keeps data and hence makes forecasts at the level of the leg and fare class. It controls inventory using leg/fare class limits, and it finds these limits using a leg/fare class EMSR optimization.
- EMSR Heuristic Bid Price (HBP) keeps data and hence makes forecasts at the level of the leg and value bucket. It controls local (single-leg) inventory using leg/fare class limits, and it controls connecting inventory using a heuristic bid price rule based on these leg fare class limits. It also finds the leg limits using a leg EMSR optimization.
- Displacement Adjusted Virtual Nesting (DAVN) keeps data and makes forecasts at the level of the Origin-Destination itinerary and the fare class (ODIF); optimization also

takes place at the ODIF level, making use of a deterministic or probabilistic network LP. DAVN then controls inventory using leg and value bucket limits which it finds using a leg/bucket EMSR optimization.

- Probabilistic Bid Price (ProBP) keeps data and makes forecasts at the level of the Origin-Destination itinerary and the fare class (ODIF). It controls all inventory using a bid price rule based found via a probabilistic network convergence.

3.1.1 EMSRb—Expected Marginal Seat Revenue

As noted previously, EMSR methods were introduced to airline seat inventory control by Littlewood (1972) and then more widely applied by Belobaba: a version known as EMSRa was introduced in Belobaba (1987), and was further developed into the version known as EMSRb (Belobaba 1992). The differences between these methods lie in heuristics used in their precise formulations, and more details can be found in Mak (1992) and Belobaba (1992). EMSRb is now a very widely used leg-based seat inventory control methodology.

EMSRb is a decision rule for choosing how many seats on each flight leg will be protected for a given fare class against the lower-valued fare class/es below it; or conversely, choosing the limit on the number of seats on each leg that can be sold in a fare class, such that enough seats are saved for the higher-valued fare class/es above it.

EMSRb does not provide a “distinct” inventory solution- i.e., it does not allocate a certain number of seats to each fare class separately: it is instead a 'nested' inventory solution. It defines a combined protection level for a fare class by determining a booking limit on all of the less-valued fare classes below it. In a nested system, there is no discrete limit on the higher fare class seats, for example a full fare request can never be turned down if there is any availability at all. More seats than originally predicted are allowed to be sold in higher fare classes if there is demand for them (see Figure 3-1).

Non-nested algorithm formulations are used when using network optimizations to find “optimal” inventory solutions, an approach which is appropriate for a constrained resource allocation problem. However, as noted previously, it is not practical for seat inventory control because flight leg inventories (i.e., physical seats) can be shared between different marketed “products.” Another factor to note is that nested systems provide better revenue performance than non-nested systems, with non-nested systems overprotecting seats for higher fare classes (Williamson 1988: 116).

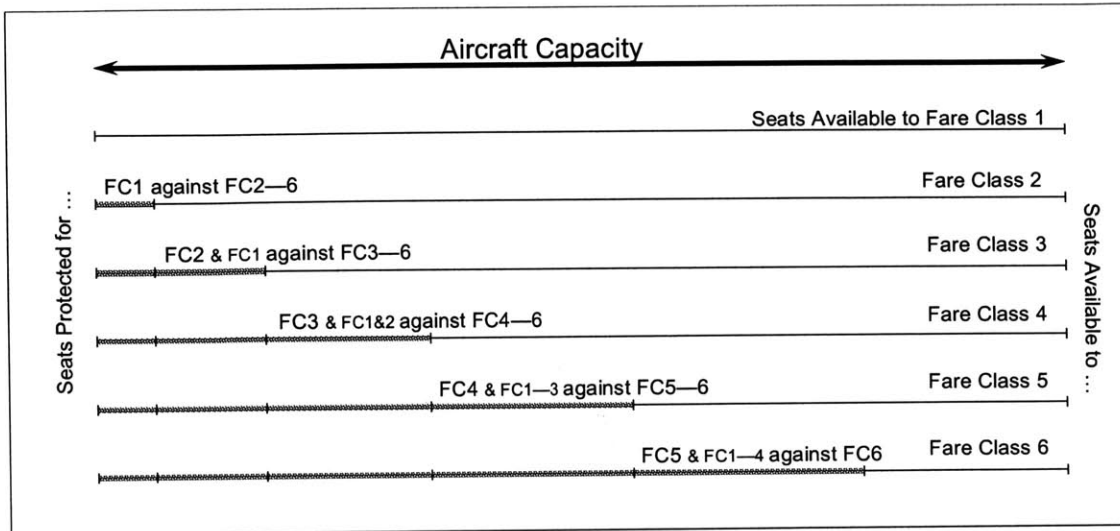
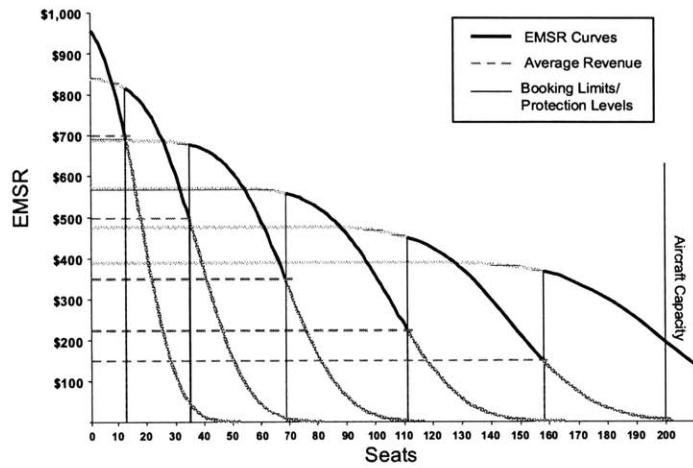


Figure 3-1: Nested Inventory

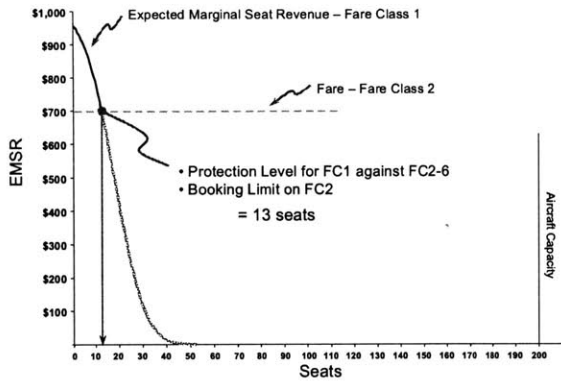
EMSRb booking limits are determined by trading off the expected revenues from different fare classes. In a nested inventory, it is necessary to limit the access that lower fare classes have to seats, and so it is necessary to calculate protection levels for higher classes. Seats are protected for a fare class as long as the seats' expected revenue is greater than the fare in the next lower fare class. The protection level for a class is set at seat n , the seat where the expected revenue (i.e., {average revenue from the fare class} x {probability of selling seat n in that fare class}) is still greater than or equal to the fare in the next lower fare class.

For example, this means that the first n seats will be protected for the highest fare class. It doesn't prevent more than n seats being sold in the highest fare class, it just prevents those n seats from being sold in lower fare classes. The booking limit on the next lower class is set at the point where its fare is just less than the expected revenue of the higher fare class – i.e. remaining capacity minus the protection level n (see Figure 3-2(b)).

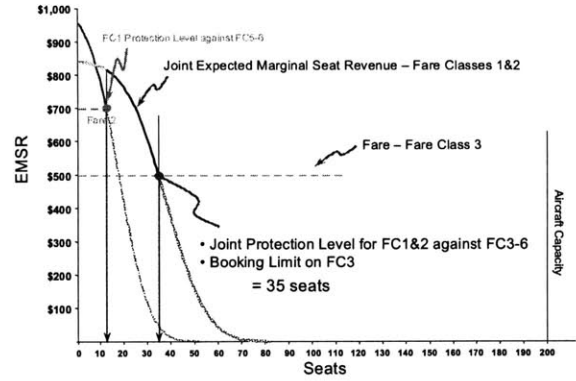
For seat $n+1$, the seat where the expected revenue ({average revenue from the fare class} x {probability of selling seat $n+1$ in that fare class}) is no longer greater than the fare in the next lower fare class, there is no longer any advantage to protecting that seat for the highest fare class.. This is a consequence of nesting: the first n seats are being protected explicitly for the higher fare class, but more seats are still available to it if demand materializes. However since the probability of this is sufficiently low, the next seats are also made available to the next lower fare class.



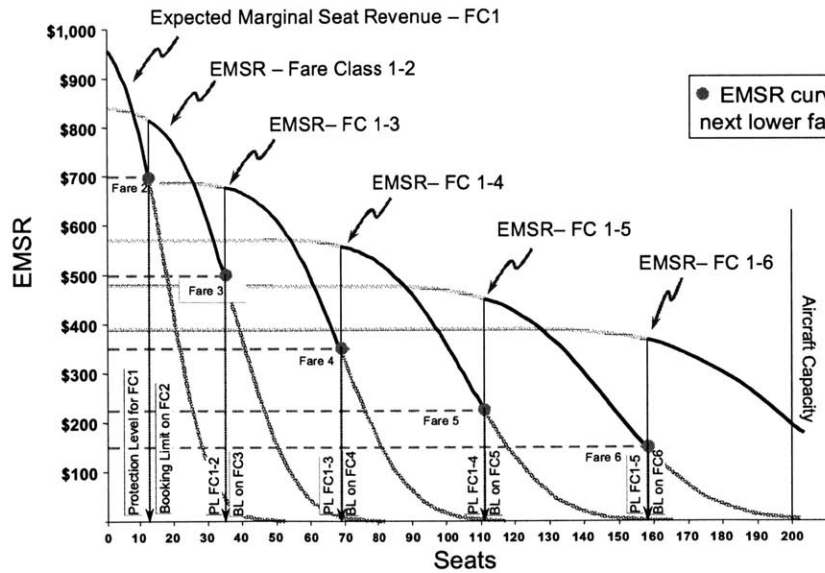
(a) Expected Marginal Revenue Curves & Average Revenues Determine Booking Limits



(b) Finding the Protection Level for Fare Class 1



(c) Finding the Joint Protection Level for Fare Classes 1&2



(d) Example Protection Levels and Booking Limits for a 6 Fare Class Fare Structure

Figure 3-2 : EMSRb Process

Since the inventory is nested, a protection level is not found for the second highest value fare class on its own. Rather, a joint protection level is calculated for the two fare classes together. And so (see Figure 3-2(d)):

- The protection level for the highest fare class determines the booking limit on the second fare class;
- The combined protection level for the two highest fare classes is a booking limit on the third fare class;
- and so on in a similar vein; until:
- The booking limit for the lowest fare class is simply the aircraft capacity minus the joint protection for higher fare classes.

The size of the booking limit on the lowest value fare class is a function of the lowest fare, but is not related to demand for the lowest fare class, i.e., there could be unlimited demand for the lowest fare class. This is the object of EMSRb - to protect the right number of higher value seats from the large amount of demand for the cheapest products. Hence, the leg-based fare class management system will accept a request for a seat in a particular fare class if the booking limit for the fare class has not yet been reached.

From the above it is clear that leg-based fare class control using EMSRb requires information about the probability of selling a seat on a leg in a given fare class. These probabilities come from the booking forecasts, which essentially involves projecting leg/fare-class demand based on historical bookings (the forecasting process will be described in more detail below). When implemented, EMSRb is used at the start of the booking process to establish initial limits based only on historical demand. Then, as the date of the flight approaches, the booking limits are revised periodically taking into account information about the bookings that have actually been received.

Under leg EMSRb fare class control, forecasting and optimization and inventory control are all carried out on each individual leg. Since it is a leg/fare-class optimization, it relies upon leg/fare-class forecasts, and thus a database of bookings and sales by leg and by fare class needs to be maintained. Given most airline reservation systems, this information is easily available, making EMSRb relatively simple to implement. Another practical advantage of EMSRb is that its optimization outputs are individual flight leg fare class booking limits. From an implementation point of view, these are intuitive to analysts and can be easily interpreted (and also adjusted if necessary).

However, a disadvantage of leg EMSRb is the broader problem of forecasting demand

by fare-class. In its traditional formulation, leg-based fare class control using EMSRb relies upon the assumption of independent demand for different fare classes. We have seen that this is not a trivial assumption in general, and in recent years has been closely related to the problem of spiral-down in revenues.

The next practical problem with leg-based fare class control is that the inventory on the leg is not just shared between fare-classes, it is also shared between different itineraries. Hence, calculating expected revenue based on average fares in a designated fare class is imprecise when the fares in a fare class can be very different for different itineraries. Furthermore, this also leads to the problem of network effects, which will be addressed below.

Nevertheless, leg-based fare class revenue management is widely used. Many studies have showed its benefits over not using any automated revenue management at all. Simulations performed in Belobaba and Wilson (1997) show revenue gains of 4-8% in a single market 2-airline competitive scenario. Network simulations in Cusano (2003) show benefits of 6-8% in a 2-airline competitive scenario. Estimates of its efficacy in real airline environments indicate that revenues increase by about 2-5% (Belobaba 1998b). Belobaba and Wilson (1997) also show that implementing fare class revenue management provides a competitive advantage for an airline, helping it to improve its fare mix by protecting a sufficient amount of seats for high-fare passengers, and that the industry as a whole sees increased revenues when other airlines also implement EMSRb, meaning that the extra revenue is coming out of the consumer surplus rather than being captured from competing airlines (1997:9).

Implications for use in simplified fare environments

One of the negative consequences of operating in less-differentiated fare environments for EMSRb is the effect caused by fare compression – i.e., the reduction of the fare ratios of the products on offer. Fare compression can result in the reduction of protection levels for higher-priced fare products.

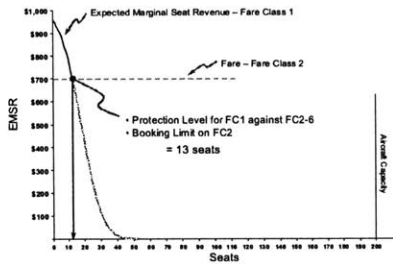
As the fare levels get closer together, the expected revenue of a fare class reaches the level of the fare for the next lower fare class sooner, i.e. the next cheaper product becomes more “valuable” with respect to the higher priced product. As the relative value of the cheaper product increases, the protection level for fare classes above it gets smaller since the optimizer is now less willing to protect for the higher classes, and so – all else being equal – the booking limits for lower fare classes increase.

For example, consider the case described in Table 3-2. Assuming the same mean demand and demand variability for six fare products, the effect of lowering and compressing fares can be seen by comparing the original protection levels and the new protection levels. In the original case, the fare ratio between the highest and lowest fares is 6.67, and in the compressed case the fare ratio is 4. It can be seen that fewer seats are being protected for the higher fare classes as the fare structure is compressed. This is also illustrated in Figure 3-3.

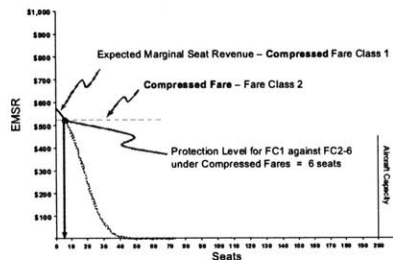
Table 3-2 : Fare Compression and EMSRb Protection Levels

Fare Class	Mean Demand	Standard Deviation	Original Fares	Fare Ratios	Original Protection Levels	Compressed Fares	Fare Ratios	New Protection Levels
FC1	18	10	\$1000	6.67	13	\$600	4	6
FC2	20	9	700	4.67	35	525	3.5	27
FC3	31	8	500	3.33	69	450	3	66
FC4	37	10	350	2.33	111	300	2	105
FC5	41	12	225	1.5	158	225	1.5	153
FC6	53	15	150	1	200	150	1	200

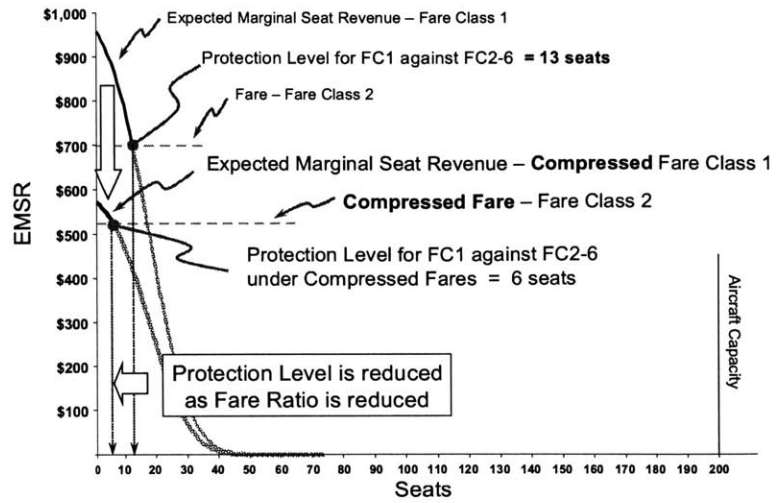
In addition, we have already seen that the assumption of distinct and independent demands for different fare products is a problem in general, and it becomes more so in less-differentiated fare environments. As restrictions are removed and there is less inducement to buy high-priced products, fewer instances of these purchases are seen in booking histories, leading the optimizer to protect fewer seats for high fare classes. This means that EMSRb in its traditional formulation has difficulty operating in less-differentiated fare environments. Some approaches addressing this problem will be described in the forecasting section below.



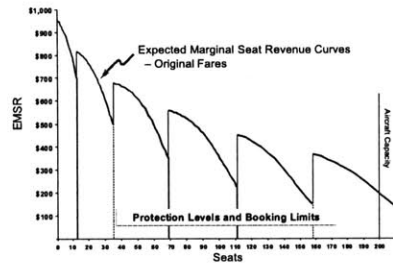
(a) Fare Class 1 Protection Level



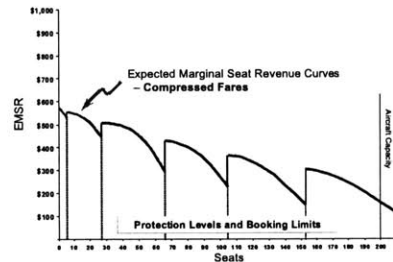
(b) Fare Class 1 Protection Level when Fares are Compressed



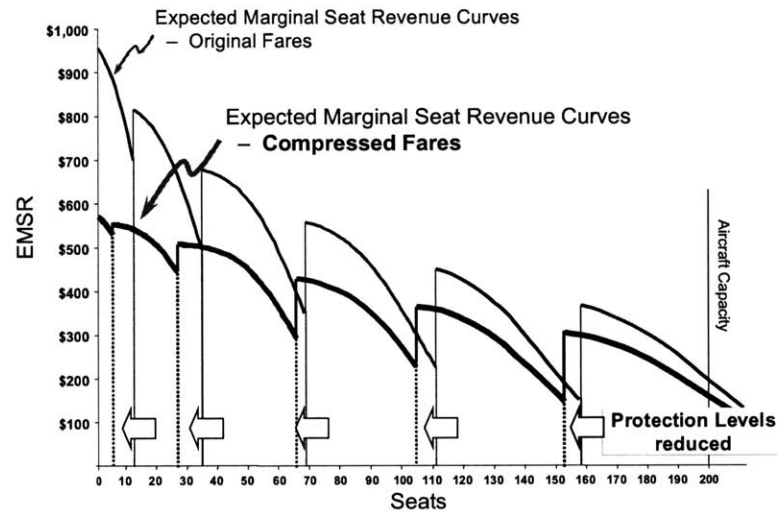
(c) Comparison between Fare Levels, EMSR Curves and Booking Limits when Fares are Compressed



(d) EMSR Curves



(e) EMSR Curves when Fares are Compressed



(f) Comparison between Fare Levels, EMSR Curves and Booking Limits when Fares are Compressed

Figure 3-3: EMSRb with Fare Compression

3.1.2 HBP—EMSR Heuristic Bid Price

The objective of leg-based seat inventory control is to maximize the revenues on each flight leg; the question of maximising revenues system-wide is not addressed. Standard leg-based EMSRb cannot distinguish between the relative values of local and connecting passengers, and this is especially significant in cases where the fares have the same fare class designation but are very differently priced. Sometimes it can be preferable to reject a booking request for a high fare class local fare and to accept a request for a low fare class connecting fare – if the low fare class connecting fare is more expensive than the high priced local fare and will thus provide more revenue to the network. Hence the importance of Origin-Destination control. In this thesis I am investigating three revenue management methods that can be applied for O-D control. They are methods that incorporate information about network passenger flows and network revenues in order to consider the different network revenue contributions of different O-D requests. The parameter that these three methods try to ascertain is a fare’s “network revenue value”.

Bucketing

The first conceptual tool that is necessary if network value is to be taken into account is the practice of bucketing. “Value buckets” are virtual booking classes that are based on the value of the fares (Smith and Penn 1988). Within them they contain fares from the traditional booking classes, but now ranked by each itinerary’s value rather than by its fare class designation (see Table 3-2). This allows better comparison between local and connecting itineraries, where for example a local Fare Class 1 fare might be cheaper than a connecting Fare Class 3 fare.

Table 3-3 : Example “Greedy” Value Buckets

Value Bucket	Fare Range	Equivalent Booking Class (examples)
1	> \$1500	...
2	1250 - 1500	...
3	1000 - 1250	Connecting FC3 fare
4	700 - 1000	...
5	500 - 700	Local FC1 fare
6	320 - 500	...
7	200 - 320	Local FC2 fare
8	< 200	...

Now that the fares are bucketed by value, the inventory is no longer nested according to fare class but rather according to these “value buckets”. As such, the processes involved in seat allocation are carried out neither at the level of fare class nor in terms of official fare class designation. Instead, virtual fare classes defined by these value ranges are used for nesting and for trading off requests for inventory. This is now referred to as “Virtual” nesting.

Nesting fares according to value buckets, and then performing EMSRb on the buckets rather than on fare classes is a rudimentary form of Origin-Destination control. However using this value bucket method always gives preference to the most expensive fares, which will tend to be the long-haul itineraries. This can be undesirable because individual itineraries are only being weighed against one another – i.e., a long-haul connecting itinerary is being compared with a single local itinerary, and hence this method can’t consider the fact that the connecting itinerary might displace passengers on **both** of its constituent local legs. The request for a connecting itinerary will be accepted unless its price is less than the EMSR value of both of the local legs. But the sum of the fares on these two local legs may in fact be higher than the fare for the connecting itinerary. Therefore, in considering which booking requests to accept, it would be more advantageous if a penalty of some kind could be applied to the connecting fare, based on the probability that the passenger purchasing it would be displacing (higher revenue) local passengers.

Displacement Valuation

The value bucket method uses the full amount of the fare to measure a connecting itinerary's value to the network: this is a “greedy” method, in that it gives priority to the highest revenue itineraries and yet can be blind to potentially greater network revenue opportunities. In comparison, bid-price control penalises the connecting itinerary by valuing it with only a proportion of the full fare—if there is a high probability that two local passengers will be displaced. It is thus a less “greedy” way of weighing up a connecting request's value to the network, and results in higher revenues. In bid-price control, the network value of the connecting itinerary is adjusted in relation to “bid prices” for each of its legs. Bid prices are here calculated heuristically using EMSRb to find the expected value of the last currently available seat on the flight-leg: the value off this last seat is called the “leg bid price”. The leg bid price represents the displacement cost of a connecting passenger on the leg in the connecting itinerary.

Under the Heuristic Bid Price method, the decision rule for accepting or rejecting the connecting itinerary involves comparing the itinerary's fare with a “total bid price” for the

total itinerary. This total bid price is a function of the two leg bid prices and of the probability of displacing a local passenger on each of the legs. If the fare is greater than the bid price then the connecting request can be accepted. If the fare is less than the bid price then the connecting request can be rejected in the expectation that carrying two local passengers instead will provide result in higher revenue for the network.

An early version of HBP is described and derived in detail in Belobaba (1998b). This version was tested in a network simulation in Cusano (2003) who found that it provided consistent revenue gains over using EMSRb. In that early formulation, a connecting itinerary request is accepted if:

$$\begin{aligned} \text{Fare} &> \text{ItineraryBidPrice} \\ \text{Fare} &> \max [\text{leg1EMSR}, \text{leg2EMSR}] + \text{displacementFactor} \times \min [\text{leg1EMSR}, \text{leg2EMSR}] \end{aligned}$$

where the displacement factor is defined as probability of displacing a local passenger on both legs. For example if historical booking records show that the proportion of local passengers on each leg is 50%, then the displacement factor will be 0.25.

The formulation of HBP has since been updated. The version used in this thesis is now a function of the expected revenue penalty of displacing a local passenger on leg 1, the expected revenue penalty of displacing a local passenger on leg 2, and the expected revenue if no locals are displaced. In the latest version of HBP, a connecting itinerary request is accepted if:

$$\begin{aligned} \text{Fare} &> \text{Itinerary Bid Price} \\ \text{Fare} &> \text{propLocal1} \times \text{leg1EMSR} + \text{propLocal2} \times \text{leg2EMSR} \\ &+ (1 - \text{propLocal1}) \times (1 - \text{propLocal2}) \times \max [\text{leg1EMSR}, \text{leg2EMSR}] \end{aligned}$$

where propLocal is defined as the observed historical local booking fraction plus an increment to this observed local booking fraction. This increment is an estimate of the unconstrained local proportion of demand.

In the older formulation, the itinerary bid price consisted of the maximum leg bid price plus a proportion of the minimum leg bid price. The new formulation is more elaborate, consisting of proportions of both leg bid prices and the likelihood that local passengers will not be displaced on both legs.

Under HBP, forecasting is performed at the leg level. For local itineraries, optimization and inventory control are carried out at the leg level via EMSRb, and for connecting itineraries at the path level via the total itinerary bid prices. But since all forecasting is still done at the leg level, HBP like EMSRb only requires records of bookings and sales by leg and by fare class. Once again, this makes HBP relatively simple to implement from this perspective. Further, the optimization outputs are individual flight leg fare class booking limits, or bid prices that are based on these leg booking limits, which again makes this method intuitive to use.

Of course, HBP is hindered by the same problem that burdens leg-based EMSRb – the difficulty of forecasting distinct fare-class demands in a less-differentiated fare environment. Nevertheless, HBP is an elegant method that allows a certain amount of O-D control without the need for path based forecasting and multiple itinerary/fare class booking limits, but rather uses a single bid-price value for each request and a simple accept/deny decision rule. It is an improvement over leg-based control in that it can protect seats for high revenue connecting passengers but is not so greedy that it accepts connecting requests when it is expected that local passengers will be displaced on both legs in a connecting itinerary, and it is easier and cheaper to implement than a network optimization based technique.

3.1.3 DAVN—Displacement Adjusted Virtual Nesting

In contrast to HBP, DAVN is an O-D method that does make use of network optimization. DAVN also weighs requests for connecting itineraries by adjusting their fares using a displacement cost, but in this case the displacement cost is not calculated via EMSRb but rather is found from a network optimization.

Under DAVN, a deterministic LP is solved in order to determine which legs in the network are constrained, based on demand forecasts for each path and fare class. The shadow prices on these legs become the displacement cost of accepting a connecting itinerary request which makes use of that leg. When the legs are not constrained, i.e. when the displacement costs are low, the higher priced connecting requests can be accepted. But when one or both of the legs is constrained, i.e. when the displacement costs are high, preference can be given to the local itinerary request.

Once the shadow prices have been calculated, the “leg” fares are nested. In this case they are not nested by fare class designation or by total fare value, but rather they are nested

by buckets according to their displacement adjusted fares (see Table 3-3). Hence this method is called Displacement Adjusted Virtual Nesting.

Table 3-4 : Example Displacement Adjusted Value Buckets

Value Bucket	Fare Range	Equivalent Booking Class (examples)
1	> \$1500	...
2	1250 - 1500	...
3	1000 - 1250	...
4	700 - 1000	...
5	500 - 700	Local FC1 fare
6	320 - 500	Connecting FC3 fare – displacement cost
7	200 - 320	Local FC2 fare
8	< 200	...

$$\text{Displacement Adjusted Leg Fare}_{|local\ itinerary} = \text{Local Fare}$$

$$\text{Displacement Adjusted Leg Fare}_{|leg1\ of\ a\ connecting\ itinerary} = \text{Connecting Fare} - \text{Shadow Price}_{|Leg1}$$

$$\text{Displacement Adjusted Leg Fare}_{|leg2\ of\ a\ connecting\ itinerary} = \text{Connecting Fare} - \text{Shadow Price}_{|Leg2}$$

The protection levels and booking limits are then calculated on a leg basis using EMSRb on these displacement adjusted value buckets, and the decision to accept a booking is similarly based on whether or not the booking limit for that value bucket has yet been reached. The workings of DAVN are outlined in Belobaba (2002), and detailed derivations and simulation studies can be found in Lee (1998).

Under DAVN, seat allocation optimization and inventory control are performed at the leg level via EMSRb for all itineraries. But forecasting is carried out at the path level, with an optimization being performed at the network level in order to find the leg shadow prices. This means that a database of bookings and sales by OD itinerary and by fare class needs to be maintained, which can make DAVN more difficult to implement. A real airline can have many thousands of possible paths on offer, and most of these paths will be rarely traversed and thus their booking histories will show mean demands of close to zero: on top of this, looking at these tiny demands by fare product makes the expected demands even smaller. Using very small demand inputs such as these leads to precision problems with the forecasting and optimization routines. Further, the optimization

outputs from DAVN are LP shadow prices, which are not particularly intuitive to interpret, and so the implementation of DAVN requires more informed and highly skilled and analysts.

These drawbacks are exacerbated in the less-differentiated fare environment. Demand forecasting in DAVN is done at the level of path/fare class, and if spiral down is a problem for leg/fare class forecasting, it can be more pronounced in path forecasting where there are already so few bookings being recorded in high fare classes for many itineraries.

3.1.4 ProBP—Probabilistic Bid Price

The final O-D control method that will be looked at briefly is the Probabilistic Bid Price method. ProBP uses a “probabilistic network convergence” algorithm to perform its network optimization which estimates the marginal value of the last available seat on each leg in the network. This is known as the “critical” EMSR value, and the network value of the connecting itinerary fare is adjusted by prorating it in relation to this critical EMSR value. The decision to accept or to reject the connecting fare is again based on comparing the fare with a bid price, but now the bid price is calculated based on a proration of the path fare, rather than a combination of the leg bid prices as in HBP. This method was formulated and tested in Bratu (1997), and further testing was carried out in Cusano (2003).

ProBP is a complex method and is difficult to implement: like DAVN it requires forecasts by path and fare class. This would be a problem in any competitive environment, given the difficulties of forecasting independent fare class demand for thousands of possible O-D markets. It becomes more of a problem in a less-differentiated fare environment, where the difficulties of forecasting fare class demand in general are increased.

3.2 Forecasting Demand

The utility of revenue management is concentrated around high demand flight legs – if demand is low and capacity is not limited then protecting inventory for higher priced products becomes moot. Thus it is important to be able to estimate demand. In this thesis, two approaches to forecasting will be tested: “pick-up” forecasting, and “hybrid” forecasting.

All forecasting is based on booking histories, i.e., a historical sample which consists of bookings and boardings for previously departed flights and current bookings for upcoming flights. In many cases, however, the observed demand can be less than the true demand. If a flight leg or fare class is sold out, and demand is rejected for want of supply, no records capture this. The “true” demand for the flight leg remains unobserved, and so the historical data often needs to be detruncated before it can be used for forecasting.

Detruncation is the process of deducing actual demand from the observed demand. There are occasions when the true total demand cannot be accommodated on a flight leg and as such the booking limit might be reached for one or more of the fare classes on offer. In these cases, the true demand cannot be observed, and thus needs to be estimated. Detailed analyses of different detruncation methods can be found in Wickham (1995) and Skwarek (1997). In this study, the method known as “booking curve detruncation” will be used. This method assumes that booking patterns over time are similar for all flights. It estimates the true total demand for flights where at least one of the fare classes was booked to capacity by extrapolating patterns observed on flights where that fare class was not closed. Once the booking histories have been unconstrained, they can be used to estimate future demand.

3.2.1 Pick-up Forecasting

Pick-up forecasting is an elaborate form of time series forecasting. It uses averaged values of previous unconstrained bookings, and also uses the change in bookings over time—i.e., the number of passengers that are “picked up” from time period to time period. The demand forecast for a time period is then a function of the actual bookings plus the average demand that is expected to materialize based on the previous pick-up rates. Flights that are used in calculating these averages can be flights from the same day of the week, same time of day, same season, or whatever the appropriate analogous time period is judged to be. For details see Zickus (1998) and Gorin (2000).

3.2.2 Hybrid Forecasting

Hybrid forecasting is based on two concepts: the first is the “Q-forecasting” process developed by Belobaba and Hopperstad (2004), and the second is the notion of “yieldable and priceable demands” (Boyd and Kallesen 2004). The philosophical underpinning of this approach is the notion that there are two types of demand for air transportation: demand for a specific product with specific restrictions (or better perhaps to say demand

for a specific product lacking specific restrictions); and demand for the cheapest available product regardless of restrictions. Hybrid forecasting attempts to forecast these two demands separately and then combine them to derive total forecasts.

Boyd and Kallesen (2004: 172) define two models of demand. The first model, called “yieldable demand”, assumes that customers are interested in a specific product and will purchase that product even if cheaper products are on offer. The second model is referred to as “priceable demand”, and this assumes that customers only purchase the lowest priced product that is available. In PODS and in this thesis, “yieldable demand” is referred to as “product oriented demand”, and “priceable demand” is referred to as “price oriented demand”. Boyd and Kallesen point out that this way of modelling demand offers a solution to the problem of assuming distinct and independent demand for different fare products.

Under the traditional demand assumptions, a “full-fare” customer was thought to be someone who wanted to purchase the full-fare product. Hence, demand forecasts for full-fare products were based on previous full-fare bookings. In this formulation, we are now able to reconceptualize demand such that a “full-fare” customer can be someone who is **willing to pay** for the full-fare product, but will buy a cheaper product if it is available. Therefore, previously observed full-fare bookings are no longer sufficient to estimate full-fare demand. What is required is a model that will estimate the probability that a passenger would buy a higher priced product – the probability of sell-up.

Q-forecasting attempts to solve this problem by first forecasting a single total demand for travel at the lowest price (“Q” class), and then estimating the probabilities that passengers will sell-up to higher fare classes (sell-up probabilities). Together these can be used to forecast the potential demand for the entire set of fare products. Hence – critically – this method does not rely on the assumption of independent product demand. A full description of the definitions, formulas and derivations involved can be found in Belobaba and Hopperstad (2004), and in Cléaz-Savoyen (2005) who also performed extensive simulation tests using this forecasting method.

Briefly, probabilities of sell-up and observed historical bookings are converted into “Q-equivalent” bookings (demand for the lowest priced product). This demand is summed and detruncated where necessary. Pick-up forecasting methods are applied to this total base demand. Finally, the probabilities of sell-up are used to re-partitioned the total demand into demand for the various higher priced products. And so, for example, if total demand for a leg is estimated to be 130 passengers, and if the probability of a passenger

selling-up from the lowest fare-class to the highest fare class is estimated to be 8.5%, then the estimated demand for Fare Class 1 will be 11 passengers.

Naturally, the challenge here is to estimate the sell-up probabilities. There are various methodologies for estimating passenger willingness-to-pay and probabilities of buying higher priced fare products (Cléaz-Savoyen 2005). The methodology used in this thesis is called the Inverse Cumulative method. It estimates sell-up rates by time-frame, assuming that the probability of sell-up increases as the departure of the flight draws closer. It performs a regression to find the sell-up probabilities for each fare class, based on observed “pseudo-bookings” and observed instances of sell-up. The number of “pseudo-bookings” for a class is the sum of bookings observed in all higher classes. More detailed explanations of the Q-forecasting mechanisms can be found in Cléaz-Savoyen (2005: 50ff.).

Now combining these two concepts, we can use regular pick-up forecasting to estimate the “product oriented demand”, and then use Q-forecasting to estimate the “price oriented demand”. This is known as “hybrid” forecasting. In the simulations performed in this thesis, passengers are considered to be “price oriented” if they are observed purchasing the cheapest available option – i.e., the next lower class is closed for some reason, making the next lower priced unavailable. Conversely, if they purchase a product while the next lower class is not closed – if there **was** a cheaper option available – then they are considered to be “product oriented”.

There are two points to clarify: for “product oriented” passengers, we need to define what is meant by “availability”; and for “price oriented” passengers we need to be careful to check whether or not the next lower fare classes were closed due to advance purchase restrictions.

There are three ways to decide whether or not the next cheaper option was “available”:

- Path Rule the next lower class is available on that path
- Airline Rule the next lower class is available on any path provided by that airline
- Market Rule the next lower class is available on any path at all in the market

For “price oriented” passengers, if a booking is observed in the lowest available class, and if the next lower class had been closed due to AP, then that passenger may not necessarily be “price oriented”. Even if the lower class had still been available, they may have purchased the more expensive product regardless. So there are also three ways to think

about these cases. We can designate passengers who buy the cheapest available products as:

- always “price oriented”
- always “product oriented”
- “product oriented” only if the available class and the next lower class are differentiated

And so, with three different availability rules and three different AP classifications, we have a total of nine different options for dividing passengers into “price” and “product” oriented demand. It is also worth noting that the forecasters are aware of the advance purchase rules such that demand is never forecast for classes that have already been closed due to AP.

3.3 Summary

With this understanding of the four revenue management systems and the two forecasting approaches, I am now going to test these methods in simulation. We have seen that fare simplification and fare compression have adverse effects on the functioning of EMSRb—and hence on the revenue management methods that employ it in their optimizations. I am going to investigate the relative efficacies of the four methods in less-differentiated and mixed-fare environments in order to discern which are better suited to cope with simplification. I will also test them in combination with the different forecasters to see if and how the new forecasting methodology can help the traditional methods to work better. Two network structures will be used for these simulation tests, and will be described in the next chapter.

4 Simulation Environment

In this chapter I will describe some specifics of the PODS simulation environment used for the experiments in this thesis. First I will define some important parameters that are used in modelling airline passengers and are relevant to this thesis; then I will describe the network and fare environments that were used, including a description of their developments. Two simulation environments were used in this thesis—a 2-hub network with two very evenly-matched competitors; and a 4-hub network with four less evenly-matched competitors. I will also include here discussion of some adjustments that were made to the passenger choice model's assumptions with respect to willingness-to-pay and sell-up. But first, some general background to PODS.

4.1 PODS

The Passenger Origin-Destination Simulator (PODS) is used to test and evaluate revenue management systems. It can be used to elucidate the interrelationships between different effects (such as observed passenger booking behaviour, forecasting, and inventory control methodologies) in an all-else-equal simulation environment. It simulates the airline booking process in two parts: on one side there is a passenger decision model, which generates passengers who want to buy journeys; on the other side there is a revenue management system simulator, consisting of a database of historical bookings, a forecaster, and a seat allocation optimizer. The revenue management portion controls the availability of products, and the generated passengers choose amongst these products based upon their preferences and constrained by what is made available by the revenue management system.

The network environment used for simulation is defined by creating Origin-Destination markets, and assigning to them market distances; paths, travel times, and itineraries; average daily market demand; and market fares in a range of fare classes. Some previous network environments have been described in Chapter 2 (e.g. Wilson 1995; Lee 1998; Gorin 2004). In the subsequent sections of this chapter I will be deriving two new network environments, which will be called Network D6 and Network R.

Very thorough descriptions and explanations of the operations of the PODS program

can be found in Wilson (1995), Lee (1998), and Carrier (2003). They outline the functions of the various components, the passenger choice assumptions that are made, and the inputs required to run the simulation. These include such details as O-D market demands; the number of airlines and their leg capacities; passenger market segments and arrival patterns. A diagram of the simulation flow is provided in Figure 4-1, after Cléaz-Savoyen (2005). It illustrates the booking request process, from the arrival of a potential

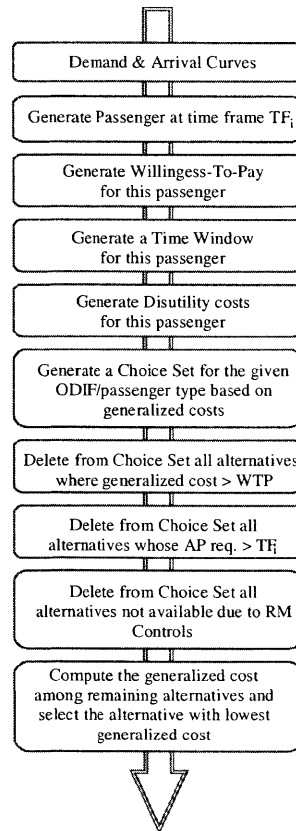


Figure 4-1 : PODS Simulation Flow

passenger, to the eventual decision made by the revenue management system. The reader is referred to the previous studies for more information about the entire process. However, for the purposes of this study, a few features of the passenger choice model and passenger willingness-to-pay will be described in some detail. These are the characterization of passenger price elasticities, and the effects of restrictions on generalized costs.

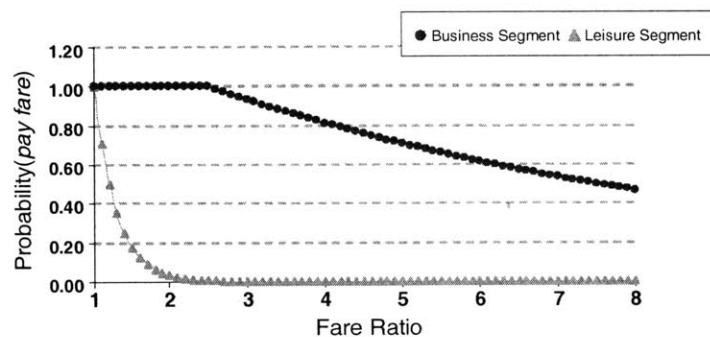
4.1.1 Willingness-To-Pay Considerations

In PODS, hypothetical passengers are generated at random during the booking process. They are either designated to be “business” or “leisure” passengers. They are randomly assigned a maximum willingness-to-pay for their itinerary, which is specified by two parameters: the Basefare and the emult (elasticity multiplier).

The Basefare is a characteristic of each market. It is provided to PODS as an input. It is defined as the fare at which all potential travellers will be willing to book. Each generated passenger in the market is then given a maximum willingness-to-pay. This willingness-to-pay is modelled in PODS as the probability that a random passenger will pay a particular given fare. The mean willingness-to-pay is defined by a negative exponential which is a function of fare ratio (the ratio between the highest fare and the lowest fare in the market) and passenger price elasticity:

$$\text{Probability}(\text{pay at least fare}) = \min\left[1, e^{\left(\frac{-0.6931}{emult-1}\right) \times \left(\frac{\text{fare} - \text{Basefare}}{\text{Basefare}}\right)}\right]$$

- Passenger Willingness-To-Pay Definitions**
- The probability that any randomly generated passenger will pay a given fare is calculated using the equation above
 - It is defined with respect to the market fares, as well as an input passenger choice parameter called the *emult* (elasticity multiplier)
 - By definition, all passengers will pay the Basefare (fare ratio=1)
 - The *emult* is the elasticity multiplier of the Basefare which defines the fare that 50% of travellers are willing-to-pay



• In this example, assumed *emult* values are 1.2 for Leisure passengers, and 3 for Business passengers

Figure 4-2: Willingness-To-Pay

In the example shown in Figure 4-2, all hypothetical business passengers are all willing to buy when the fare ratio is 2.5 or less, and 50% of them are willing to buy when the fare ratio is 5. Hypothetical leisure passengers are all willing to buy when the fare ratio is 1, and half of them are willing to buy when the fare ratio is 1.5. Very few leisure passengers are willing to buy when the fare ratio is 2. Some further implications of these elasticity quantities will be discussed in Section 4.2.4 below.

The other feature to note is PODS' modelling of fare product restrictions. Now that each passenger's willingness-to-pay has been defined, the passengers choose amongst the available products that are not more expensive than their maximum willingness-to-pay. The products are compared according to their generalized costs—that is, the sum of the product's nominal fare plus each passenger's valuation of the restrictions attached to the product.

When the network environment is defined, the set of fare products available in each market is also defined. This includes designating: the fare level (i.e., the nominal or "published" fare); the last time frame that the fare is available for sale (i.e., the Advance Purchase restriction); and which, if any, fare restrictions are associated with the product.

Table 4-1 : Example PODS Fare Structure

Fare Product	Published Fare	Advance Purchase Period	Strong Restriction Applied?	Restriction 2 Applied?	Restriction 3 Applied?
Y	\$400	0 days	0	0	0
B	\$300	7 days	Yes	0	0
M	\$200	14 days	Yes	Yes	0
Q	\$100	21 days	Yes	Yes	Yes

The current standard in PODS simulation is to apply up to three possible restrictions to the fare products. The first restriction is a strong restriction, and referred to as the Saturday Night Stay restriction. The other two restrictions are weaker, meaning that their perceived cost to passengers is less. The parameters at present in PODS assume that a Saturday Night Stay has a mean equivalent cost to "business" passengers of 2.7 times the Basefare, and 1.75 times the Basefare to "leisure" passengers. The other restrictions are each valued on average at 0.9 times the Basefare and 0.25 times the Basefare to "busi-

Table 4-2 : Example Generalized Costs

Fare Product	Published Fare	Perceived Cost to "Business" Passengers	Perceived Cost to "Leisure" Passengers
Y	\$400	$\$400 = \400	$\$400 = \400
B	\$300	$\$300 + 2.7 \times \$100 = \$470$	$\$300 + 1.75 \times \$100 = \$375$
M	\$200	$\$200 + 2.7 \times \$100 + 0.9 \times \$100 = \510	$\$200 + 1.75 \times \$100 + 0.25 \times \$100 = \350
Q	\$100	$\$100 + 2.7 \times \$100 + 0.9 \times \$100 + 0.9 \times \$100 = \$550$	$\$100 + 1.75 \times \$100 + 0.25 \times \$100 + 0.25 \times \$100 = \$325$

ness” and “leisure” passengers respectively. Hence, as we can see in Table 4-2 and Figure 4-3, the attractiveness of a product depends on its level of restriction. In line with the theory of product differentiation and market segmentation, the products with nominally lower fares become less attractive to business passengers when the valuation of restrictions is taken into account.

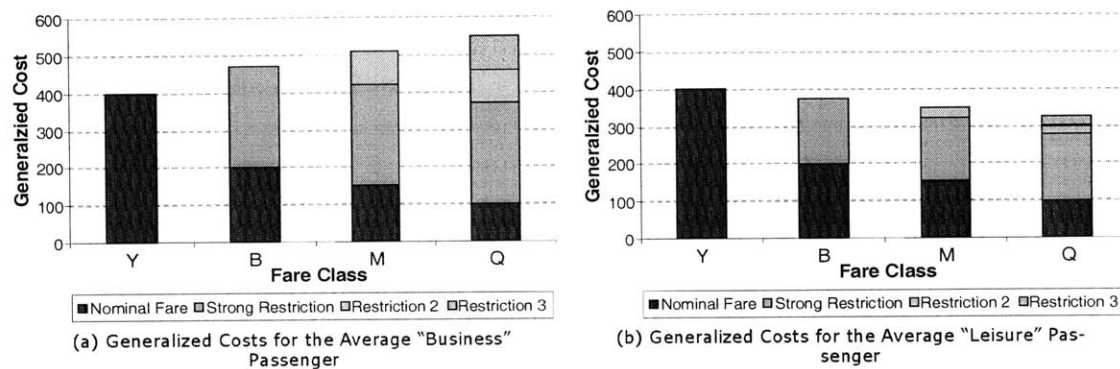


Figure 4-3 : Example PODS Generalized Costs

More details about the precise derivations and formulations of the PODS disutility models can be found in Wilson (1995) and Lee (2000: 45ff.). This brief outline is relevant to the present study because the aim is to investigate the performance of revenue management in less-differentiated fare environments – that is, with fewer restrictions on the fare products, and hence with a changed relationship between the generalized costs and the nominal fares. The issue of restrictions and generalized costs will be returned to in the next sections. Next, however, I will be describing the simulated network environments.

4.2 Network D

4.2.1 Characterizing Network D

The standard PODS simulation environment thus far has been known as Network D. I will characterize this environment in this section. Network D was developed originally in Lee's study of passenger disutility costs (Lee 2000: 42f.). It was a simplified representation of the US domestic environment, with two centrally located hubs and twenty spoke cities on either side, as shown in Figure 4-4. Its aim was to provide a wide variety of path choice options in order to calibrate the parameters of the disutility models. As such, it included

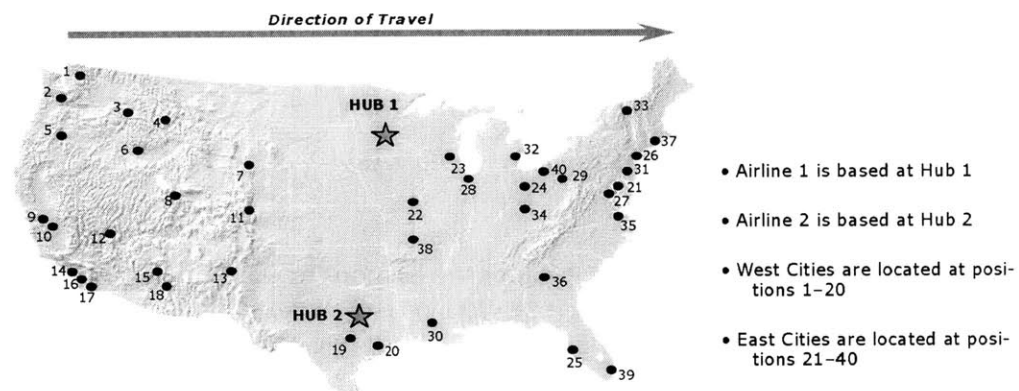
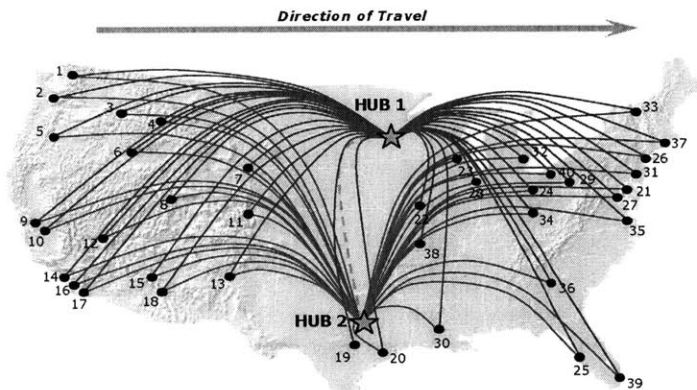


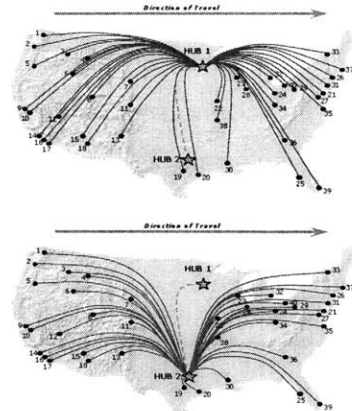
Figure 4-4: Network D Layout

three daily west-to-east connecting banks through each hub, as well as inter-hub flights by each carrier. The simulated passengers connecting between spoke cities could choose from three paths per day provided by each airline for a total of six possible paths per day. Connecting passengers had choice in terms of departure and arrival times, total trip time, and favourite airline. Passengers travelling in each airline's hub local markets also had access to other options by travelling on alternative connecting paths via the other airline's hub (see Figure 4.7(b)).

The two airlines each serve 482 Origin-Destination markets in each bank. With three banks and two airlines this gives 2892 total paths in the network. Of the 482 markets, each airline serves 42 local markets and 440 connecting markets (see Figures 4-5–4-7):

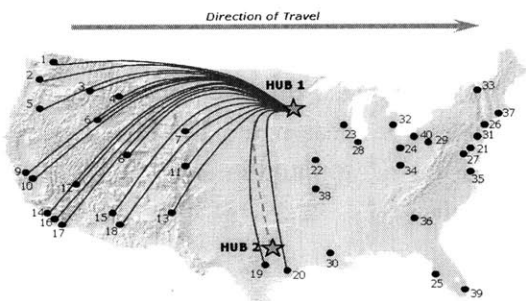


(a) Network D Route Map — Airline 1 connecting bank via Hub 1; Airline 2 connecting bank via Hub 2; inter-Hub flight

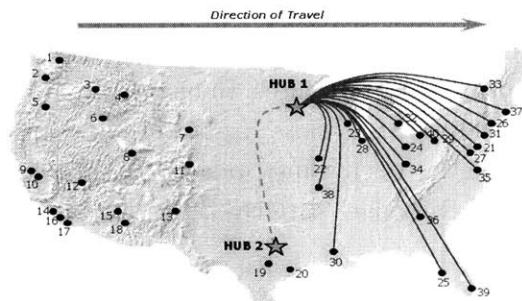


(b) Airline 1 Connecting Paths
(c) Airline 2 Connecting Paths

Figure 4-5 : Connecting Paths

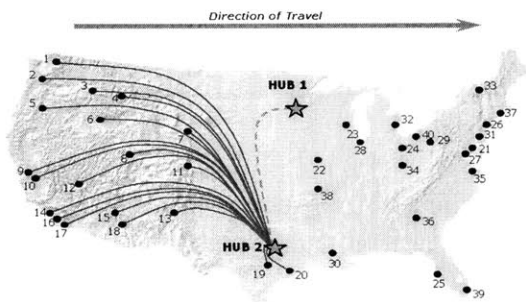


(a) Airline 1 Local Markets — West Cities into Hub 1; Hub 2 into Hub 1

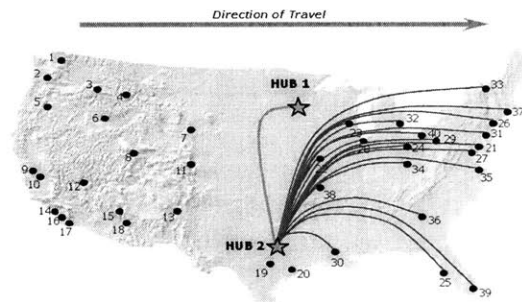


(b) Airline 1 Local Markets — Hub 1 to East Cities; Hub 1 to Hub 2

Figure 4-6 : Airline 1 Local Markets/Paths



(a) Airline 2 Local Markets — West Cities into Hub 2; Hub 1 into Hub 2



(b) Airline 2 Local Markets — Hub 2 to East Cities; Hub 2 to Hub 1.
Airline 2 here serves Airline 1 Local markets — it provides alternate paths via Hub 2

Figure 4-7 : Airline 2 Paths

- 40 local markets connecting the spoke cities and its hub;
- 2 local markets between the two hubs;
- 400 connecting markets between the twenty east and twenty west spoke cities;
- each airline also competes for the other airline’s local traffic by providing an alternative path between spoke cities and the other hub connecting via its own hub.

Table 4-3 : Network D Legs and Paths

Legs	Paths (Origin-Destination Itineraries)
42 legs in each route network	482 total O-D markets
3 banks – 126 legs flown per airline per day	3 banks – 1446 paths provided by each airline per day
252 legs flown in Network D per day	2892 paths provided in Network D per day

Network D is a symmetric network – Airline 1 and Airline 2 operate very similar schedules and offer identical fares. In its original formulation, it modelled the US fare environment with a four fare-class structure. The four fare classes were designated Y, B, M, and Q—in order of decreasing price and increasing restriction—as outlined in Table 4-4. Note the “triangular” pattern of the increasing restrictions:

Table 4-4 : Fare Structure—Classic Network D

Fare Class	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
Y	0 days	0	0	0
B	7 days	Yes	0	0
M	14 days	Yes	Yes	0
Q	21 days	Yes	Yes	Yes

The fares for each Origin-Destination market were calculated as a function of the O-D market distances, with input from some of the PODS Consortium airlines (Lee 2000: 45f). Over the next few years the fares saw a certain amount of adjustment and modification during various other simulation studies. Studies that were undertaken involved changing the market fare ratios, inverting some of the fares, and reducing and simplifying the fares in selected markets (e.g. see Cusano 2003, Gorin 2004). By the time work began

on the present study, the fare levels could be characterized as follows:

Table 4-5 : Fare Levels—Classic Network D

Fare Class	ALL MARKETS	AIRLINE 1 HUB LOCAL MARKETS	AIRLINE 2 HUB LOCAL MARKETS	CONNECTING MARKETS
	Average Fare (Standard Deviation)	Average Fare (S.D.)	Average Fare (S.D.)	Average Fare (S.D.)
Y	\$357.70 (\$116.01)	\$433.68 (\$116.07)	\$341.88 (\$115.68)	\$351.53 (\$113.47)
B	183.90 (57.07)	236.06 (96.58)	163.24 (51.51)	180.67 (48.91)
M	133.38 (35.18)	134.15 (37.16)	119.54 (38.45)	134.72 (34.41)
Q	97.96 (24.64)	82.71 (18.08)	86.46 (20.14)	100.70 (24.77)
Fare Ratio (SD)	3.74 (1.14)	5.32 (1.22)	4.06 (1.34)	3.54 (0.96)

Note the unusually high fare and fare ratio values seen in Airline 1’s local markets in comparison to Airline 2’s local markets and the industry as a whole. This phenomenon can likely be attributed to historical biases in the fare data which was based on input provided by members of the PODS Consortium.

For these fares, correlation with market distance was quite weak. The relationship between market distance and the Y fare (highest fare) and the Q fare (lowest fare) is graphed in Figure 4-8. The R^2 value of Q with respect to distance is 0.1898, and the R^2 value of Y with respect to distance is 0.0667. While a direct linear relationship between fares and distance is neither desirable nor realistic, nevertheless there is still some association between these quantities. However, these very low R^2 values indicate that the variation in fare levels is barely—if at all—related to variation in market distance:

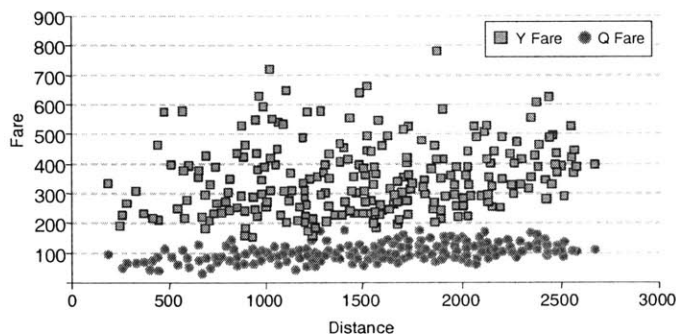


Figure 4-8 : Market Fares by Market Distance—Classic Network D

This simulation network as it has been described here thus far will be referred to as “Classic Network D”. PODS simulation was performed in this environment. Each airline used a revenue management system that consisted of: fare class yield management based on EMSRb; pick-up forecasting; and booking curve detruncation. The relative performance of each airline is:

Table 4-6(a) : Performance Metrics—Classic Network D

Network	Revenue Share		ASM Share		RPM Share	
	Airline 1	Airline 2	Airline 1	Airline 2	Airline 1	Airline 2
Classic Network D	50.4%	49.6%	49.2%	50.8%	49.1%	50.9%

Table 4-6(b) : Revenue Performance—Classic Network D

Network	Benefit of FCYM over FCFS		
	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%

Note that although Airline 2 provides a greater share of the capacity and carries more than half of the traffic, Airline 1 still earns a larger share of the revenue, with 1.8% more revenue than Airline 2. This is a reflection of its market and fare advantage as referred to above. In Table 4-6(b), the benefit of using fare-class yield management (FCYM) over not using any revenue management (First Come First Served—FCFS) can be seen for both airlines and for the industry as a whole.

The further benefits to Airline 1 of using O-D control in this environment have been described in Cusano (2003: 62) as follows: further revenue gains of 1.55% when implementing DAVN or ProBP; and 0.86% when implementing HBP. Figure 4-9 shows the

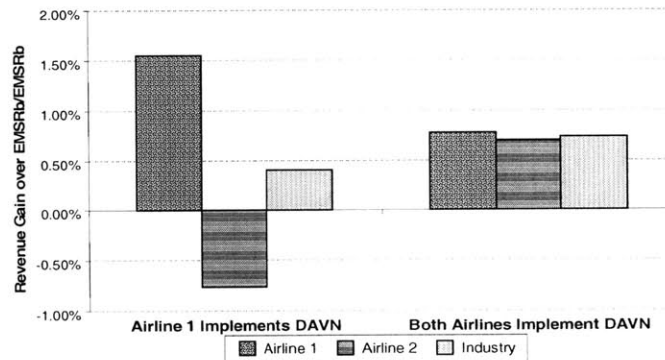


Figure 4-9 : Benefits of DAVN—Classic Network D

1.55% gain to Airline 1 with respect to EMSRb when implementing DAVN. Airline 2 sees a revenue drop of 0.76%, however the industry as a whole benefits by 0.4% when Airline 1 moves to DAVN. If both airlines implement DAVN, then they see even revenue gains – 0.77% and 0.69% gains with respect to both using EMSRb. In this case the industry as a whole benefits more, with a total gain of 0.73%.

Finally, looking at the network fare class mix in Figure 4-10, we can see a “boat-shaped” distribution of passengers over the four fare classes. That is, the lowest fare class sees the largest average load, the highest fare class sees the second highest load, and the two intermediate fare classes see lower loads. This result is characteristic of load distributions in differentiated fare environments.

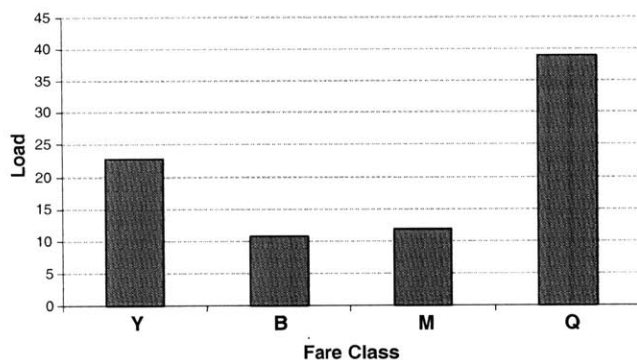


Figure 4-10: Fare Class Mix—Classic Network D—ALF=84

So much for Classic Network D. In the wake of the widespread industry fare structure simplification, fare compression, and revenue spiral down that has been described in the previous chapters, it seemed timely to review this modelling environment. In studying the potential for revenue management methods in less-differentiated fare environments, it is necessary to develop a modelling environment in which airline performance more closely resembles the revenue and load results that airlines in current less-differentiated environments are seeing. Traditional US airlines have estimated that fare simplification has cost them around a sixth of their revenues. Feedback from PODS Consortium airlines and falling industry yields indicate that the “boat-shaped” distribution is no longer an accurate representation of contemporary purchasing behaviour. More seats are being sold in lower fare classes, and far fewer high-fare class products are being sold. And so in light of these issues, the development of a network model that is more appropriate to contemporary market conditions will now be described.

4.2.2 Developing a Simplified Network Environment

The first step I took to develop a model that is more appropriate to contemporary conditions was to alter the fare structure. The fare structure that was outlined above in Table 4-4 is a four fare-class fully differentiated structure, meaning that each product is differentiated in terms of restrictions and advance purchase requirement, and the “triangular” restriction structure assures that cheaper products get progressively more restricted. In consultation with industry partners in the PODS Research Consortium, the basis of a six fare class fare structure was developed. In this structure there is no “Saturday Night Stay” requirement, the level of differentiation is lessened, and the fare ratio is reduced as shown in Table 4-7.

Table 4-7 : Fare Structure—Network D6 Flat

Fare Class	Fare Ratio	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
FC1	3.4	0 days	0	0	0
FC2	2.48	0 days	0	Yes	0
FC3	1.6	7 days	0	Yes	Yes
FC4	1.4	14 days	0	Yes	Yes
FC5	1.2	14 days	0	Yes	Yes
FC6	1	21 days	0	Yes	Yes

In this fare structure, there is one fully unrestricted product (Fare Class 1), and one product with only a single light restriction (Fare Class 2). The remaining products are not differentiated from each other in terms of restrictions. They have increasing levels of advance purchase restriction, but even so Fare Class 4 and Fare Class 5 are only differentiated from each other in terms of price.

For this new formulation of Network D, the existing Y/B/M/Q fares were modified to create six fare levels. The market Basefares were held constant, and the new lowest fare for each market (Fare Class 6) was set at \$10 below the Basefare. Next, each O-D market was assigned a fare ratio of 3.4 between Fare Class 1 and Fare Class 6. The intermediate fares were related to Fare Class 6 as defined in Table 4-5. That is, if the Fare Class 6 fare was \$125, then the other fares would be \$150, \$175, \$200, \$310 and \$425. Because of the six fare levels and the constant fare ratio across every single market, this version of the network will be referred to as “Network D6 Flat”. Table 4-8 shows the new average fares in this network, and compares them to the fares in Classic Network D. The figures

Table 4-8 : Fare Levels—Network D6 Flat

Classic Network D Fare Class	ALL MARKETS Average Fare (Standard Deviation)	Network D6 Flat Fare Class	ALL MARKETS Average Fare (Standard Deviation)
Y	\$357.70 (\$116.01)	FC1	\$299.07 (\$83.77)
B	183.90 (57.07)	FC2	218.14 (61.10)
M	133.38 (35.18)	FC3	140.74 (39.42)
Q	97.96 (24.64)	FC4	123.14 (34.49)
		FC5	105.55 (29.57)
		FC6	87.96 (24.64)
Fare Ratio (SD)	3.74 (1.14)	Fare Ratio (SD)	3.40 (0.00)

(Figure 4-11 and 4-12) show some of the consequences of fare simplification on passenger choice. In Classic Network D, the lower fare-class products are less attractive to business passengers because of the perceived costs of the restrictions that are associated with them. However, now that the fares have been lowered and the Saturday Night Stay

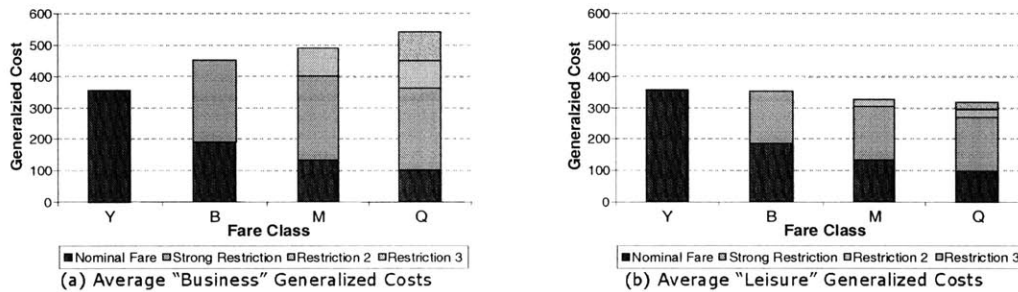


Figure 4-11 : Generalized Costs—Classic Network D

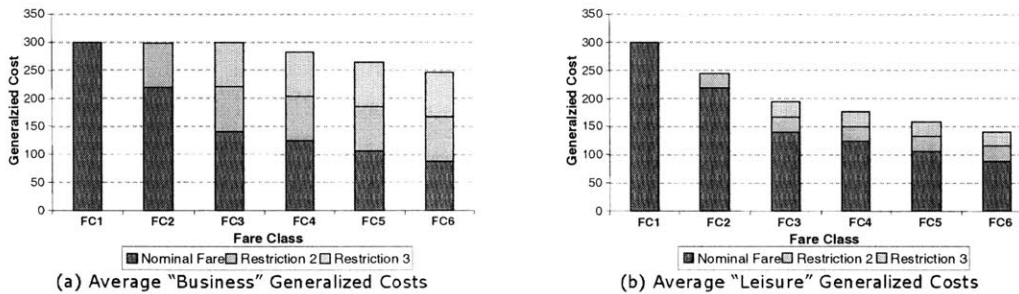


Figure 4-12 : Generalized Costs—Network D6 Flat

restriction has been removed, we can see that the three most expensive products have more or less the same perceived total value for business passengers: the restrictions on Fare Class 2 and Fare Class 3 are not sufficiently severe as to make those products unattractive. In Classic Network D, any business passenger who is willing to pay for Y Class would buy it, because it is effectively “cheaper” than B Class. Now in Network D6 Flat, any business passenger who is willing to pay for Fare Class 1 will still buy it, but only if Fare Classes 2 & 3 are not available. Now they will certainly buy down if they can, and the main differentiating factor between the top three products becomes the advance purchase restriction. And so, we can see that removing restrictions and compressing fares can have the effect of increasing consumer surplus to the detriment of airlines.

Performing PODS simulation as before with both airlines using EMSRb, the following results were obtained:

Table 4-9(a) : Performance Metrics—Network D6 Flat

Network	Revenue Share		ASM Share		RPM Share	
	Airline 1	Airline 2	Airline 1	Airline 2	Airline 1	Airline 2
Classic Network D	50.4%	49.6%	49.2%	50.8%	49.1%	50.9%
Network D6 Flat	49.9%	50.1%	49.1%	50.9%	49.1%	50.9%

Table 4-9(b) : Revenue Performance—Network D6 Flat

Network	Benefit of FCYM over FCFS			Revenue Loss wrt Classic Network D		
	Airline 1	Airline 2	Industry	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%	–	–	–
Network D6 Flat	32.3%	30.3%	31.3%	-21.4%	-19.7%	-20.5%

With the new fares, the market imbalance between the airlines is less. Airline 2 carries slightly more traffic than it did in Classic Network D, but it now also has a greater share of the revenue, with 0.44% more revenue than Airline 1. In the less-differentiated environment, note also the benefit of using fare-class revenue management. Table 4-9(b) shows that having removed much of the segmentation power from the fare structure, the benefit of using EMSRb to control leg traffic has increased to 31%. Note that this is in marked contrast to conventional wisdom, which holds that the benefit of fare-class yield management is about 4-6% over FCFS. Nevertheless, reducing fares and removing restrictions has an obvious and expected negative effect on overall revenues, and there is a revenue degradation on the order of 20% with respect to Classic Network D.

The next step in the modelling process was to change fare ratios so that they were not constant across all markets. The fare ratios were perturbed at random about their mean of 3.4 so that they ranged between 2.5 and 4.3. Average fares were almost unchanged, and the average fare ratio was naturally still 3.4. In this network, called “Network D6 Perturbed” PODS simulation showed that revenue and traffic shares were not very different from those observed in Network D6 Flat; the values in Table 4-9(a) apply equally to both networks. However, note from Table 4-10 that the revenue loss in Network D6 Perturbed with respect to Classic Network D is not as great as for Network D6 Flat. In other words, perturbing the fare ratios leads to revenues that are about 1.1% higher than were seen in Network D6 Flat.

Table 4-10 : Revenue Performance—Network D6 Perturbed

Network	Benefit of FCYM over FCFS			Revenue Loss wrt Classic Network D		
	Airline 1	Airline 2	Industry	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%	–	–	–
Network D6 Flat	32.3%	30.3%	31.3%	-21.4%	-19.7%	-20.5%
Network D6 Perturbed	30.6%	28.4%	29.5%	-20.6%	-18.8%	-19.7%

At this stage, it was thought to be a good idea to revise the fares in a more unified manner. Thus far the fare structure had been changed into a less-differentiated six product structure, and revenues had fallen in accordance with expectations and reflecting industry events. However, the revenue drop was at the high end of industry estimates of the consequences of fare simplification. The fares were still poorly correlated with distance, and some of them were unreasonably low, which may have been contributing to this. In addition, the fare class mix was no longer boat-shaped, but the distribution was not particularly intuitive either: now the load in Fare Class 5 was higher than in the cheaper Fare Class 6, and high loads were still being seen in the two most expensive fare classes, contradicting industry observations.

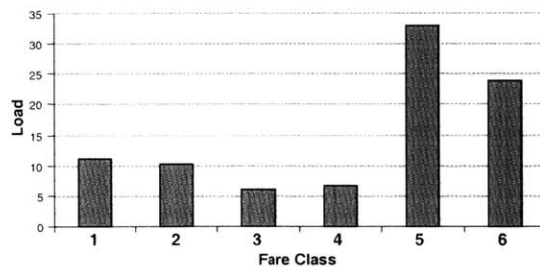


Figure 4-13 : Fare Class Mix—Network D6 Perturbed—ALF=90.8

4.2.3 Re-building Network D—“Network D6”

The goal thus far has been to make the simulation environment more reflective of what has been observed in the marketplace with respect to revenue changes and fare class mix. We would expect to see a more modest loss in revenue than is seen in Network D6 Perturbed, and a fare class mix that shows fewer high priced products and more low priced products being sold.

The fares that were arrived at for Network D6 Perturbed were still not well correlated with market distance. R^2 values were 0.1898 for the Basefares and 0.1414 for the highest fare (see Figure 4-14). The Basefares ranged between \$28 and \$175: the lowest base fares were found in Airline 2’s local markets and were unrealistically low. Therefore, the base fares were transformed so that they ranged between \$60 and \$170, and the fares were arbitrarily re-allocated among the O-D markets so that their relationship to market distance was stronger. Then the market fare ratios were once again generated at random,

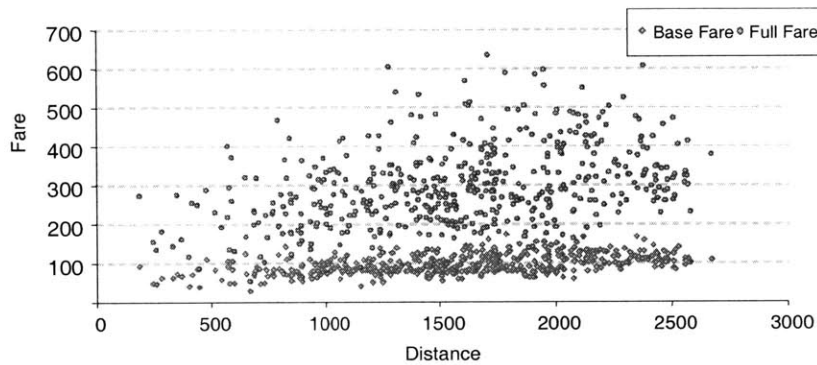


Figure 4-14 : Market Fares by Market Distance—Network D6 Perturbed

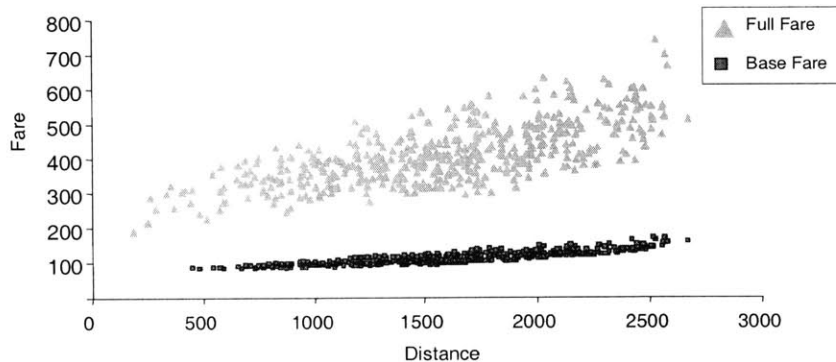


Figure 4-15 : Market Fares by Market Distance—Network D6

this time ranging between 3.2 and 5.0 with an average network fare ratio of 4.1. In each market, Fare Class 6 was defined to be 90% of the Basefare, and the other fare levels were calculated in accordance with the fare ratio for that market. Now the R² values were 0.7847 for the Basefares and 0.4775 for the highest fare (see Figure 4-15).

The figures show that the new fares are now much better correlated with market distance. The new average fares in each fare class are now higher, and it can be seen in Table 4-11 that the bias towards Airline 1's local markets has been removed, and the average fare levels and fare ratios in each airline's local markets are now more even.

Table 4-11 : Fare Levels—Network D6

Network D6 Flat Fare Class	ALL MARKETS	Network D6 Fare Class	ALL MARKETS	AIRLINE 1 HUB LOCAL MARKETS	AIRLINE 2 HUB LOCAL MARKETS
	Average Fare		Average Fare	Average Fare	Average Fare
FC1	\$299.07 (<i>\$83.77</i>)	FC1	\$412.86 (<i>\$83.41</i>)	\$350.84	\$359.62
FC2	218.14 (<i>61.10</i>)	FC2	293.34 (<i>56.56</i>)	248.69	254.42
FC3	140.74 (<i>39.42</i>)	FC3	179.01 (<i>31.59</i>)	150.98	153.79
FC4	123.14 (<i>34.49</i>)	FC4	153.03 (<i>26.23</i>)	128.77	130.92
FC5	105.55 (<i>29.57</i>)	FC5	127.05 (<i>21.17</i>)	106.56	108.04
FC6	87.96 (<i>24.64</i>)	FC6	101.06 (<i>16.71</i>)	84.36	85.17
Fare Ratio (<i>SD</i>)	3.40 (<i>0.00</i>)	Fare Ratio (<i>SD</i>)	4.09 (<i>0.51</i>)	4.17	4.22

Simulating this Network in PODS, we find the following performance as outlined in Tables 4-12. Traffic shares are now similar to traffic shares in Classic Network D, but now the revenue shares are more even. Airline 1 now has a smaller advantage:

Table 4-12(a) : Performance Metrics—Network D6

Network	Revenue Share		ASM Share		RPM Share	
	Airline 1	Airline 2	Airline 1	Airline 2	Airline 1	Airline 2
Classic Network D	50.4%	49.6%	49.2%	50.8%	49.1%	50.9%
Network D6 Flat&Perturbed	49.9%	50.1%	49.1%	50.9%	49.1%	50.9%
Network D6	50.1%	49.9%	49.1%	50.9%	49.1%	50.9%

The benefit of using fare class revenue management to control leg traffic is now around 23%. The new fares have reduced somewhat the effects of fare simplification, and the revenue degradation with respect to Classic Network D is now only around 13% – which is now at the lower end of industry estimates.

Table 4-12(b) : Revenue Performance—Network D6

Network	Benefit of FCYM over FCFS			Revenue Loss wrt Classic Network D		
	Airline 1	Airline 2	Industry	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%	–	–	–
Network D6 Flat	32.3%	30.3%	31.3%	-21.4%	-19.7%	-20.5%
Network D6 Perturbed	30.6%	28.4%	29.5%	-20.6%	-18.8%	-19.7%
Network D6	24.0%	22.5%	23.2%	-13.4%	-12.4%	-12.9%

Looking at the fare class mix, however, we still do not see hoped for results. There are far fewer seats being sold in Fare Class 1, but there is a spike in the distribution at Fare Class 2, and more seats are sold in Fare Class 3 than in Fare Class 4. Neither of these

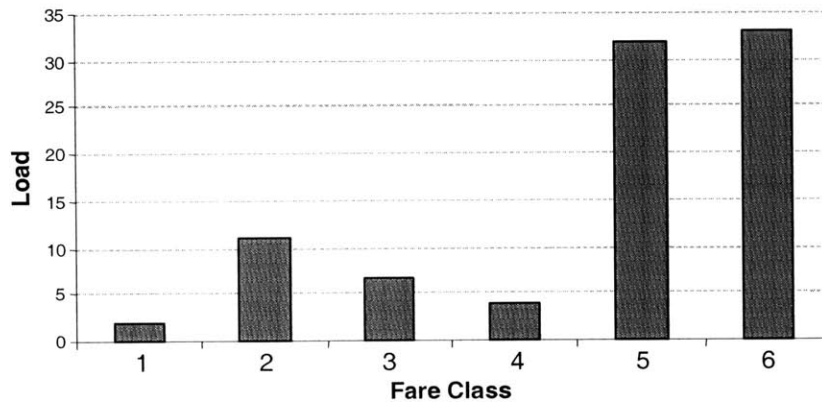


Figure 4-16 : Fare Class Mix—Network D6 Perturbed—ALF=88

aspects of the distribution are consistent with industry observations. We expect to see more seats being sold in each progressively cheaper fare class. Hence, the next step was to look at the passenger willingness-to-pay and sell-up behaviour that has been assumed in this modelling environment, to see if they are reasonable in light of current conditions.

4.2.4 Price Elasticity Multipliers

Passenger price elasticity behaviour is defined in PODS via two parameters: *emult* and *bfbmult*. The *bfbmult* parameter defines the relationship between the price elasticities of the two market segments (business and leisure). As described above, the *emult* parameter defines the price elasticities within each of the market segments.

The *bfbmult* parameter is the *Basefare business multiplier*. If the Basefare is the fare that **all potential travellers** would be willing to pay, then (*bfbmult*) x (Basefare) defines the fare **all potential “business” travellers** would be willing to pay.

The *emult* parameter is the *elasticity multiplier* of the Basefare. If the Basefare is the fare that all potential travellers would be willing to pay, then (*emult*) x (Basefare) defines the fare that 50% of all potential travellers would be willing to pay.

There are different *emults* and for business and for leisure passengers: we multiply the Basefare by the leisure *emult*, and we multiply the “business” Basefare by the business *emult*. That is:

- (Basefare) x (leisure *emult*) gives the fare at which 50% of all potential “business” travellers would book;
- (*bfbmult*) x (Basefare) x (business *emult*) gives the fare at which 50% of all potential “business” travellers would book.

The values of these parameters govern the relationship between the fares on offer and the willingness-to-pay of the simulated passengers. This relationship is defined in terms of fare ratios. The values used in Classic Network D were *bfbmult* = 2.5; business *emult* = 3; leisure *emult* = 1.2. This meant that:

- all leisure passengers were willing to pay a fare that is equal to the Basefare;
- all business passengers were willing to pay a fare that is 2.5 times the Basefare;
- on average, 50% of business passengers were willing to pay a fare that is 7.5 times the Basefare; and
- on average, 50% of leisure passengers were willing to pay a fare that is 1.2 times the Basefare.

However, since the fare ratios in Network D6 are different to those in Classic Network D, it is necessary to revisit these parameters.

In Network D6, the average value of the fare ratio between Fare Class 1 and Fare Class 6 is 4.1, and it ranges between 3.2 and 5.0. This means that the average value of the fare ratio between Fare Class 1 and the Basefare is 90% of this, i.e., it ranges between 2.88 and 4.5. Given the current elasticity multipliers, this means that 50% of business passengers are willing to pay a fare that is well in excess of the highest fare available in the network. The relationship between *emult*, fare ratio, and business passenger willingness-to-pay is shown in Figure 4-17. It shows that as the business *emult* is reduced from 3 to 2 to 1.5, the fare that 50% of business passengers are willing to pay is reduced accordingly from 7.5 times the Basefare to 3.75 times the Basefare. This relationship is superimposed over the

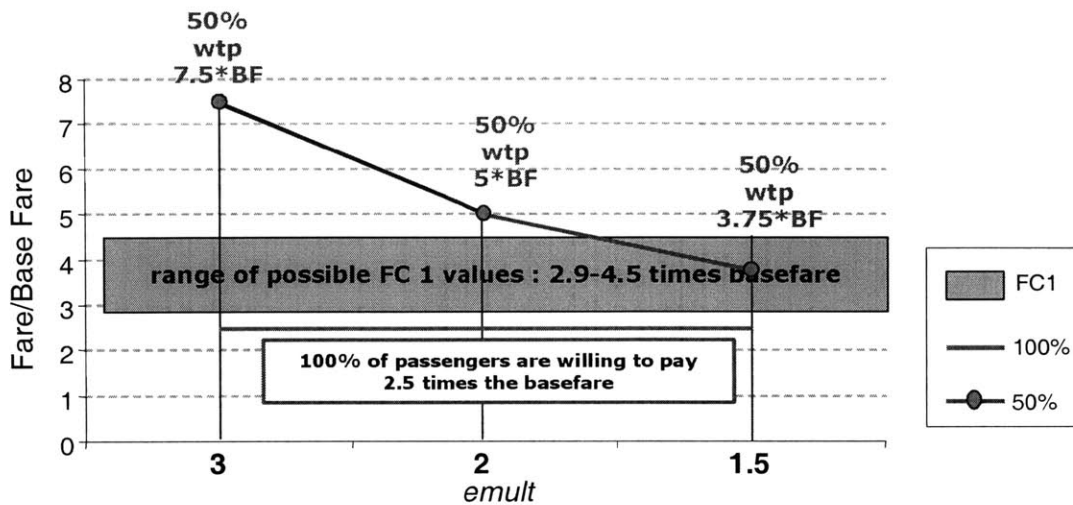


Figure 4-17: Business *emult* and Network D6 Fare Ratios

range of Fare Class 1 fares that are used in Network D6. We can see that this network's price elasticity assumptions are not appropriate given these new fare levels. Therefore, it is necessary to adjust these model parameters. A business *emult* of 3 is too high – it implies that 50% of business passengers are willing to pay a fare that is on the order of double the highest network fare. On the other hand, a business *emult* of 1.5 seems too low – it implies that 50% of business passengers are not willing to pay for a good deal of the Fare Class 1 products on offer. Hence, a business *emult* of 2 was selected. With this value, on average 50% of business passengers are willing to buy all of the Fare Class 1 products on offer in this network.

Similarly, looking at the value of the current leisure *emult* parameter in Figure 4-18, we can see that 50% of leisure passengers are only willing to pay for Fare Class 5 under a leisure

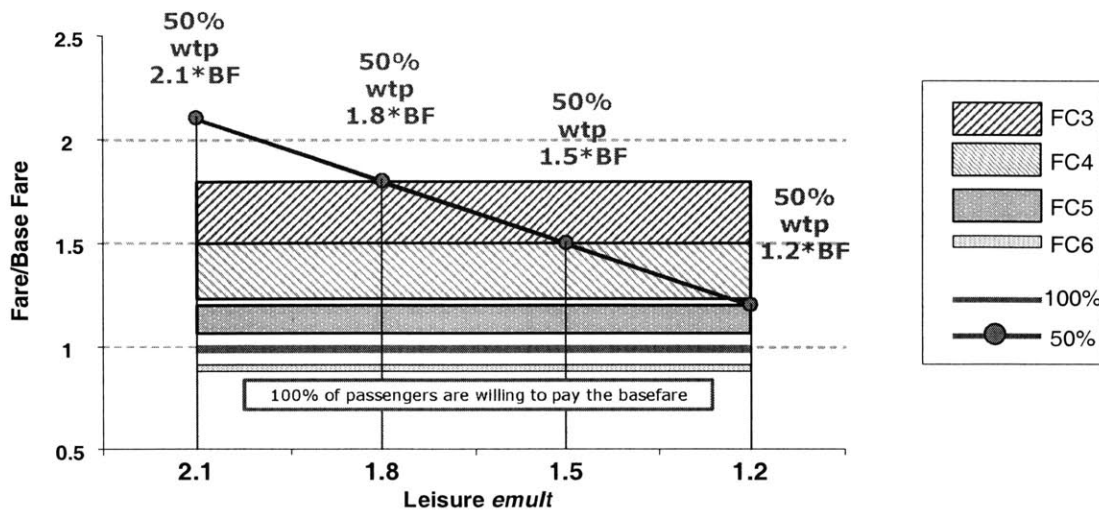


Figure 4-18 : Leisure emult and Network D6 Fare Ratios

emult of 1.2. If the leisure emult is assumed to be 1.5 or 1.8 then we would see that 50% of leisure passengers would be willing to buy Fare Class 4 and Fare Class 3 respectively. It seems too optimistic to expect that half of all passengers would be willing to pay for the third most expensive product (Fare Class 3), so the leisure emult of 1.8 is too high. Conversely we could expect that half of all passengers would be willing to buy higher than merely the second cheapest product (Fare Class 5), and so the leisure emult of 1.2 seems too low. Hence, a leisure emult of 1.5 was selected. With this value, 50% of leisure passengers are willing to buy up to Fare Class 4. This means that if the cheaper products are not available for whatever reason, on average 50% of all passengers will be willing to sell-up into Fare Class 4.

Having changed the emult and bsbmult parameters, we are almost ready to return to consider the network. But first, one more parameter needs reconsidering. Under the new fares in Network D6 the average load factor (ALF) for each airline over the entire network is about 88. This is very high, and not seen in the US environment. Thus, the total network demand in Network D6 needs to be scaled back somewhat. In PODS, demands are defined individually for each O-D market, but there is also a multiplier that can be used to scale demand over the entire network. This parameter was thus lowered from a factor of 1.0 to a factor of 0.9.

Simulating in PODS, we find that revenue and market shares are almost unchanged. Looking at revenue management and overall revenue performances (Table 4-13) we can

see that the benefit of fare class revenue management is around 19%. Revenue degradation with respect to Classic Network D is around 16%, which is about the level expected.

Table 4-13 : Revenue Performance—Network D6 (new elasticity parameters)

Network	Benefit of FCYM over FCFS			Revenue Loss wrt Classic Network D		
	Airline 1	Airline 2	Industry	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%	–	–	–
Network D6 Flat	32.3%	30.3%	31.3%	-21.4%	-19.7%	-20.5%
Network D6 Perturbed	30.6%	28.4%	29.5%	-20.6%	-18.8%	-19.7%
Network D6	24.0%	22.5%	23.2%	-13.4%	-12.4%	-12.9%
Network D6 (new mults.)	19.2%	17.9%	18.6%	-16.7%	-15.6%	-16.2%

Looking at the average fare class mix in the network — the relative distribution in Fare Classes 3 & 4 is improved; however the Fare Class 2 spike is still in evidence. And so, the final adjustment to the new network environment will be to make the Fare Class 2 product less attractive to passengers. In the current network formulation, the differentiation between Fare Classes 1 & 2 relies on the price difference, and a single weak restriction.

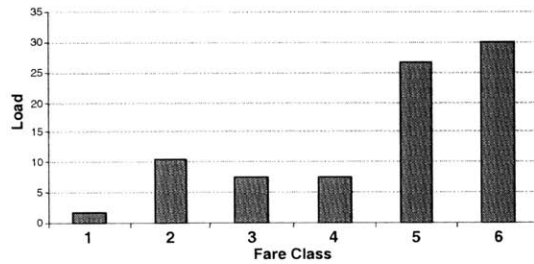


Figure 4-19 : Fare Class Mix—Network D6 (new mults.)—ALF=84.5

Looking at the generalized costs now with the new fares and the new business Basefare multiplier (Figure 4-20), we can see that the perceived costs of the fare products to business passengers are such that the products with higher nominal prices are never more attractive than the cheaper products. In other words, these low levels of restriction do not provide significant disincentives to business passengers. From the graph, it is easy to

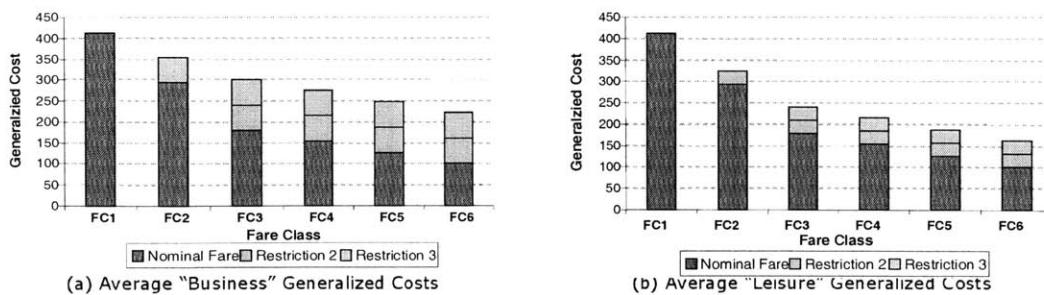


Figure 4-20 : Generalized Costs—Network D6

see that the average passenger would not choose Fare Class 1 if Fare Class 2 were available. Fare Class 2 is nominally cheaper and cheaper in terms of perceived cost, and since there is no advance purchase restriction on Fare Class 2, it is likely to be available most of the time. Thus, a small advance purchase restriction of 3 days was added to Fare Class 2 in an attempt to force some sell-up into Fare Class 1.

Table 4-14 : Fare Structure—New Network D6

Fare Class	Fare Ratio	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
FC1	3.4	0 days	0	0	0
FC2	2.48	3 days	0	Yes	0
FC3	1.6	7 days	0	Yes	Yes
FC4	1.4	14 days	0	Yes	Yes
FC5	1.2	14 days	0	Yes	Yes
FC6	1	21 days	0	Yes	Yes

Finally, the new network, called “Network D6” is ready. Simulating in PODS and looking at the average fare class mix (Figure 4-21), we can see that the load in Fare Class 1 has increased at the expense of the load in Fare Class 2. Now there is a more realistic distribution, with the loads carried in each fare class decreasing with increasing price, and the highest load being carried in Fare Class 6.

Looking at the airlines’ revenue performance (Table 4-15), we can see that the total revenue degradation is around 16%, and the benefit of using fare class revenue management is still around 19%.

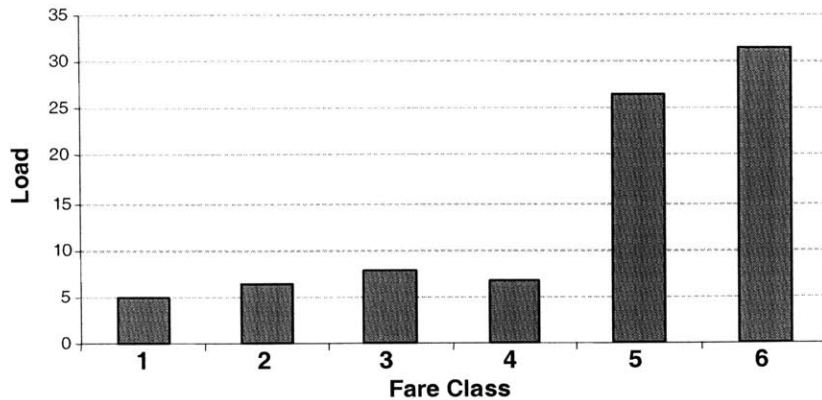


Figure 4-21 : Fare Class Mix—Network D6—ALF=84

Table 4-15 : Revenue Performance—New Network D6

Network	Benefit of FCYM over FCFS			Revenue Loss wrt Classic Network D		
	Airline 1	Airline 2	Industry	Airline 1	Airline 2	Industry
Classic Network D	8.6%	5.7%	7.2%	–	–	–
Network D6 Flat	32.3%	30.3%	31.3%	-21.4%	-19.7%	-20.5%
Network D6 Perturbed	30.6%	28.4%	29.5%	-20.6%	-18.8%	-19.7%
Network D6 (new mults.)	19.2%	17.9%	18.6%	-16.7%	-15.6%	-16.2%
Network D6	19.5%	18.2%	18.8%	-16.5%	-15.4%	-16.0%

4.2.5 Summary—Network D6

In summary, the new PODS modelling environment is called “Network D6”. It has 6 fare classes, consisting of two differentiated high fare class products and four less differentiated lower fare class products. Market fare ratios range between 3.2 and 5.0. Passenger elasticity behaviour is assumed to be such that the relationship between the price elasticities of the business and leisure market segments (*bfbmult*) is a factor of 2.5; and the price elasticities within each of the market segments (*emults*) are such that 50% of all passengers are willing to pay a fare that is 1.5 times the Basefare, and 50% of all business passengers are willing to pay a fare that is 5 times the Basefare. In this environment, the benefit of

using fare-class revenue management over not using any revenue management at all is around 19% for both airlines and for the industry as a whole.

Table 4-16 : Fare Levels & Structure—New Network D6

Network D6 Fare Class	Average Fares	Average Fare Ratio	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
FC1	\$412.86 (<i>\$83.41</i>)	3.4	0 days	0	0	0
FC2	293.34 (<i>56.56</i>)	2.48	3 days	0	Yes	0
FC3	179.01 (<i>31.59</i>)	1.6	7 days	0	Yes	Yes
FC4	153.03 (<i>26.23</i>)	1.4	14 days	0	Yes	Yes
FC5	127.05 (<i>21.17</i>)	1.2	14 days	0	Yes	Yes
FC6	101.06 (<i>16.71</i>)	1	21 days	0	Yes	Yes

This new network appears to be a better representation of current market conditions. With its two very similar airlines, it is suitable for use in simulating symmetric environments, and to test the effects of changes in revenue management methods when all is held equal and when the competitors are evenly matched. However, there is also value in having a non-symmetric modelling environment. In the next section, such an environment will be developed.

4.3 Network R

Network D6 has now been fully characterized and adapted for use in less-differentiated fare environments. We can appreciate the particular advantages of its symmetric structure, in particular its utility in performing “all else equal” simulation experiments. However, Network D6 is missing two important aspects of the current competitive landscape: namely, asymmetry, and mixed fare structures. Asymmetry is important because the largest airline marketplaces do not tend to see much symmetric head-to-head competition in all of their markets. In order to investigate the usefulness of revenue management in more complex and uneven environments, it will be helpful to build an asymmetric modelling environment. Similarly with mixed fare structures: nowhere in the world has there been complete fare simplification uniformly across all markets, but rather there exists a mixture of traditional and less-differentiated fare structures in many networks. So while network D is valuable for abstract all else equal studies, a new network is needed that might be a better representation of competitive realities in deregulated markets such as the US or the EU. Hence, I will now briefly characterize this network, which will be called “Network R.”

Network R is a 4-hub network. It is not symmetric like Network D6, but rather contains four airlines with different sizes and scopes. Three of airlines (Airlines 1, 2, and 4) are “traditional,” and the other airline is a “low-cost carrier.” The main airline, Airline 1 based in MSP, serves every O-D market in the network. It has a close competitor, Airline 2 based in ORD, which shares most of its O-D markets. Airline 2 is geographically situated near to Airline 1, and so its paths and travel times are quite similar. Third, there is a low-cost competitor, Airline 3 based in MCI, which competes in about half of Airline 1’s O-D markets. Finally, there is a smaller “traditional” airline, Airline 4 based in DFW. It is intermediate in size between Airline 3 and the large traditional airlines, and is removed geographically from all of them. Table 4-17 summarizes each Airline’s constituent legs and O-D markets.

Table 4-17 : Network R Legs and Paths

Airline 1 MSP	572 O-D Markets; 1457 Total Paths
	<ul style="list-style-type: none"> • 24 Origin Cities (20 west spokes; the 4 hubs) and 24 Destination Cities (20 east spokes; the 4 hubs)
	<ul style="list-style-type: none"> • 46 local markets
	<ul style="list-style-type: none"> • 460 spoke-originating connecting markets; 66 other-hub originating connecting markets
	<ul style="list-style-type: none"> • 525 O-D markets in Bank 1; 357 O-D markets in Bank 2; 572 O-D markets in Bank 3; 3 hub-bypass legs

Table 4-17 : Network R Legs and Paths

Airline 2 ORD	548 O-D Markets; 1458 Total Paths
	<ul style="list-style-type: none"> • 24 Origin Cities (20 west spokes; the 4 hubs) and 23 Destination Cities (19 east spokes; the 4 hubs)
	<ul style="list-style-type: none"> • 45 local markets
	<ul style="list-style-type: none"> • 440 spoke-originating connecting markets; 63 other-hub originating connecting markets
	<ul style="list-style-type: none"> • 548 O-D markets in Bank 1; 356 O-D markets in Bank 2; 548 O-D markets in Bank 3; 6 hub-bypass legs
Airline 3 MCI	296 O-D Markets; 907 Total Paths
	<ul style="list-style-type: none"> • 15 Origin Cities (11 west spokes; the 4 hubs) and 20 Destination Cities (16 east spokes; the 4 hubs)
	<ul style="list-style-type: none"> • 33 local markets
	<ul style="list-style-type: none"> • 209 spoke-originating connecting markets; 54 other-hub originating connecting markets
	<ul style="list-style-type: none"> • 296 O-D markets in each bank 1; 19 hub-bypass legs
Airline 4 DFW	428 O-D Markets; 1044 Total Paths
	<ul style="list-style-type: none"> • 18 Origin Cities (14 west spokes; the 4 hubs) and 24 Destination Cities (20 east spokes; the 4 hubs)
	<ul style="list-style-type: none"> • 40 local markets
	<ul style="list-style-type: none"> • 322 spoke-originating connecting markets; 66 other-hub originating connecting markets
	<ul style="list-style-type: none"> • 364 O-D markets in Bank 1; 248 O-D markets in Bank 2; 428 O-D markets in Bank 3; 4 hub-bypass legs

All of the airlines provide connecting services through their hubs, and they also offer non-stop hub bypass flights. In addition, the low-cost competitor also offers some head-to-head competition in four of Airline 1's local markets. The airline route networks are illustrated in Figures 4-22–4-25.

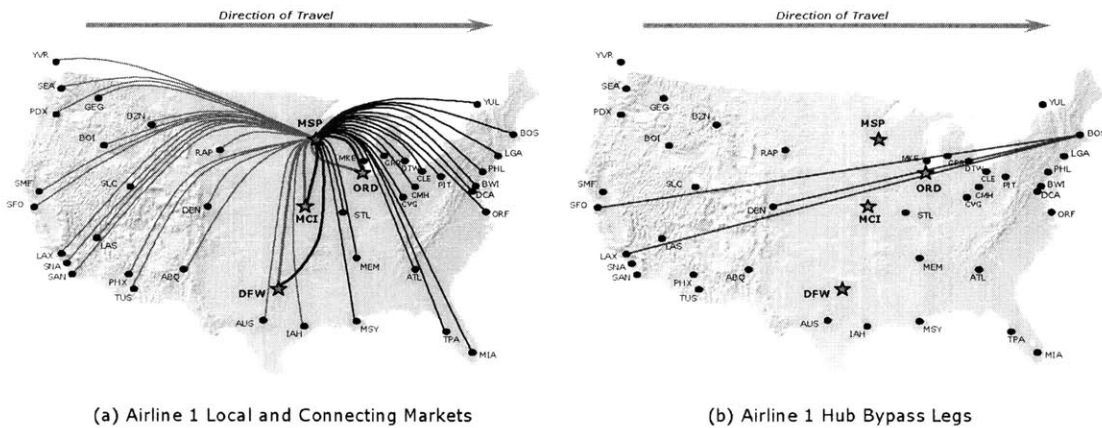
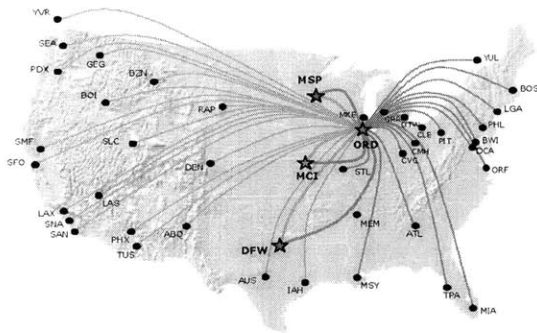
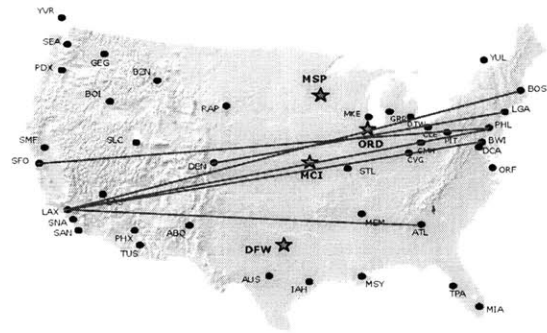


Figure 4-22 : Airline 1—MSP Local Markets/Paths

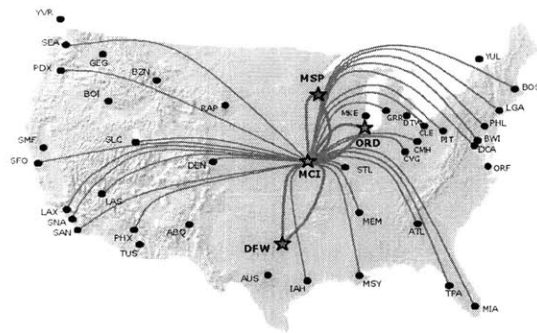


(a) Airline 1 Local and Connecting Markets

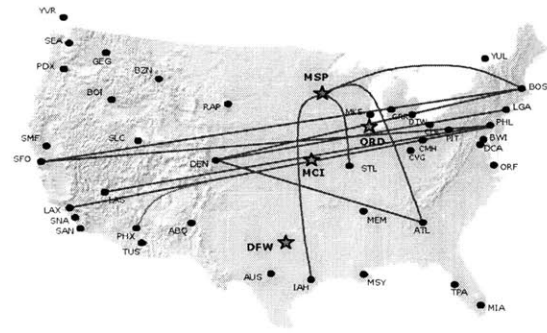


(b) Airline 1 Hub Bypass Legs

Figure 4-23 : Airline 2—ORD Local Markets/Paths

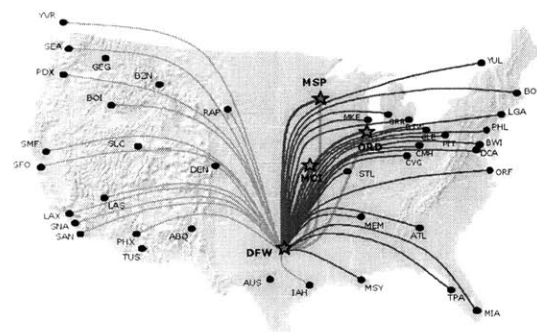


(a) Airline 1 Local and Connecting Markets

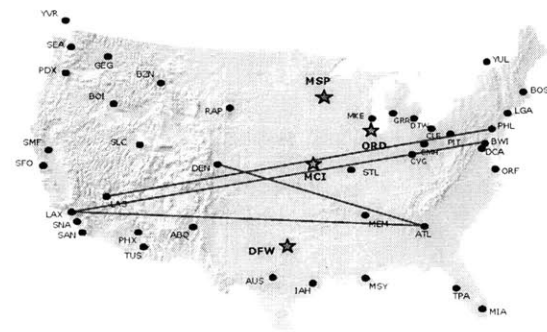


(b) Airline 1 Hub Bypass Legs. Note also the direct competition in four of Airline 1's local markets

Figure 4-24 : Airline 3—MCI Local Markets/Paths



(a) Airline 1 Local and Connecting Markets



(b) Airline 1 Hub Bypass Legs

Figure 4-25 : Airline 4—DFW Local Markets/Paths

Airline 3 competes in 296 of the network’s 572 O-D markets. It uses a less-restricted fare structure with compressed fares, and all of the “traditional” airlines match it in the markets in which it competes with them. In the remaining markets, the traditional airlines use a restricted six class fare structure—hence the traditional airlines can be said to be operating “mixed-fare” networks, meaning that they have both differentiated fare products and less-differentiated fare products on offer on many of their legs.

Table 4-18(a) : Fare Levels—Network R Non-LCC Markets

Network R Non-LCC Markets	Average Fares	Average Fare Ratio	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
FC1	\$812.93 (<i>\$148.08</i>)	5	0 days	0	0	0
FC2	562.91 (<i>97.78</i>)	3.47	3 days	0	Yes	0
FC3	323.76 (<i>50.47</i>)	2	7 days	Yes	Yes	0
FC4	269.41 (<i>40.13</i>)	1.67	14 days	Yes	Yes	0
FC5	215.06 (<i>30.28</i>)	1.33	14 days	Yes	Yes	Yes
FC6	160.71 (<i>51.59</i>)	1	21 days	Yes	Yes	Yes

Table 4-18(b) : Fare Levels—Network R LCC Markets

Network R LCC Markets	Average Fares	Average Fare Ratio	Advance Purchase Period	Strong Restriction	Restriction 2	Restriction 3
FC1	\$429.24 (<i>\$91.34</i>)	4.1	0 days	0	0	0
FC2	304.78 (<i>61.83</i>)	2.91	3 days	0	Yes	0
FC3	185.73 (<i>34.44</i>)	1.78	7 days	0	Yes	Yes
FC4	158.67 (<i>28.57</i>)	1.52	14 days	0	Yes	Yes
FC5	131.62 (<i>23.06</i>)	1.26	14 days	0	Yes	Yes
FC6	104.56 (<i>18.23</i>)	1	21 days	0	Yes	Yes

In Tables 4-18, we can see the differences between the two fare structures. The average

fare ratio in the Non-LCC markets is higher than for the LCC markets. In the Non-LCC markets it ranges between 4 and 6; in the LCC markets it ranges between 3.2 and 5 (as in Network D6). The level of restriction is also different: in the Non-LCC markets each of the fare products is differentiated either in terms of restriction or advance purchase requirement, or both; in the LCC markets, the same set of restrictions is used as for Network D6.

Looking at the generalized cost charts for the Non-LCC markets (Figure 4-26), we can see that the Fare Class 2 product generally has a lower perceived cost than the Fare Class 3 product both for the business and leisure segments. Fare Class 2 is still more attractive than Fare Class 1, meaning that Fare Class 1 would only ever be purchased if Fare Class 2 was not available. In the less-differentiated LCC markets (Figure 4-27), we can see that perceived cost decreases as price decreases in all cases, so that the more expensive products are never more attractive to consumers, which is the same as for Network D6.

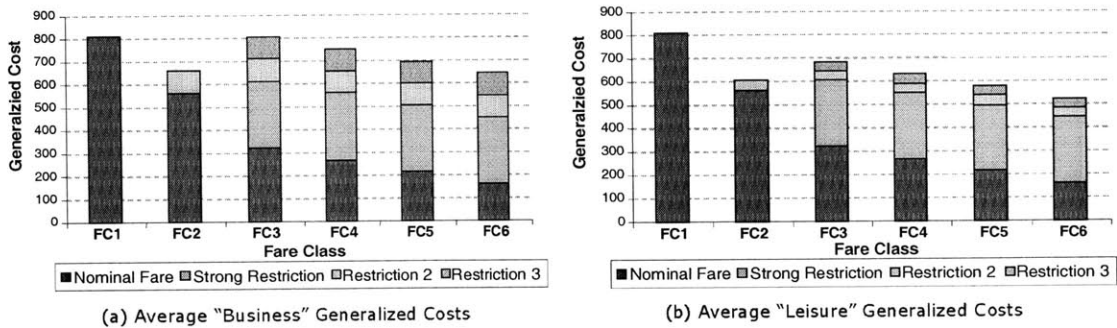


Figure 4-26: Generalized Costs—Network R Non-LCC Markets

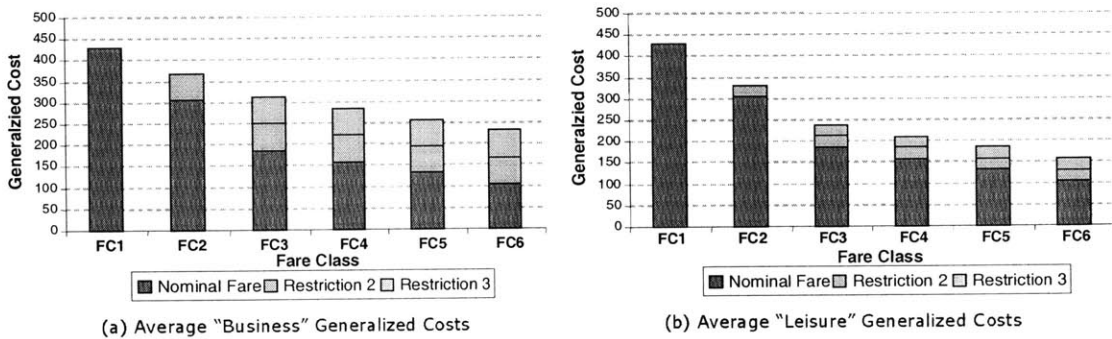


Figure 4-27: Generalized Costs—Network R LCC Markets

Finally, the Network R environment is simulated in PODS, with each airline using an EMRSb leg-based revenue management system. Looking at the relative revenue and traffic performance of each airline (Table 4-19(a)), we can see that as expected Airlines 1 & 2 carry the largest shares of the traffic and earn the largest shares of the revenue (~30%). Airline 3 is smaller, and Airline 4 is intermediate. Fare class yield management benefits all airlines (Table 4-19(b)). It seems to benefit the low-cost carrier the most, since the fare structure in all of its markets is not very differentiated, and hence it relies on revenue management to segment its demand. Revenue management is again seen to benefit the industry as a whole.

Table 4-19(a) : Performance Metrics—Network R

Network		Revenue Share					
Network R		Airline 1	Airline 2	Airline 3	Airline 4		
		31.6%	32.2%	15.8%	20.4%		
ASM Share				RPM Share			
Airline 1	Airline 2	Airline 3	Airline 4	Airline 1	Airline 2	Airline 3	Airline 4
29.7%	31.9%	16.9%	21.6%	29.8%	31.9%	16.9%	21.5%

Table 4-19(b) : Revenue Performance—Network R

Network	Benefit of FCYM over FCFS				
Network R	Airline 1	Airline 2	Airline 3	Airline 4	Industry
	10.3%	9.6%	19.3%	12.5%	11.8%

4.4 Conclusion

In this chapter, I have presented two new PODS simulation environments—Network D6 and Network R. Their aim is to model contemporary conditions in the world airline industry. By using both Network D6 and Network R, we can now start to understand the insights that both have to offer. On the one hand, Network D6 can be used to isolate the effects of different revenue management methods and forecasters in a head-to-head “all else equal” environment. On the other hand, Network R can be used to investigate competitive effects of revenue management in a more complex network.

Hence, in Chapter 5 I will investigate the performance of some revenue management

systems in these two network environments. I will test the optimization and forecasting methods that were described in Chapter 3 in these two different environments that have been outlined above. By considering the revenue management systems and the modelling environments together, we can get a broader picture of the strengths and weaknesses of the different methods in different competitive situations.

5 Results

Having developed the two simulation environments in the previous chapter, we can now use them to investigate the benefits of the revenue management systems and forecasting methodologies that were described in Chapter 3. In this chapter, the implementation of each of the four methods—EMSRb, DAVN, HBP, and ProBP—will be simulated in Network D6 and in Network R. In the first two sections of this chapter, each of the methods will be tested in combination with pick-up forecasting; in the subsequent two sections they will be tested in combination with hybrid forecasting.

5.1 Network D6—Pick-up Forecasting

In this section, we take a systematic look at the performance of the three Origin-Destination inventory control methods in Network D6, with respect to leg-based revenue management. The baseline environment has both of the competitors using Fare Class Yield Management—FCYM. This is leg-based EMSRb optimization with pick-up forecasting, as introduced Chapter 4: recall that leg-based revenue management provided each airline with benefits of about 19% over First Come First Served (Table 4-15). We will use this environment to look at the incremental benefits of different O-D control methods in a context in which fares are compressed and fare products have less differentiation.

The FCYM/FCYM case is used as a baseline because leg-based revenue management is very widely employed in industry, making FCYM a good benchmark for comparison. This baseline is also useful because it allows us to look at the “pure” benefits of implementing O-D control. Since the network is symmetric—the airlines offer very similar service in terms of schedule—we can construct a situation in which the predominant distinction between the airlines is their revenue management system. This allows us to look into the possibilities for revenue management to aid airlines that are seeing revenue degradation in less-differentiated environments, for example by making adjustments to their forecasting methods, or by implementing O-D control.

Recall from Section 4.2.1 the incremental benefits of O-D control methods in Classic Network D (Figure 5-1). In this section, we will look how the efficacy of these methods changes in a less-differentiated fare environment.

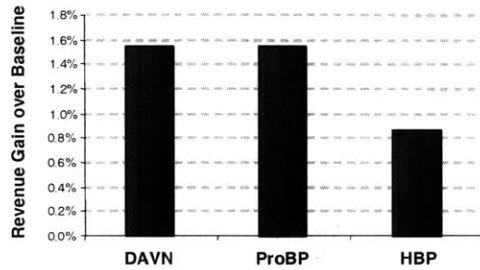


Figure 5-1 : Incremental Benefits of O-D Control—Classic Network D

5.1.1 DAVN

In Classic Network D, it was found that implementing DAVN helps the airline that makes that move, and while it also increases total industry revenues, it does cause the other airline to see some revenue loss. Referring to Step 1 in Figure 5-2, when Airline 1 implements DAVN, it sees a revenue gain of 1.55% and Airline 2 a loss of 0.76%. Then in Step 2, when Airline 2 implements DAVN, it sees a revenue gain of 1.46% and Airline 1 sees a loss of 0.77%. In total, when both airlines use DAVN, their cumulative revenues are 0.77% and 0.69% higher than when they both used FCYM. At both steps in the process, the industry captures additional revenue from the consumer surplus for a total industry gain of 0.73%.

Now looking at the implementation of DAVN in the new less-differentiated environment, we see a quite different pattern. In Step 1, when Airline 1 implements DAVN, it sees a revenue gain of 0.74%, but now Airline 2 also sees a gain of 1.15% (Figure 5-3). Note that the revenue gain to the passive airline is larger than the gain to the airline that makes the move. In this environment, the percentage gain to each airline is less than for Classic Network D, but the gain to the industry overall is greater (0.94%). Then, in Step 2, when Airline 2 implements DAVN, it actually sees a loss compared to the case where only Airline 1 was using DAVN, seeing a 0.82% drop in revenues. Airline 1 is able to capture some additional revenue, but total industry revenues are higher when only one airline uses DAVN.

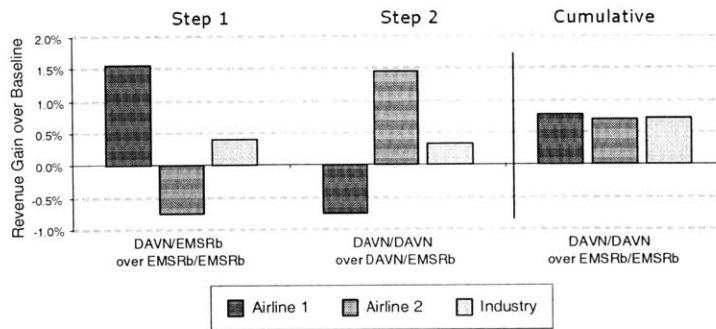


Figure 5-2: Incremental Benefits of DAVN—Classic Network D

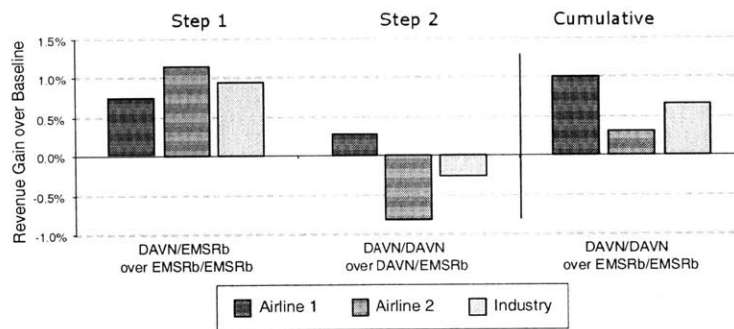


Figure 5-3: Incremental Benefits of DAVN—Network D6

Bear in mind that when both airlines use DAVN, their revenue is still higher than for the case in which they both used FCYM: the cumulative benefits are 1.01% for Airline 1 and 0.31% for Airline 2, and the industry gains 0.67%. Counterintuitively however, Airline 2 and the total industry revenues are better off in the situation in which only Airline 1 implements DAVN. Airline 1 also benefits in this case, but its best scenario occurs when Airline 2 implements DAVN as well.

Table 5-1: DAVN Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
DAVN	+0.74%	82.8	0.1022	FCYM	+1.15%	83.1	0.0980
DAVN	+1.01%	82.2	0.1031	FCYM	+0.31%	81.2	0.0994

Looking at performance measures, we can see in Table 5-2 that revenues and yields increase in Network D6 with respect to the FCYM baseline case, for both airlines over both steps. We can also see that the network average load factors (ALF) are lower, since the DAVN revenue management system is choosing a better mix of higher yield traffic.

Counterintuitive Effects of DAVN

Why has the pattern of revenue gains and losses changed in the new fare environment? The answer lies in the passive airline's revenue sources. Each airline's revenue can be said to come from a combination of three sources:

- "First Choice" revenue, which is earned when the product purchased is the passenger's first choice;
- "Sell-up" revenue, which is earned when the passenger's first choice of product is not available, and they agree to buy a more expensive product provided by the same airline; and
- "Spill-in" revenue, which is earned when the passenger's first choice of product is not available, and they then purchase a product offered by a different airline, i.e., they are "spilled" to the other airline.

When Airline 1 introduces O-D control, Airline 2 loses "First Choice" revenue, and it gains "Sell-up" revenue and "Spill-in" revenue. In Network D6, it gains much more spill-in revenue than it did in Classic Network D. In Classic Network D, the contribution of Spill-in revenue to Airline 2's total revenue rose from 4% in the baseline case to 5% in the case where Airline 1 implemented DAVN. In Network D6, the contribution of Spill-in revenue to Airline 2's total revenue rose from 7% in the baseline case to 13% in the case where Airline 1 implemented DAVN.

Why is the Spill-in now more lucrative to the passive airline? When Airline 1 introduces DAVN, the average fares paid to the passive airline tend to fall, as Airline 1 is able to better choose which booking requests to accept. In Network D6 we find that the value of the passengers being carried in Airline 2's intermediate fare classes is higher than it was for Classic Network D. Fares paid for Fare Classes 4 & 5 after Airline 1 implements DAVN are on average \$1 higher, and fares paid for Fare Class 3 are on average \$5 higher.

This implies that passengers who are being rejected from Airline 1 by DAVN are finding availability on Airline 2, since Airline 2's leg-based revenue management system does not have the same power to choose amongst the passengers in terms of origin-destination

paths. Moreover, since the fares in this network are lower and have smaller differentials and very slight differentiation, passengers are now more willing to buy up into a higher fare class, and therefore the “spilled” passengers are more lucrative than they were in Classic Network D.

Since the revenues of both airlines are higher with respect to the FCYM baseline case when both airlines have implemented DAVN, we can see that DAVN is still functioning as expected. It is still rejecting passengers that are of lower network value in favour of higher value passengers. We would expect that Airline 2 would see a drop in yield, and hence in revenue, when Airline 1 introduces DAVN, since Airline 1 is able to reject low yield connecting passengers whereas Airline 2 using FCYM is not. However, from the above we can see that in this new fare environment, the losses that Airline 2 experiences in step 1 when Airline 1 introduces DAVN are actually outweighed by the revenue gains it sees from spilled passengers who are selling-up. Hence, it sees a gain in yield, and its average network load factor remains steady (see Table 5-2).

In step 2, when Airline 2 also introduces DAVN, we can see similar results, with Airline 2 now actually losing revenue. This indicates that the benefits to Airline 2 of using DAVN are actually less than the benefits of spill-in from Airline 1’s implementation of DAVN. We can surmise that although DAVN is still doing its job, the feedback effects in this symmetrical 2-airline environment can overwhelm its incremental benefits.

Beyond the feedback and spill-in, another factor to note is the amount of buy-down overall in this fare environment. DAVN’s assumption of independent class demand is problematic here because the distinctions between the fare products are so weak, and it is evident that passengers are more willing to buy up or buy down, irrespective of official fare class designation. Therefore, the disjunction between the forecasting philosophy and the realities of passenger purchasing behaviour are harming DAVN’s performance.

5.1.2 ProBP

As was the case for DAVN, implementing ProBP in Classic Network D boosts the revenues of the airline that makes the move (+1.55%), and causes the other airline to see some revenue loss (-0.81%). But, as was the case for DAVN, in Network D6 this traditional pattern of gains and losses is not seen (Figure 5-4).

Looking at ProBP in the new less-differentiated environment, when Airline 1 implements

ProBP it sees a benefit of 0.43%, but now Airline 2 sees a much larger benefit—in fact Airline 2’s revenue gain of 1.05% is now almost 2.5 times larger than Airline 1’s revenue gain.

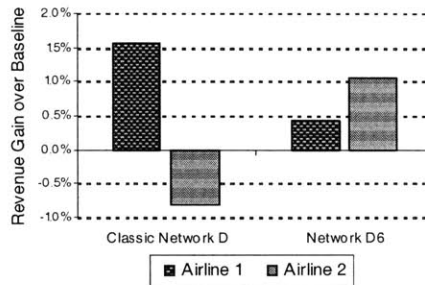


Figure 5-4 : Incremental Benefits of ProBP—Classic D & Network D6

In Table 5-3, we can see that revenues and yields increase with respect to the FCYM base-line case for both airlines, and that the average load factors are reduced. Similar to the DAVN case, ProBP is clearly still doing its job and is choosing a higher revenue mix of traffic. But it is suffering from the same feedback effects that were evident in the performance of DAVN.

Table 5-2 : ProBP Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
ProBP	+0.43%	82.5	0.1022	FCYM	+1.05%	83.1	0.0978

Decline in ProBP Performance

Another point to note is that in Classic Network D the gain to Airline 1 when implementing ProBP was 1.55%, which was the same as the gain to Airline 1 when implementing DAVN. However in Network D6, the gain to Airline 1 when implementing DAVN is 0.74%, whereas the gain to Airline 1 when implementing ProBP is only 0.43%. Why has ProBP’s performance suffered so much in comparison to DAVN?

To answer this question, a test was carried out comparing the performance of these two methods across Network D6 and in a slightly modified version of Network D6 called

Network D6 “Restricted.” In this version of Network D6, more restrictions are added to some fare products, to create more product differentiation (see Table 5-4). Fare Classes 1 & 2 are the same, but the four lower fare classes are now each differentiated from each other, either in terms of restrictions or advance purchase period or both. In all other respects – i.e., fares and fare ratios – Network D6 and Network D6 “Restricted” are the same.

Table 5-3 : Fare Structure—Network D6/Network D6 “Restricted”

Fare Class	Advance Purchase Period	Network D6			Network D6 “Restricted”		
		Strong Restriction	Restriction 2	Restriction 3	Strong Restriction	Restriction 2	Restriction 3
FC1	0 days	0	0	0	0	0	0
FC2	3 days	0	Yes	0	0	Yes	0
FC3	7 days	0	Yes	Yes	Yes	Yes	0
FC4	14 days	0	Yes	Yes	Yes	Yes	0
FC5	14 days	0	Yes	Yes	Yes	Yes	Yes
FC6	21 days	0	Yes	Yes	Yes	Yes	Yes

The results of this test are shown in Figure 5-5, and we can see that when more restrictions are added to the fare products, the performances of DAVN and ProBP are once again comparable, providing revenue gains of +0.63% and +0.59% respectively. Hence we can conclude that ProBP is more adversely affected by the removal of restrictions. With some level of differentiation it performs as well as DAVN, but when differentiation is removed it has difficulty with spiral down. As noted in Chapter 3, ProBP is a sensitive revenue management method. It controls inventory at the path level, depending on path-class forecasts and controlling exclusively with leg-specific bid-prices. Perhaps it is this reliance on such fine resolution forecasting that makes it particularly vulnerable to spiral down in less-differentiated fare environments, in which it is very difficult to forecast the distinct demands for undifferentiated fare products by path and by fare class.

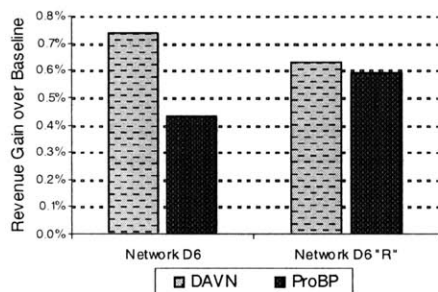


Figure 5-5 : DAVN and ProBP Performance

5.1.3 HBP

We turn now to look at the performance of the new formulation of HBP in the less-differentiated environment. It is important to note that this HBP is not strictly comparable to the version tested in previous studies (e.g. Cusano 2003). It has been reformulated since those tests were done, and now includes an increment to its local demand forecast which estimates unconstrained local demand (as described in Chapter 3). Nevertheless, Figure 5-6 illustrates the performance of the older HBP in Classic Network D, and the performance of the new HBP in Network D6.

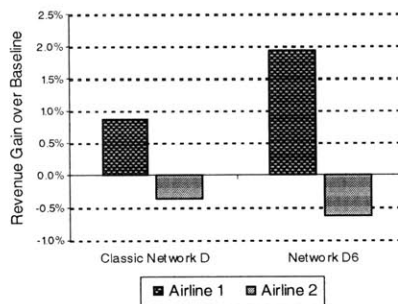


Figure 5-6 : Incremental Benefits of HBP—Classic D & Network D6

For the HBP case, the pattern of revenue gains and losses is now the same in both environments. Indeed, the gain to Airline 1 from implementing the new HBP in Network D6 is greater than was the gain to Airline 1 from implementing the older HBP in Classic Network D. Looking at yields and load factors for the HBP case, Airline 1’s loads now increase and its yields decrease; whereas Airline 2’s loads decrease while its yield increases. HBP is accepting more traffic in this environment and is not producing as much spill – i.e., HBP is not giving as much traffic away to its competitors. This is related to HBP’s use of leg forecasting. The resolution of its forecasts is larger, making it less liable to inflate its forecasts. Thus it doesn’t over-forecast for higher fare classes, and it allows more lower fare class traffic, increasing total revenues.

Table 5-4 : HBP Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	–	83.6	0.1004	FCYM	-	83.2	0.0967
HBP	+1.93%	87.8	0.0974	FCYM	-0.63%	80.9	0.0989

5.1.4 Summary

Considering the performance of all three methods with respect to the FCYM baseline (Figure 5-7), we can see that the new HBP is clearly the best performer in this less differentiated environment. This is likely a result of its not using path-based forecasting or optimization. Forecasting by fare class is difficult in the less differentiated environment even at a leg level, and leads to spiral down. These results indicate that it is even more of a problem at the path level, because of the much smaller mean observed demands for a large proportion of the many path/fare class combinations, and the resulting tendency for the forecasts to overestimate demand. This seems to a serious factor hindering the performance of ProBP in particular.

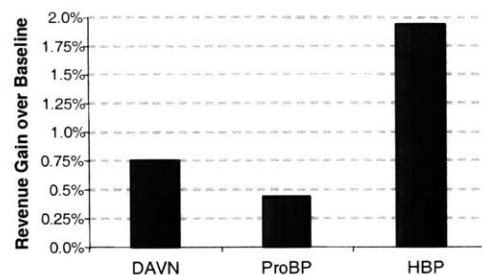


Figure 5-7: Incremental Benefits of O-D Control—Network D6

HBP is also the only method to display the traditional pattern of revenue gains and losses when implemented by Airline 1, despite the high feedback in this high demand symmetric environment. It is able to accept more traffic and is not allowing its competitor to capture a larger share of the revenue. In contrast, feedback effects appear to dominate the performance of DAVN and ProBP in this environment, in which passengers are willing to buy up or buy down rather than reliably buying particular fare products in particular fare classes, as forecast by the path/class forecaster.

Table 5-5 : O-D Control Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
DAVN	+0.74%	82.8	0.1022	FCYM	+1.15%	83.1	0.0980
ProBP	+0.43%	82.5	0.1022	FCYM	+1.05%	83.1	0.0978
HBP	+1.93%	87.8	0.0974	FCYM	-0.63%	80.9	0.0989

Consequently, we now turn to look at the performance of these methods in Network R. Network R has many more paths and offers more capacity than Network D6; the route networks offered by the airlines are somewhat different from each other; and its “traditional” airlines offer a mixed fare structure. Hence, the feedback effects that were seen in Network D6 may be alleviated in the more complex environment of Network R.

5.2 Network R—Pick-up Forecasting

After looking at the performance of these revenue management systems in Network D6 with respect to a symmetric FCYM base case, we now turn to a more complex network. In this section we will look at the effects of implementing O-D control in the more elaborate non-symmetric competitive environment. To this end, the base case will not call for all of the airlines to use FCYM, but will also be more elaborate. The focus airline will once again be Airline 1, and it will be used to look at the incremental benefits of DAVN, ProBP, and HBP. Its baseline method will be FCYM. The “traditional” competitors (i.e., Airlines 2 & 4) will use DAVN, and the low-cost carrier (Airline 3) will use a threshold revenue management method called AT80¹. In this way Airline 1 faces close competition in most of its connecting markets from a large competitor with an advanced revenue management system (Airline 2); and it also sees discount competition in just over half of its connecting markets and direct competition in some of its local markets from a smaller competitor with a relatively simplistic revenue management system (Airline 3). Airline 4 then functions to provide more capacity in the market so as to avoid feedback problems.

The performance of this baseline case (EMSRb/DAVN/AT80/DAVN) with respect to no airlines using any revenue management at all is shown in Table 5-7. As might be expected, these benefits are greater than those listed in Table 4-19(b), which compared all airlines using leg-based revenue management with First Come First Served. This is because the Network R baseline offers a more sophisticated revenue management environment than one in which all of the airlines use FCYM.

Table 5-6 : Benefits of Revenue Management—Network R

Network	Benefit of Baseline Case over FCFS				
Network R	Airline 1	Airline 2	Airline 3	Airline 4	Industry
	10.9%	12.9%	20.4%	14.8%	13.6%

1. AT80 is a dynamic threshold revenue management method that progressively closes down fare classes on a leg as they reach target load thresholds. Its total target leg load factor is 80. See Cléaz-Savoien (2005: 31; 101f) for more details regarding threshold revenue management methods and AT80.

We now turn to the performance of the three methods in Network R.

5.2.1 DAVN

In Classic Network D, recall that implementing DAVN helped the airline that made the move with a revenue gain of 1.55%, and led to a revenue loss for the other airline. Now, when Airline 1 implements DAVN in Network R, it also sees a revenue gain of 1.51%. Airline 2 sees a loss of 1.55%, and the other two airlines also see revenue losses of 1.17% and 0.9% (Figure 5-8). These smaller losses are reflective of their diminished exposure to Airline 1’s network.

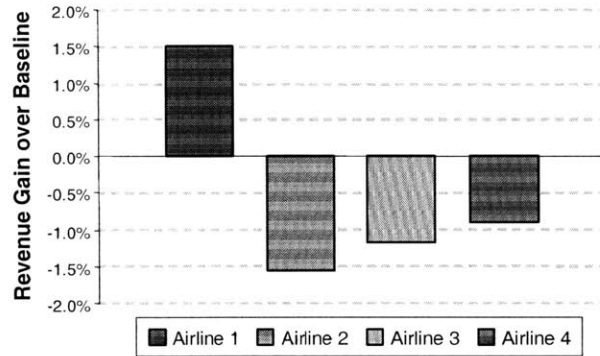


Figure 5-8 : Incremental Benefit of DAVN—Network R

Looking at performance measures, we can see in Table 5-8 that Airline 1’s load increases with its revenue, and that its yield decreases somewhat. This indicates that it is carrying more traffic, but it is carrying more connections—which by definition tend to be lower yield traffic. Revenues, loads and yields all decrease for Airline 2 & Airline 3, indicating that Airline 1 is improving at their expense. Airline 4 sees lower a load but a slightly higher yield, implying that it is losing some of its cheaper connecting traffic to the other airlines.

Table 5-7 : DAVN Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	DAVN/ D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	DAVN	+1.51%	84.9	0.1311
DAVN	-	86.4	0.1240	DAVN	-1.55%	85.3	0.1237
AT80	-	78.6	0.1236	AT80	-1.17%	78.0	0.1230
DAVN	-	85.9	0.1161	DAVN	-0.90%	85.0	0.1164

From these results, we can see that DAVN’s efficacy is greater in Network R. This is a result of the increased capacity and more diverse competition, and also the higher level of fare product differentiation afforded by the mixed-fare network.

5.2.2 ProBP

In contrast with the situation in Network D6, when Airline 1 implements ProBP in Network R it now sees a revenue gain of 0.65%. The other airlines see losses of 1.12%, 0.91%, and 0.41% (Figure 5-9), reflecting again their relative exposures to Airline 1’s network.

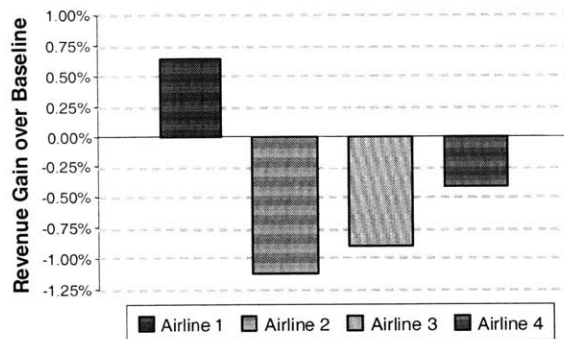


Figure 5-9 : Incremental Benefit of ProBP—Network R

Looking at the performance measures, we can see in Table 5-9 that Airline 1’s load now decreases as revenue increases, and its yield increases somewhat. This indicates that it is carrying a better mix of traffic, and that ProBP is functioning well. Revenues, loads and yields all decrease for Airline 2 & Airline 3 again, indicating that Airline 1 is managing to improve its traffic at their expense. Airline 4 sees a small drop in revenue and load.

Table 5-8 : ProBP Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	ProBP/ D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	ProBP	+0.65%	82.3	0.1340
DAVN	-	86.4	0.1240	DAVN	-1.12%	86.2	0.1230
AT80	-	78.6	0.1236	AT80	-0.91%	78.3	0.1229
DAVN	-	85.9	0.1161	DAVN	-0.41%	85.6	0.1161

However, from these results, we see that ProBP’s performance in Network R is still not comparable to that of DAVN. It appears that ProBP’s reliance on path forecasts and path control are still resulting in over-forecasting, making it vulnerable to fare simplification even in a mixed network such as Network R.

5.2.3 HBP

Returning to the new formulation of HBP, we find that when Airline 1 implements HBP in Network R it sees a revenue gain of 0.44%. Airline 2 & Airline 4 see losses of 0.77% and 0.33% respectively, but now the low-cost carrier Airline 3 sees a revenue gain of 1.75% (Figure 5-10).

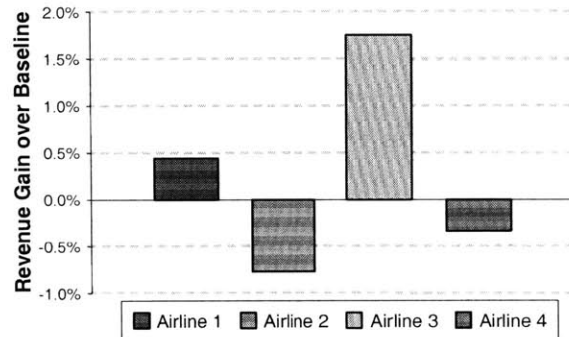


Figure 5-10 : Incremental Benefit of HBP—Network R

Looking at performance measures, we can see in Table 5-10 that Airline 1’s load now increases significantly with its revenue increase. The average load across its network rises by 4.5 points and its yield decreases by over half a cent. The low-cost airline sees an accompanying small decrease in load and a small increase in yield. This would imply that Airline 1 is being less parsimonious in deciding which booking requests to accept, and is allowing more traffic, but not necessarily traffic in higher classes. This gives Airline 3 the opportunity to capture a higher-yield traffic mix—traffic that was denied to it when Airline 1 implemented DAVN or ProBP. The other “traditional” airlines see their loads decrease and their yields increase, indicating that they are carrying a better mix of traffic; nevertheless, their loss in traffic is still sufficient to result in revenue losses.

Table 5-9 : HBP Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	HBP /D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	HBP	+0.44%	87.8	0.1254
DAVN	-	86.4	0.1240	DAVN	-0.77%	85.1	0.1250
AT80	-	78.6	0.1236	AT80	+1.75%	78.3	0.1262
DAVN	-	85.9	0.1161	DAVN	+0.33%	85.2	0.1168

From these results, we can see that HBP's performance in Network R is satisfactory, but not as impressive as its performance in Network D6. It is the weakest of the three methods in Network R. Its strength in Network D6 is here preventing it from being as effective, in that it is perhaps under-protecting for the higher fare classes and accepting too much low yield traffic.

5.2.4 Summary

Considering the performance of all three methods with respect to the E/D/A/D baseline (Figure 5-11 and Table 5-11), we can see that DAVN is clearly the best performer. DAVN's performance here is comparable to its performance in Classic Network D, and it is not experiencing the same feedback problems as was seen in the Network D6 environment. On the other hand, given standard forecasting methods, ProBP's performance clearly still suffers in this less differentiated mixed-fare environment. Lastly, the new HBP is not as effective in Network R as it was in Network D6. This may be because HBP is not as systematic as DAVN in terms of judging the network value of the different booking requests. HBP doesn't over-protect for higher fare classes, and so it allows more overall traffic than DAVN, thus filtering out feedback. However, it is not as effective when there are more path options and more room for subtle trade-offs, for these same reasons—it under-protects for higher fare classes and allows too much traffic in the lower fare classes. And so, to reiterate: HBP under-protects for the higher fare classes given standard pick-up forecasting, just as DAVN and ProBP over-protect for them. Hence, we turn to a different forecasting method.

Table 5-10 : O-D Control Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	DAVN/ D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	DAVN	+1.51%	84.9	0.1311
DAVN	-	86.4	0.1240	DAVN	-1.55%	85.3	0.1237
AT80	-	78.6	0.1236	AT80	-1.17%	78.0	0.1230
DAVN	-	85.9	0.1161	DAVN	-0.90%	85.0	0.1164
ProBP /D/A/D	Revenue Change	ALF	Yield	HBP /D/A/D	Revenue Change	ALF	Yield
ProBP	+0.65%	82.3	0.1340	HBP	+0.44%	87.8	0.1254
DAVN	-1.12%	86.2	0.1230	DAVN	-0.77%	85.1	0.1250
AT80	-0.91%	78.3	0.1229	AT80	+1.75%	78.3	0.1262
DAVN	-0.41%	85.6	0.1161	DAVN	+0.33%	85.2	0.1168

In Sections 5.1 and 5.2, we have seen that fare simplification has implications for the effectiveness of these three O-D revenue management methods. We have seen the problems of spill-in and feedback when fare levels are lowered and compressed, but this was not seen to be a problem in all network situations – it was only a problem in intense symmetric competitive situations. More significantly, we have seen the difficulties faced by methods that rely on path and fare class forecasting in environments with less differentiated fare products. Therefore, in the next two sections, the revenue management methods will be tested in conjunction with Hybrid Forecasting (HF), in an attempt to overcome the assumption of independent path and fare class demands.

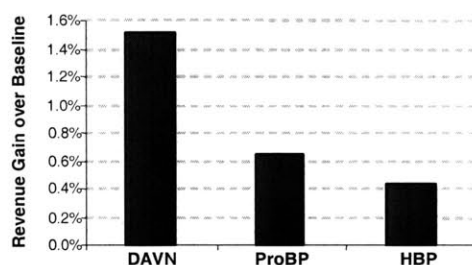


Figure 5-11 : Incremental Benefits of O-D Control—Network R

5.3 Network D6—Hybrid Forecasting

In the previous sections, we compared the performance of the different revenue management methods in the two competitive environments. In each case, the forecasting component of the revenue management system was pick-up forecasting, and we saw the effect on revenue management performance that results from reduced fare product differentiation. The problems associated with all the methods’ assumptions of distinct and independent demands for different path-class products have been emphasized previously. To this end, in this and in the following section, we will investigate the effects of using Hybrid Forecasting (HF) in combination with EMSRb, DAVN, ProBP, and HBP.

As described in Chapter 3, hybrid forecasting is a method of forecasting the demand for fare products based on the assumption that there are two demands: “product oriented demand” and “price oriented demand”. A number of different ways of defining these demands was outlined in Chapter 3. For the purposes of these tests, passengers will be classified as “product oriented” if they purchase a ticket when seats were available in the

next lower class on that same path (“Path Rule”). Passengers will always be defined as “price oriented” if they are observed purchasing the cheapest available option, regardless of whether the next lower class was closed due to advance purchase restriction or do to revenue management system intervention. These options have consistently been found to be the most effective. See Reyes (2006) for more detailed analyses of hybrid forecasting in Network D6.

This section contains a brief examination of the added benefits of hybrid forecasting in Network D6. The FCYM and DAVN results presented in this section are based on results from Reyes (2006). The Network D6 results don’t use estimates of sell-up rates, but rather, for simplicity, they use input sell-up rates (FRAT5 of “C”). For more details on the mechanics of input FRAT5s see Cléaz-Savoyen (2005).

5.3.1 FCYM

As a first comparison, we look at the effect that hybrid forecasting has on leg-based FCYM. In adding HF to the implementation of FCYM, a revenue gain of 2.63% is observed over EMSRb with pick-up forecasting. This gain is larger than any of the incremental gains realized from implementing O-D control using pick-up forecasting.

The hybrid forecaster has the effect of raising Airline 1’s load and its yield, and has a large beneficial effect on the performance of FCYM, helping it to overcome the problems associated with forecasting distinct demands for fare products.

Table 5-11 : FCYM+HF Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
FCYM+ HF	+2.63%	84.3	0.1022	FCYM	-0.41%	82.8	0.0968

5.3.2 DAVN

When Airline 1 implements DAVN with HF it sees further significant gains in revenue above those seen when it implemented DAVN with pick-up forecasting; in this case Airline 2 sees revenue losses. The feedback problem seems to have been averted, and the

“traditional” pattern of revenue gains and losses has been restored. Airline 1 is accepting more traffic, and its load rises by 4.5 points and its yield decreases (Table 5-13). The hybrid forecasting methodology is much better suited to forecasting in this environment. It is not over-protecting for high fare classes, and is more lenient in terms of closing down lower classes. Thus it is more effective at increasing revenues.

The gain to Airline 1 is 4.01% over the FCYM baseline case; an incremental gain of 3.3% over the DAVN/FCYM case is provided by hybrid forecasting.

Table 5-12 : DAVN+HF Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
DAVN	+0.74%	82.8	0.1022	FCYM	+1.15%	83.1	0.0980
DAVN+HF	+4.01%	88.1	0.0991	FCYM	-1.08%	80.9	0.0984

5.3.3 ProBP

When Airline 1 implements ProBP with HF, the “traditional” pattern of revenue gains and losses is restored as it was for DAVN+HF. Similar to the DAVN+HF case, Airline 1’s average load rises by about 4.5 points and its yield decreases (Table 5-14). Although its performance trends are similar to those of DAVN, it is still not as effective in boosting revenues.

The gain to Airline 1 is 3.31% over the FCYM baseline case; an incremental gain of 2.9% over the ProBP/FCYM case is provided by hybrid forecasting.

Table 5-13 : ProBP+HF Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
ProBP	+0.43%	82.5	0.1022	FCYM	+1.05%	83.1	0.0978
ProBP+HF	+3.31%	87.1	0.0995	FCYM	-1.02%	80.7	0.0987

5.3.4 HBP

Finally, when Airline 1 implements HBP with HF, the “traditional” pattern of competitive revenue gains and losses is seen once more. The benefit to Airline 1 of using HBP+HF is almost twice that of using HBP with pick-up forecasting. In addition, adding hybrid forecasting lowers Airline 1’s load factor and raises its yield with respect to the HBP/FCYM case (Table 5-15). This shows that HBP’s strengths in terms of using leg-based forecasting and optimization, and the forecasting advantages provided by HF, combine to produce an effective revenue management combination in the less differentiated fare environment.

The gain to Airline 1 is now 3.71% over the FCYM baseline case; an incremental gain of 1.7% over the HBP/FCYM case is provided by hybrid forecasting.

Table 5-14 : HBP+HF Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
HBP	+1.93%	87.8	0.0974	FCYM	-0.63%	80.9	0.0989
HBP+HF	+3.71%	85.4	0.1019	FCYM	-0.67%	82.4	0.0970

5.3.5 Summary

Considering the benefit of hybrid forecasting to all four revenue management methods with respect to the FCYM baseline (Figure 5-13 and Table 5-16), we can see that DAVN with hybrid forecasting is the best performer in the less differentiated environment. The

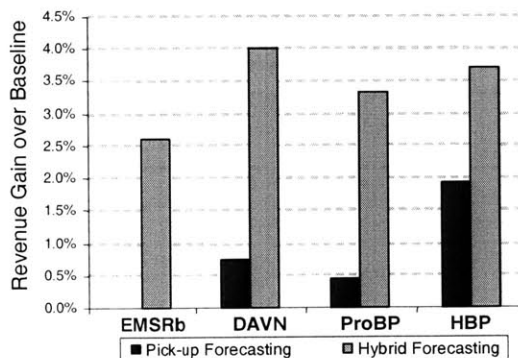


Figure 5-12: Incremental Benefits of Hybrid Forecasting—Network D6

performance of HBP with hybrid forecasting is also very good, showing that HBP is well suited to the less differentiated symmetric network environment in general, with pick-up forecasting or with hybrid forecasting. The performance of ProBP is certainly aided greatly by hybrid forecasting, but it is still not as effective as DAVN or HBP, both of which are easier methods to implement and to operate. The difficulties related to revenue management systems that depend exclusively on path-based processes are still in evidence here.

Table 5-15 : Hybrid Forecasting Changes wrt Baseline—Network D6

Airline 1	Revenue Change	ALF	Yield	Airline 2	Revenue Change	ALF	Yield
FCYM	-	83.6	0.1004	FCYM	-	83.2	0.0967
FCYM+ HF	+2.63%	84.3	0.1022	FCYM	-0.41%	82.8	0.0968
DAVN+ HF	+4.01%	88.1	0.0991	FCYM	-1.08%	80.9	0.0984
ProBP+ HF	+3.31%	87.1	0.0995	FCYM	-1.02%	80.7	0.0987
HBP+H F	+3.71%	85.4	0.1019	FCYM	-0.67%	82.4	0.0970

5.4 Network R—Hybrid Forecasting

We now turn to examine the benefits of hybrid forecasting in Network R. The definitions of price and product oriented demand will be the same as were used when looking at Network D6. In this section however, the hybrid forecaster will make use of PODS' "Inverse Cumulative" method for estimating passenger sell-up rates. This is obviously a more realistic approach than using input sell-up rates. No airline would make use of arbitrarily assigned, assumed rates of sell-up when trying to estimate demand and determine booking limits. Clearly the only realistic solution is to use a method that can estimate passenger price elasticity based on previously seen bookings. Hence the sell-up estimator will be used in Network R. At this stage, we are also discontinuing testing of ProBP, due to its weaker performance with less-differentiated fare environments. Thus, in this section, we will investigate the effects of using hybrid forecasting (HF) in Network R in combination with FCYM, DAVN, and HBP.

5.4.1 FCYM

Once again, as a first comparison we look at the assistance that hybrid forecasting can provide to leg-based revenue management in Network R. In adding HF to the implementation of EMSRb, an additional revenue gain of 1.71% is observed over EMSRb with pick-up forecasting. This gain is once again larger than the incremental benefit realized when implementing DAVN with pick-up forecasting (+1.51%).

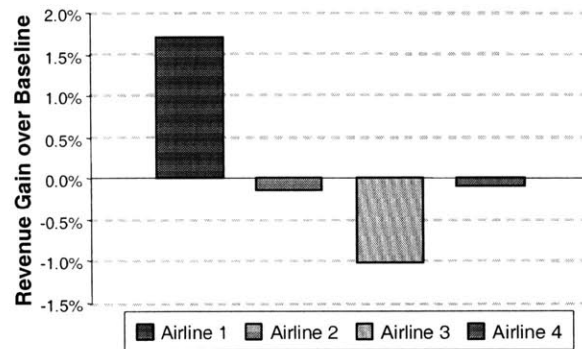


Figure 5-13 : Incremental Benefit of FCYM+HF—Network R

Adding the hybrid forecaster now has the effect of lowering Airline 1’s load and raising its yield, meaning that it is doing a better job of forecasting demand and protecting seats. It also results in revenue drops for each of the competing airlines.

Table 5-16 : FCYM+HF Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	FCYM+HF/ D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	FCYM+HF	+1.71%	81.8	0.1363
DAVN	-	86.4	0.1240	DAVN	-0.14%	86.9	0.1231
AT80	-	78.6	0.1236	AT80	-1.02%	78.5	0.1224
DAVN	-	85.9	0.1161	DAVN	+0.10%	86.4	0.1154

5.4.2 DAVN

When Airline 1 implements DAVN+HF in Network R it again sees significant gains in revenue in comparison to the case where Airline 1 implemented DAVN with pick-up

forecasting. In contrast to the FCYM+HF case, the addition of hybrid forecasting now has the effect of raising Airline 1's load and lowering its yield with respect to the case where DAVN was used with pick-up forecasting (Table 5-18). With hybrid forecasting in use, Airline 1 is not over-protecting for the higher fare classes, allowing it to carry more traffic, thus raising its revenues.

Table 5-17 : DAVN+HF Changes wrt Baseline—Network R

Baseline	Revenue Change	ALF	Yield	DAVN /D/A/D	Revenue Change	ALF	Yield
FCYM	-	83.2	0.1319	DAVN	+1.51%	84.9	0.1311
DAVN	-	86.4	0.1240	DAVN	-1.55%	85.3	0.1237
AT80	-	78.6	0.1236	AT80	-1.17%	78.0	0.1230
DAVN	-	85.9	0.1161	DAVN	-0.90%	85.0	0.1164
				DAVN+HF/ D/A/D	Revenue Change	ALF	Yield
				DAVN+HF	+3.45%	86.8	0.1308
				DAVN	-1.74%	85.4	0.1233
				AT80	-1.05%	78.0	0.1232
				DAVN	-1.14%	85.4	0.1155

The gain to Airline 1 is 3.45% over the FCYM/D/A/D baseline case; an incremental gain of 1.91% over the DAVN/D/A/D case is provided by hybrid forecasting.

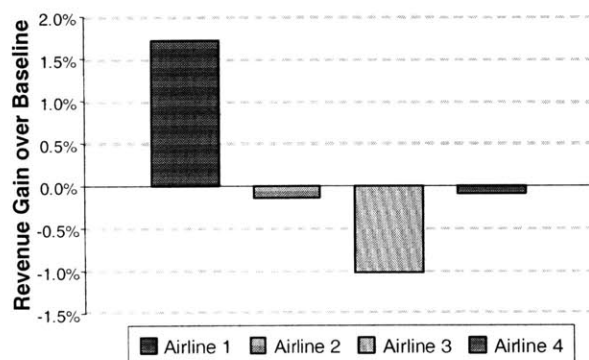


Figure 5-14 : Incremental Benefit of DAVN+HF—Network R

5.4.3 HBP

Finally, when Airline 1 implements HBP+HF in Network R it too sees a significant boost to its revenue. The revenue gain is associated with a drop in load and a large increase in yield. Recall that when Airline 1 implemented HBP with pick-up forecasting, it saw a 4.5 point load increase, and a 6.5-cent yield decrease. Now, when it implements HBP with hybrid forecasting, the load increase is reversed, but its yield increases by a cent (see Table 5-19). That is, Airline 1's load is the same as for the FCYM with pick-up forecasting case, but its yield is almost half a cent higher. This indicates that it is now forecasting demand much more effectively than it did when using pick-up forecasting, allowing it to improve its revenue by improving its yield. As for the Network D6 case, this result implies that the combination of HBP with hybrid forecasting is also an effective revenue management combination in the mixed fare environment.

Table 5-18 : HBP+HF Changes wrt Baseline—Network R

Baseline	ALF	Yield	HBP/D/A/D	Revenue Change	ALF	Yield
FCYM	83.2	0.1319	HBP	+0.44%	87.8	0.1254
DAVN	86.4	0.1240	DAVN	-0.77%	85.1	0.1250
AT80	78.6	0.1236	AT80	+1.75%	78.3	0.1262
DAVN	85.9	0.1161	DAVN	+0.33%	85.2	0.1168
			HBP+HF/D/A/D	Revenue Change	ALF	Yield
			HBP+HF	+2.98%	83.2	0.1357
			DAVN	-1.46%	86.3	0.1223
			AT80	-0.59%	78.2	0.1234
			DAVN	-0.44%	86.7	0.1146

The gain to Airline 1 is 2.98% over the FCYM/D/A/D baseline case; an incremental gain of 2.53% over the HBP/D/A/D case is provided by hybrid forecasting.

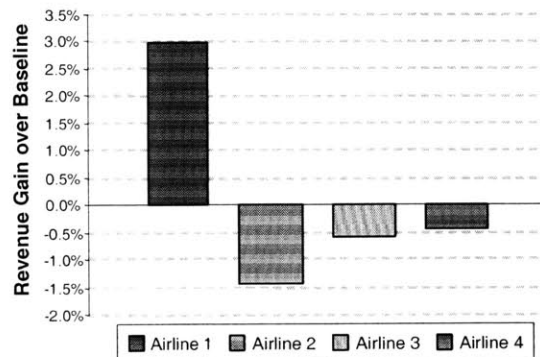


Figure 5-15 : Incremental Benefit of HBP+HF—Network R

5.4.4 Summary

Considering the benefit of hybrid forecasting to these three revenue management methods with respect to the FCYM baseline (Figure 5-17), we can see that DAVN with hybrid forecasting is still the best performer in this mixed-fare environment. The performance of HBP with hybrid forecasting is once again also very good, confirming the previous observations that HBP is well suited to less-differentiated environments.

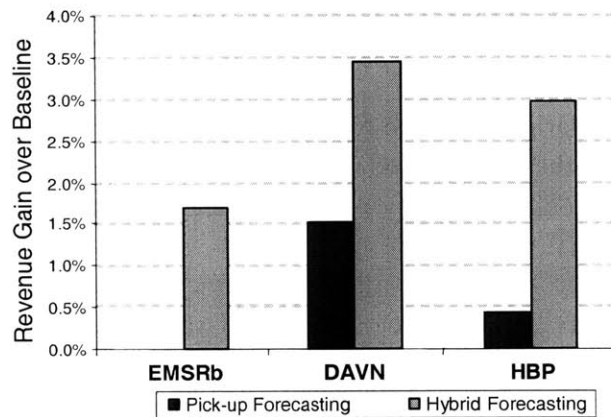


Figure 5-16 : Incremental Benefits of Hybrid Forecasting—Network R

5.5 Summary of Results

In this chapter, I have presented the results obtained from testing four revenue management methods in the two new PODS simulation environments. The methods were tested in conjunction with either pick-up forecasting or hybrid forecasting.

In the sections on pick-up forecasting, we observed some of the problems associated with revenue management in less-differentiated fare environments. Performing forecasting under the assumption of distinct demands for fare class products leads to poor revenue management system performance, and can produce counterintuitive results. We saw that compression and the removal of restrictions affected the methods which rely on path forecasts in particular, with ProBP suffering the most from the removal of restrictions and the consequent spiral down. ProBP's performance was not good in either network environment.

Another consequence of fare compression and the removal of restrictions was seen in Network D6, where passengers who were rejected from one airline were seen to be willing to buy higher fare class products. However, the pick-up forecasting philosophy is not well suited to deal with the sell-up phenomenon, and the gains realized from DAVN and ProBP were overshadowed by the effects of passengers shifting between fare classes.

However, although the performance of DAVN was hurt in an environment in which there is uniform fare simplification across all markets, it was able to recover its effectiveness in the higher capacity mixed network. ProBP, however, saw its performance diminished in both situations.

HBP was seen to be particularly well suited to high demand, highly symmetrical environments, and less effective in the larger mixed network. It does not over-forecast for the higher fare classes, and therefore allows more traffic and earns higher revenues.

This was an advantage to it in the high demand environment of Network D6 where the other methods were over-protecting and closing lower fare classes; but it was less of an advantage in the larger Network R, where it accepted too much traffic and did not protect enough for the high fare classes.

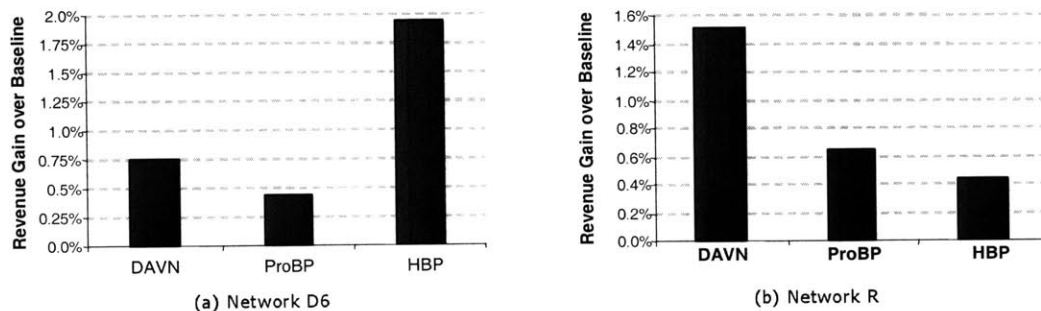


Figure 5-17 : Incremental Benefits of O-D Control

In the sections dealing with hybrid forecasting, we found that the hybrid forecasting philosophy is much better suited to these less-restricted fare environments. It doesn't inflate the demand forecasts for the expensive fare products, and as such it doesn't close down the lower fare classes as aggressively, allowing more traffic to be carried.

DAVN's performance was once again best overall, but that HBP's performance was also aided greatly by the hybrid forecasting methodology. When Airline 1 used DAVN with

pick-up forecasting, its loads were too low, and hybrid forecasting was seen to raise its loads and lower its yields by allowing more bookings. Conversely, when Airline 1 used HBP with pick-up forecasting, its loads were too high, and hybrid forecasting was seen to lower its loads and raise its yields by to be more selective about the traffic it carries, as a result of its better forecasts.

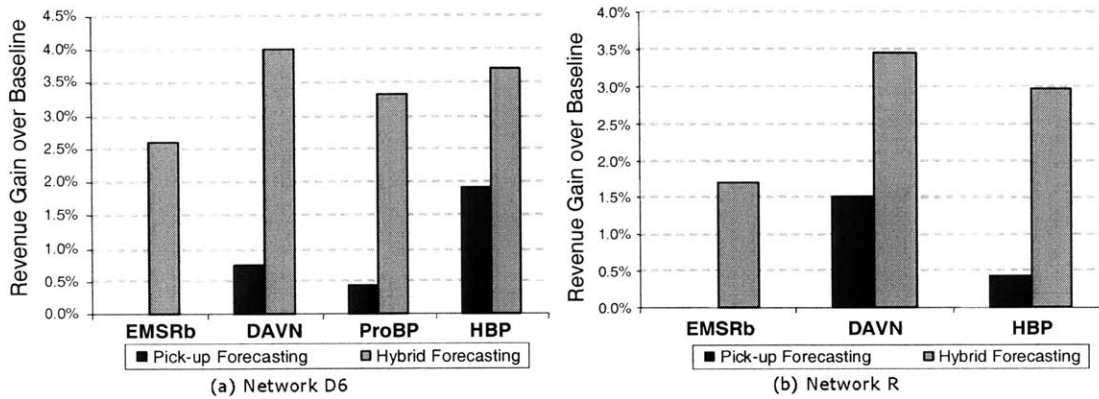


Figure 5-18 : Incremental Benefits of O-D Control—Hybrid Forecasting

From these results, we could conclude that HBP is a method of great interest, and qualifies for further investigation. Its use of leg-based forecasts makes it suitable for use in less differentiated environments, where the difficulties of path forecasting are exacerbated by fare simplification. Its performance is broadly comparable to that of DAVN and it has the significant advantage of being cheaper and simpler to implement. The combination of HBP and hybrid forecasting appears to have the potential to be a very effective and practical revenue management system for use in less-differentiated and mixed fare environments.

6 Conclusion

This thesis has been concerned with the possibilities and potential for airline revenue management in contemporary less-differentiated fare environments. I have described briefly the revenue situation facing the airline industry today. In the current environment we have seen falling yields as a consequence of compressed fares and the removal of restrictions on many fare products. I have also outlined some of the previous work on revenue management in general, and in particular the growing field concerned with revenue management in less-restricted environments.

After describing some seat allocation optimization and demand forecasting methodologies, I developed two new simulation environments for use in the Passenger Origin-Destination Simulator (PODS). These environments have been designed particularly for use in investigating the possibilities for revenue management in less-differentiated environments. They include widespread fare compression and fare product simplification on the one hand (Network D6); and incorporating a traditional fare structure with a less-differentiated fare structure on the other hand (Network R). Network D6 is appropriate for use in modelling intense symmetric competitive environments, whereas Network R is more elaborate and can be used to model more complex competitive environments.

Four revenue management methods and two forecasting methodologies were tested in these two environments. These are FCYM, DAVN, ProBP, and HBP; and Pick-up and Hybrid Forecasting. From these results we saw the difficulties of using traditional revenue management in less-differentiated or mixed fare environments. The assumption of independent and distinct demand for path/fare-class products is not consistent with passenger behaviour in the current environment. Passengers are not purchasing fare products, but rather they are purchasing journeys, and without compelling differentiation between fare products they are flexible in terms of which fare class designation they will purchase. Methods which forecast future demand based on previously observed sales in particular fare classes do not and can not perform well in the new environments. They need to be modified in order to account for passenger flexibility before they will be able to function well in the changed fare structures.

This inconsistency between behaviour and revenue management systems' assumptions has consequences for their ability to set protection levels effectively. On the one hand, spiral down leads to low fare class forecasts that are too large; at the same time, the

observed histories for many path/fare class combinations can be very small but highly variable, and this results in forecasts that are too large. Therefore, a forecasting methodology that is not based on historically observed demand in particular fare classes is particularly necessary in the new fare environment.

If we consider that there are only two different demands for air travel—demand for flexibility and demand for the lowest price—we come to the hybrid forecasting methodology which forecasts demand for “product oriented” and “price oriented” travel. Results from testing this methodology were very encouraging in both the less-differentiated environment and the mixed-fare environment.

We saw that pick-up forecasting created problems by either over- or under-forecasting demand. In some situations this resulted in airlines accepting too much low yield traffic, or conversely rejecting too much traffic overall. These problems were largely alleviated when the hybrid forecaster was employed. We saw that DAVN was less aggressive in closing down fare classes and could thus carry more traffic and earn higher revenues. We also saw that HBP was able to reject more low-yield traffic and earn higher revenues by improving its load. In all cases, the hybrid forecaster produced better forecasts and resulting in improved revenue for the airlines that implemented it.

Future Research

Following on from the results found in this thesis, some further research directions can be seen. Firstly, there is much further work to be done with the hybrid forecasting methodology. Since the hybrid forecaster produces separate estimates for “price” and “product” oriented demand, it would be instructive to look more closely at these forecasts and to see how they vary in different markets.

For example, we could look in more detail at Network R to see how the forecasts vary across different types of markets. The network can be divided into markets where Airline 1 is in competition with Airline 2 only and they both offer a traditional fare structure; and markets where Airline 1 is in direct competition with Airline 3 and they both offer simplified fare structures. It would be interesting to compare the performance of hybrid forecasting in these different types of markets.

Methods of estimating passenger willingness-to-pay and estimating the probabilities of passenger sell-up are also critically important avenues of research. The work in this thesis has been based either on the assumption of independent demand for different fare prod-

ucts, or on the assumption that there is one demand for the most flexible product and another demand for cheapest product. Further investigation of passenger choice behaviour and the changing price elasticity of demand in the new less-differentiated fare environments will be fundamental to the development of more sophisticated ways of determining passenger willingness-to-pay.

Another direction could be to investigate the feedback phenomenon observed in Network D6—namely, under what circumstances do the feedback effects dominate the benefits of O-D revenue management. The relationship between demand and capacity might be a factor in this phenomenon.

Finally, the relative benefits of DAVN and the new HBP are also of interest. The new HBP has shown itself to be a good performer across a wide variety of situations. It is still dominated by DAVN, but its simplicity makes it a very attractive method for further research.

7 Bibliography

Barnhart, C., Belobaba, P., Odoni, A.R., 2003, 'Applications of Operations Research in the Air Transportation Industry', *Transportation Science*, 37(4): 368-391

Belobaba, P.P., 1987, 'Air travel demand and airline seat inventory management', PhD Thesis, Massachusetts Institute of Technology

Belobaba, P.P. 1992, 'Optimal vs. heuristic methods for nested seat allocations', in *Proceedings of the AGIFORS Reservations and Yield Management Study Group*, Brussels, Belgium 1992, pp. 28-53

Belobaba, P.P. 1998a, 'Airline Differential Pricing for Effective Yield Management', in G.F. Butler & J. Peel (eds.), *Handbook of Airline Marketing*, Aviation Week Group (McGraw-Hill), pp. 349-361

Belobaba, P.P. 1998b, 'The Evolution of Airline Yield Management: Fare Class to Origin-Destination Seat Inventory Control', in G.F. Butler & J. Peel (eds.), *Handbook of Airline Marketing*, Aviation Week Group (McGraw-Hill), pp. 385-302

Belobaba, P.P. 2002, 'Airline Network Revenue Management: Recent Developments and State of the Practice', in D. Jenkins (ed.), *Handbook of Airline Economics*, 2nd ed., Aviation Week Group (McGraw-Hill), pp. 141-156

Belobaba, P.P. & Dar, M., 2005, 'RM Performance Under Simplified Fare Structures', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, May 2005, Cape Town, South Africa

Belobaba, P.P. & Hopperstad, C.A., 2004, 'Alternative RM Algorithms for Unrestricted Fare Structures', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, March 2004, Auckland, New Zealand

Belobaba, P.P. & Wilson, J.L., 1997, 'Impacts of yield management in competitive airline markets', *Journal of Air Transport Management*, 3(1): 3-9

Bhatia, A.V. & Parekh, S.C., 1973, 'Optimal allocation of seats by fare', AGIFORS Reservations and Yield Management Study Group, 1973

Bohutinsky, C.H., 1990, 'The sell up potential of airline demand', Master's Thesis, Massachusetts Institute of Technology

Botimer, T.C., 1994, 'Airline pricing and fare product differentiation', PhD Thesis, Massachusetts Institute of Technology

Boyd, E.A. & 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

Buhr, J., 1982, 'Optimal Sales Limits for 2-sector flights', in AGIFORS Twenty Second Annual Symposium Proceedings, October 1982, volume 22, pp. 291-304

Cakulev, I., Kallesen, R., Narayanan, P.R., 2005, 'Modeling in a competitive environment', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, May 2005, Cape Town, South Africa

Carrier, E., 2003, 'Modeling airline passenger choice: passenger preference for schedule in the passenger origin-destination simulator (PODS)', Master's Thesis, Massachusetts Institute of Technology

Cléaz-Savoyen, R.L., 2005, 'Airline revenue management methods for less restricted fare structures', Master's Thesis, Massachusetts Institute of Technology

Cooper, W.L., Homem-de-Mello, T., Kleywegt, A., 2004, 'Models of the Spiral-Down Effect in Revenue Management'

Curry, R.E., 1990, 'Optimum Seat Allocation with Fare Classes Nested by Origins and Destinations', *Transportation Science*, 24: 193-204

Cusano, A.J., 2003, 'Airline revenue management under alternative fare structures', Master's Thesis, Massachusetts Institute of Technology

Darot, J.F.J., 2001, 'Revenue management for airline alliances: passenger origin-destination simulation analysis', Master's Thesis, Massachusetts Institute of Technology

Doganis, R. 2005, *The Airline Business*, 2nd ed., London: Routledge

Donnelly, S., James, A., Binnion C., 2004, 'bmi's response to the changing European airline marketplace', *Journal of Revenue and Pricing Management*, 3(1): 10-17

Dorinson, D.M., 2004, 'The evolution of airline distribution channels and their effects on revenue management performance', Master's Thesis, Massachusetts Institute of Technology

Fiig, T., Isler, K., Hopperstad, C., Cléaz-Savoyen, R.L., 2005, 'DAVN-MR: A Unified Theory of O&D Optimization in a mixed Network with restricted and unrestricted fare products', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, May 2005, Cape Town, South Africa

Forsyth, P., 2003, 'Low-cost carriers in Australia: experiences and impacts', *Journal of Air Transport Management*, 9(5): 277-284

Franke, M., 2004, 'Competition between network carriers and low-cost carriers-retreat battle or breakthrough to a new level of efficiency?', *Journal of Air Transport Management*, 10(1): 15-21

Geslin, C., 2006, 'Pricing and Competition in US Airline Markets: Changes in Air Travel Demand Since 2000', Master's Thesis, Massachusetts Institute of Technology

Glover, F., Glover, R., Lorenzo, J., McMillan, C., 1982, 'The passenger mix problem in the scheduled airlines', *Interfaces*, 12: 73-79

Gorin, T.O., 2000, 'Airline revenue management: sell-up and forecasting algorithms', Master's Thesis, Massachusetts Institute of Technology

Gorin, T.O., 2004, 'Assessing low-fare entry in airline markets: impacts of revenue management and network flows', PhD Thesis, Massachusetts Institute of Technology

Holloway, S. 2003, *Straight and Level: Practical Airline Economics*, 2nd ed., Ashgate: Aldershot

- Lawton, T.C. 2002, *Cleared For Take-Off*, Ashgate: Aldershot
- Lee, A.Y.H., 1998, 'Investigation of competitive impacts of origin-destination control using PODS', Master's Thesis, Massachusetts Institute of Technology
- Lee, S., 2000, 'Modeling passenger disutilities in airline revenue management simulation', Master's Thesis, Massachusetts Institute of Technology
- Littlewood, K., 1972, 'Forecasting and Control of Passenger Bookings', in AGIFORS Symposium Proceedings, October 1972, volume 12, pp. 95-117
- Mak, C.Y., 1992, 'Revenue impacts of airline yield management', Master's Thesis, Massachusetts Institute of Technology
- Mason, K.J., 2005, 'Observations of fundamental changes in the demand for aviation services', *Journal of Air Transport Management*, 11(1): 19-25
- McGill, J.I. & Van Ryzin, G.J., 1999, 'Revenue Management: Research Overview and Prospects', *Transportation Science*, 33(2): 233-256
- Morrell, P., 2005, 'Airlines within airlines: an analysis of US network airline responses to Low Cost Carriers', *Journal of Air Transport Management*, 11(5): 303-312
- Ratliff, R., 2005, 'Demand Distributions for Revenue Management', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, May 2005, Cape Town, South Africa
- Ratliff, R. & Vinod, B., 2005, 'Airline Pricing and Revenue Management: A future outlook', *Journal of Revenue and Pricing Management*, 4(3): 302-307
- Reyes, M.H., 2006, 'Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures', Master's Thesis, Massachusetts Institute of Technology
- Richter, H., 1982, 'The Differential Revenue Methods to Determine Optimal Seat Allotments by Fare Type', AGIFORS Twenty Second Annual Symposium Proceedings, October 1982, volume 22, pp. 339-362

- Rothstein, M., 1985, 'OR and the Airline Overbooking Problem', *Operations Research*, 33(2): 237-248
- Saranathan, K. & Zhao, W., 2005, 'Yield Management in the New Fare Environment', Presentation to AGIFORS Reservations and Yield Management Study Group Meeting, May 2005, Cape Town, South Africa
- Shaw, S. 2004, *Airline Marketing and Management*, 5th ed., Ashgate: Aldershot
- Skwarek, D.K., 1997, 'Competitive impacts of yield management system components: forecasting and sell-up models', Master's Thesis, Massachusetts Institute of Technology
- Smith, B.C. & Penn, C.W., 1988, 'Analysis of Alternate Origin-Destination Control Strategies', in AGIFORS Twenty Eighth Annual Symposium Proceedings, October 1988, volume 28, pp. 123-144
- Tan, J.C., 1994, 'Leg-based heuristic methods to network seat inventory control', Master's Thesis, Massachusetts Institute of Technology
- Vanhaverbeke, T., 2006, 'Revenue Management Based on Dynamic Programming in Unrestricted and Simplified Fare Structures', Master's Thesis, Massachusetts Institute of Technology
- Wang, K., 1983, 'Optimum Seat Allocation for Multi-Leg Flights with Multiple Fare Types', in AGIFORS Twenty Third Annual Symposium Proceedings, October 1983, volume 22, pp. 225-246
- Wei, Y.J., 1997, 'Airline O-D control using network displacement concepts', Master's Thesis, Massachusetts Institute of Technology
- Wickham, R.R., 1995, 'Evaluation of forecasting techniques for short-term demand of air transportation', Master's Thesis, Massachusetts Institute of Technology
- Williamson, E.L., 1988, 'Comparison of optimization techniques for origin-destination seat inventory control', Master's Thesis, Massachusetts Institute of Technology
- Williamson, E.L., 1992, 'Airline network seat inventory control: methodologies and revenue impacts', PhD Thesis, Massachusetts Institute of Technology

Wilson, J.L., 1995, 'The Value of Revenue Management Innovation in a Competitive Airline Industry', Master's Thesis, Massachusetts Institute of Technology

Windle, R. & Dresner, M., 1999, 'Competitive Responses to Low Cost Carrier Entry', *Transportation Research Part E*, 35: 59-75

Wysong, R., 1988, 'A Simplified Method for Including Network Effects in Capacity Control', in AGIFORS Twenty Eighth Annual Symposium Proceedings, October 1988, volume 28, pp. 113-121

Zerbib, G.J-J., 2006, 'Expansion of Point-to-Point Routes by Low-Cost Carriers in Hub Networks: Traffic and Revenue Impacts', Master's Thesis, Massachusetts Institute of Technology

Zickus, J.S., 1998, 'Forecasting for airline network revenue management: revenue and competitive impacts', Master's Thesis, Massachusetts Institute of Technology