Airline Revenue Management Methods for Less Restricted Fare Structures

By

Richard L. Cléaz-Savoyen

Ingénieur des Arts et Manufactures (ECP 2003)

Submitted to the Department of Aeronautics and Astronautics
in partial Fulfillment of the Requirements for the Degree of Master of Science
in Aeronautics and Astronautics

At the

Massachusetts Institute of Technology

May 2005 [June 2005]

1st Copy

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
JUN 2 3 2005
LIBRARIES

© 2005 Massachusetts Institute of Technology
All Rights Reserved

Signature of Author ....

International Center for Air Transportation
Department of Aeronautics and Astronautics
May 10, 2005

Certified by ....

Dr Peter P. Belobaba
Principal Research Scientist
Department of Aeronautics and Astronautics
Thesis Supervisor

Accepted by ....

Jaime Peraire
Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Student

AERO

Airline Revenue Management Methods for Less Restricted Fare Structures

By

Richard L. Cléaz-Savoyen

Submitted to the Department of Aeronautics and Astronautics
on May 11, 2005, in partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

ABSTRACT

Traditional Revenue Management systems were developed to maximize airlines' revenues in restricted fare product environments, based on the assumption of independence of demand by fare class. With the rapid emergence of low-cost carriers in various parts of the world, the pricing environments have changed. The network carriers have had to deal with undifferentiated fare structures, for which demand for each fare class can no longer be assumed to be independent, given that customers systematically buy the lowest fare available in the absence of distinctions between fare products. In these fare structures, traditional Revenue Management systems lead to a "spiral down" of revenues.

The first goal of this thesis is to present two alternatives that address the problem described above, allowing airlines to partially recover from the decrease in revenue which occurs when an airline removes the restrictions in its fare structure. These alternatives are designed for implementation into O&D-Control – or network Revenue Management – methods. The alternative methods are based on the sell-up probability, which is the probability that a passenger is willing to buy a higher-fare ticket in case his/her request is denied. One of the methods ("Q-Forecasting") modifies the forecaster by estimating the demand based on the probability of sell-up, while the other ("Fare Adjustment") acts at the booking limit optimizer level, accounting again for the sell-up potential. We focus in this thesis on explaining the processes and mechanisms involved in these two methods, how they are linked and complement each other, but also on their performances based on a simulator which allows us to observe the impact of each method under various characteristics of the booking process.

Thesis Advisor: Dr Peter P. Belobaba
Title: Principal Research Scientist, Department of Aeronautics and Astronautics

ACKNOWLEDGEMENTS

My first acknowledgements are to my thesis and research advisor, Dr Peter Belobaba, in regards to all his help. He taught me the basics and then the less obvious aspects of airline Revenue Management, as well as the tortuous evolutions of the Passenger Origin Destination Simulator on which this thesis is based. He provided me great support and helped me to correct the mistakes I made.

A special "thank you" to Craig Hopperstad with whom I spent several hours on the phone – at decent hours of the day to satisfy both Boston and Seattle times – to understand the new models implemented in each version of PODS. We only met 4 times, at the PODS conferences, but our discussions have always been dense and interesting. Without the simulator Craig developed at Boeing, I could not have done the simulation part of this thesis.

Then, I would like to thank Thomas Fiig, from SAS, who helped me with the Fare Adjustment method he developed in collaboration with Karl Isler from Swissair. Thomas spent several days in January 2005 in Boston and many hours on the phone with me – from Copenhagen – to help me understand the subtleties of this new method which represents half of my thesis in terms of both theoretical descriptions and simulation analyses.

Hereafter I present my acknowledgements to the remaining members of the PODS Consortium, composed of Air France-KLM, Air New Zealand, Continental Airlines, Delta Airlines, LAN, Lufthansa, Northwest Airlines, SAS and Swissair. At each PODS conference, their representatives provided comments on my results and guidance in the future research directions. They introduced in my thesis the point of view from the real world, especially regarding the unrestricted environments and the competition with Low-Cost Carriers.

I would like to thank the amazing student team of the MIT International Center for Air Transportation. Especially: Thomas Gorin, Emmanuel Carrier, Kendell Timmers and Maital Dar – with whom I worked on PODS – who helped me understand the simulator.

Even if Europe and France are maybe a little bit far from Boston, I thank my family for helping me make it to MIT and supporting me every day.

Unless they did not help me to understand Revenue Management, last but not least, I thank all my friends at MIT, but also in Boston, in France, Germany, England, Russia, Austria, China, and Canada for their support during these two great years in New England.

AUTHOR'S BIOGRAPHY

The author comes from Grenoble, a city in the French Alps, near Lyon, where he did his undergraduate studies in science. In 2000, after a competitive exam, he entered the Ecole Centrale Paris, a French engineering "Grande Ecole", where he specialized in Aeronautics during the third year. At the end of this school, he did a 4-month training at Snecma Moteurs, where he worked as manufacturing engineer on a project related to the high pressure compressor blades of the GE90 aircraft engine. Finally, he entered MIT in September 2003 in the Aero/Astro department as a Graduate student, for a Master of Science. He quickly specialized in flight transportation and took the following classes:

Air Traffic Control (Hansman, Clarke)
Airline Industry (Belobaba)
Airline Management (Belobaba)
Logistical and Transportation Planning Methods (Odoni, Barnett, Larson)
Airport System Planning, Design and Management (Odoni, DeNeufville)
Air Transportation System Architecting (Hansman, Clarke, Murman)
Applied Probabilities (Tsitsiklis)
Airline Schedule Planning (Barnhart) – as listener

While at MIT, the author began in January 2004 working as Research Assistant with Dr Peter Belobaba on the Passenger Origin Destination Simulator (PODS) Project. For this research, he attended the PODS conferences in Amsterdam (May 2004), Minneapolis-St Paul (September 2004), Boston (January 2005) and Copenhagen (May 2005). Rapidly, he worked on topics related to unrestricted environments which would become the topic of this thesis due in May 2005. These two years at MIT will hopefully end with the AGIFORS conference in Cape Town (South Africa), a few days before Commencement.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AT	Adaptive (or Accordion) Threshold
AP	Advance Purchase
BKD	Booked
BKG	Booking(s)
BK	Booking Limit
BTC	Bookings-To-Come
BP	Bid-Price
CAP	Capacity
DAVN	Displacement Adjusted Virtual Nesting
DN	"DP New"
DP	Dynamic Programming
DWM	Decision Window Model (by Boeing)
EMSR	Expected Marginal Seat Revenue
FA	Fare Adjustment
FCFS	First Come First Served
FCYM	Fare Class Yield Management
FT	Fixed Threshold
Frat5	Fare Ratio to Q at which 50% of the passengers accept to buy-up
GVN	Greedy Virtual Nesting
HBP	Heuristic Bid-Price
KI	Karl Isler
LCC	Low-Cost Carrier
LFA	Low-Fare Airline
LP	Linear Program
MR	Marginal Revenue
NEM	New Entrant Market(s)
OD	Origin-Destination
PAX	Passengers
PODS	Passenger Origin-Destination Simulator
Pro-BP	Probabilistic Bid-Price
PUMA	Pick-Up Moving Average
QF	Q-Forecasting
Ri	Restriction number i (stronger as i is smaller)
RM	Revenue Management
RMS	Revenue Management System(s)
SIC	Seat Inventory Control
TR	Total Revenue
TF	Time-Frame
WTP	Willingness-To-Pay
YM	Yield Management

1. INTRODUCTION

The goal of this thesis is to present and evaluate the performance of recently developed methods for adapting traditional Revenue Management algorithms to less restricted fare structures. Since the end of the 1970's, airlines across the world have developed more and more sophisticated Revenue Management methods to optimize the seat inventory control and therefore maximize their revenue. These airlines based their Revenue Management algorithms on strong assumptions, such as segmentation by price differentiation and independence of demand and behavior by passenger types. The segmentation was achieved by selling the same seat at different fares depending on associated advance purchase requirements and restrictions – Saturday night stay, change fees and non-refundability of ticket.

These assumptions are now more and more violated, because of the change in consumer behavior, the emergence of the low-cost carriers and the dramatic drop in demand which was seen after the events of September 11, 2001. The traditional carriers are facing new challenges - erosion of differences between passenger types, change in fare products – and are looking for new methods in order to adapt their algorithms to these new demand patterns and new competitors. We will simulate in this thesis the possible revenue gains that can be achieved by implementing two different Revenue Management alternatives in less restricted fare structures.

The simulation will use the Passenger Origin-Destination Simulator (PODS) developed at the Boeing Company by Hopperstad to simulate as accurately as possible the booking process of an airline in a competitive environment, with various levels of control over the passenger choice model, the environment and the airline Revenue Management methods settings.

1.1. Revenue Management

Revenue Management was born in the 1970s and was at first known as Yield Management (YM). Yield refers to the average fare per passenger mile. Therefore, yield management methods were logically intended to capture as many high yield passengers as possible.

Before the deregulation of the US airline industry in 1978, airline companies had no control over fares, which were determined by the Civil Aeronautics Board (CAB) on a per-mile basis. The board allowed for fare increases to compensate for operating costs and in case of losses. This led to relatively high fares, and a good level of service for those who could afford it, but revenues and loads remained low.

Starting from deregulation, the airline companies entered a competitive environment, and yield management became more and more crucial. The goal for each of them was (and still is) to increase revenue and decrease costs, of course. The airline companies compete on an OD basis, meaning that passengers choose between airlines and products on a city-pair, composed of an Origin and a Destination, generally independent from the itinerary flown in between. Consequently, fares no longer depend on distance, but vary with respect to the OD market. In order to maximize revenues, the theoretical optimal strategy for airlines is to charge passengers at their maximum Willingness-To-Pay (WTP). In practice, this is difficult to achieve, that is why revenue management was developed, starting from the differential pricing and overbooking strategies initially introduced before deregulation.

Yield Management has been developed over the last 30 years to achieve these goals and has reached high levels of complexity to deal with complicated fare structures established by the pricing departments in each company. The aim of Yield Management remained to attract as many high yield passengers as possible while filling the aircraft and avoiding flying empty seats.

The segmentation was based on restrictions (or fare fences) such as advance purchase requirements, Saturday night stay, round trip purchase requirement, and non or partial refundability. But the complication lies in the arrival process: high yield passengers book late, whereas leisure passengers (lower willingness-to-pay) book early; Revenue Management has therefore to find a means to separate both behaviors. In order to satisfy the two goals described above, an accurate forecast of high yield demand is needed. Then, based on these forecasts, Revenue Management practices are designed to optimize the allocation of seats to a specific fare class with its associated restrictions.

Today, the level of complexity is even higher, and many different seat allocation algorithms have been developed, either flight-leg based or OD-based (most recently developed). We will describe these methods in more detail in the next chapters. Demand forecasting and seat allocation optimization, the two main aspects of Revenue Management, rely strongly on mathematical formulas and modeling approaches. These modeling approaches are changing and this thesis will deal with changes in both aspects of Revenue Management.

1.2. Changes in Airline Industry and Market Structure

Along with the increasing complexity of Revenue Management Methods, the emergence of new air service providers has changed the industry. Unlike the legacy carriers, the low-fare carriers have very low costs although they provide similar services at lower fares. And their strategy appears to be quite successful at this point, as they were among the only companies of the US airline industry to achieve a positive profit at the beginning of 2005 (JetBlue, Southwest and America West).

The most visible characteristic of the low-cost carriers compared to the legacy carriers is their simplified fare structure, with very few restrictions or advance purchase requirements. This means that they do not segment demand as much as their competitors. From the consumer point of view, there is a lower number of products offered by the low-cost carriers, but their prices change as we approach the departure date, whereas in the case of a legacy carrier, the consumer has to choose between different fare products (with distinct restrictions and advance purchase requirements for each fare product).

In the markets where the legacy carriers compete with low-cost carriers, the former have no other choice than matching their competitors to avoid losing too much revenue. They dramatically change their fare products, which means compression of fare ratios and total or partial removal of restrictions and advance purchase requirements. In these new kinds of environments, where the fundamental assumption of traditional Revenue Management – independence of demand by fare class – is violated since there is no longer the possibility to segment the demand, the traditional Revenue Management methods do not perform as well as they used to in restricted environements. As we will see, without any restrictions or advance purchase requirements, the traditional Revenue Management algorithms lead to a "spiral down" effect, and revenues drop very fast. Given the quick expansion of the low-cost carriers, there is today a pressing need for developing new algorithms which can provide an optimal seat inventory control even in these cases.

1.3. Recently Developed New Approaches

From this point forward, we will focus on the cases where the fences – restrictions and advance purchase requirements – have been removed in a series of markets for a given network carrier. We will deal with unrestricted fare product structures, without restrictions or advance purchase requirements.

1.3.1. Forecasting and Sell-up

Forecasting is based on historical data: it is the estimation by fare class and by flight of the unconstrained demand, obtained from the record of the booking process of previous "equivalent" flights – previous Fridays for instance if we try to forecast demand for a future Friday. Forecasts should be as accurate as possible, since it is a crucial step in the process of Revenue Management. In the unrestricted fare product structure, since every passenger is going to buy the lowest fare available (given that there is no distinction between products anymore), it becomes impossible to estimate, based on historical data, a demand for each fare class using traditional methods. When nothing allows us to distinguish between different fare classes, customers indeed purchase the lowest available fare and demand for highest classes does not materialize. And if it is not observed in the historical database, no demand can be forecast for highest fare classes in the future.

The new methods we will present later will deal, in terms of forecasting, with the probability of sell-up, which is the probability that a passenger is willing to buy a higher fare ticket for the same flight in case he (or she) is denied booking for the requested fare product. Accounting for this probability in the forecast and in the optimization process will allow more seats to be protected for higher yield passengers and force them to sell-up when they will be denied booking for their first request. The challenge is to know how to use these sell-up probabilities accurately, and even more difficult, how to estimate them dynamically during the forecast process (they depend on the flight and cannot be optimal if set arbitrarily) in order to include them in the Revenue Management Systems (RMS).

1.3.2. New Methods: "Q-Forecasting" and "Fare Adjustment"

In order to address the problems mentioned above, several developments have occurred in the field of revenue management algorithms for less restricted fare structures, by **Belobaba/Hopperstad**[1] on one hand, and by **Fiig/Isler**[2] on the other hand. Both of these new approaches, respectively called "Q-Forecasting" (QF) and "Fare Adjustment" (FA) methods, have been implemented and tested in the Passenger Origin-Destination Simulator (PODS). Each of them makes use of the probability that a passenger is willing to buy-up at the time of booking, but each deals with a different aspect of the problem. The first addresses the problem of forecasting a less segmented demand in order to feed a traditional optimizer. The second deals with adjusting the fares to incorporate the potential of sell-up into the seat inventory control optimizer.

1.4. Goal of the Thesis

Today, the vast majority of airline companies use leg-based forecasts and optimizers. Nevertheless, there is a slow but inevitable shift towards OD control methods which are proven to increase significantly revenues although expensive to implement. Moreover, many legacy airlines struggle with low-cost carriers in more and more markets. They need to introduce some new features into their Revenue Management Systems in order to account for the presence of low-cost carriers and compete with these airlines in a given set of markets, while keeping traditional fare structures and methods in other parts of their OD sets. Previous studies - **Cooper and Homen-de-Mello**[3] - have shown the impact of the "do nothing" strategy in these markets, and have explained the negative impact of traditional forecasting and optimization when the "old" assumptions are violated.

Therefore, in this thesis, we will present and explain the new alternatives, the way they are linked with the remainder of the Revenue Management (RM) process and the seat inventory optimization, mainly in the context of OD control methods. It is the goal of the thesis to present the revenue benefits obtained from the use of both methods, initially with inputs set arbitrarily. The next step presented here, is a way to estimate the required inputs during the forecasting process, in order to make the alternative methods autonomous and dynamically adaptive, which is a required feature given the level of automation reached in the Revenue Management System (RMS) to which they will be added.

All performance evaluations will be done using the Passenger Origin-Destination Simulator (PODS) developed by Hopperstad at the Boeing Company. The simulator will be described in detail in Chapter 3.

[1] BELOBABA and HOPPERSTAD, *Q investigations – Algorithms for Unrestricted Fare Classes*, PODS presentation, Amsterdam, 2004
[2] ISLER and FIIG, SAS O&D Low Cost Project, PODS presentation, Minneapolis-St Paul, 2004
[3] COOPER and HOMEN-DE-MELLO, *Models for the Spiral-down effect in Revenue Management*, Dec 2004

1.5. Structure of the Thesis

The thesis consists of 3 main parts: The Literature Review, The Approach to the Problem, and The Analysis of the Simulations.

In Chapter 2, the literature review will present more extensively the history of Revenue Management starting from the 1978 US deregulation, going through the increasingly sophisticated algorithms, and the revenue benefits each major improvement achieved with respect to previously used methods. Then, we will introduce some more detailed analysis done on the change in the airline industry due to the emergence of the low-cost carriers in the United States, but also in other parts of the world. We will discuss the "hybrid" legs in which different fare product structures co-exist, and the "spiral down" effect to which the traditional RMSs lead in unrestricted environments, explaining the need for alternative methods dealing with these new kinds of markets.

Chapter 3 is divided into sections related to the Passenger Origin-Destination Simulator (PODS) on one hand and sections related to the new RM methods on the other hand. The description of PODS is more an overview than a detailed presentation, given that the simulator has been extensively described in previous theses. We will present the fundamental processes, their interactions, the different Revenue Management methods we will refer to during the thesis and the specifics of PODS with respect to the topic of this thesis, such as sell-up inputs and features. Then come the theoretical parts, with at first a more detailed presentation of the sell-up concept and the mechanisms it involves. Both new alternative methods are described after that, given that they are both based on sell-up considerations. Q-Forecasting will come first, and then we will present the Fare Adjustment methods and their links with Q-forecasting. We will describe the methods assuming the required inputs are known and manually set, and we will explain the estimation processes used to dynamically compute the required inputs in the case of the Q-Forecasting method.

In Chapter 4, the simulation results, we will analyze the performance of both methods independently, at first without any estimation, and then using the estimator. Finally, we will combine both Fare Adjustment and Q-Forecasting methods. The point is to quantify the revenue benefits, but more interestingly to focus on the mechanisms involved, analyzing the loads – network wide, or in specified markets, looking at local and connecting passengers – fare mix, booking curves, sell-up curves and adjusted fares.

Finally, the conclusion – Chapter 5 – will summarize the benefits obtained from both methods in unrestricted environments and suggest directions for future work.

2. LITERATURE REVIEW

This chapter first reviews the work done to this point in Revenue Management, in traditional environments, and then presents the new kinds of environments involving less differentiated fare product structures. The conditions corresponding to this new kind of environments are due to the rapid expansion of low-cost carriers in America, Asia and Europe competing with legacy carriers in a larger and larger number of markets.

2.1. Review of Revenue Management's Traditional Assumptions and Algorithms

Revenue Management methods deal with the problem of maximizing revenues by optimizing the number of seats made available for each booking class (so-called "Seat Inventory Control", SIC). This means limiting the availability for low-fare, early bookings and simultaneously protecting enough seats for later-booking, higher fare passengers.

The first Revenue Management systems appeared in the early 1980's, and kept improving since that time, adding more and more sophisticated features to the seat inventory control and optimization. They are based on strong assumptions which remained reasonably true until recent years.

3 main techniques have been successively added to the Revenue Management Systems (RMSs): Overbooking, Fare Class Mix and OD Control.

There are two underlying assumptions on which are based the "traditional" RMSs. The first one is the independence of demand (mean and standard deviation) by fare class for a given flight. This means that the forecast predicts the demand for each fare class independently from other fare classes, based on observations of bookings on previous flights in the same fare class. It implies that the customer is supposed not to consider choosing between different fare classes. The second assumption is that fare structure is determined independently from this process by the pricing department of the airline (not described in this thesis).

Barnhart, Belobaba and Odoni[4], as well as **McGill and Van Ryzin**[5] provide very good summaries and descriptions of the evolution of RMS in the airline industry. As said before, the first RMS appeared in the early 1980's, focusing on the overbooking problem in order to increase revenue. Then, the 2nd generation provided threshold methods to control the inventory. Later, the 3rd generation added the ability to forecast and optimize by booking classes for each future flight leg departure. These 3rd generation RMSs are implemented today in the vast majority of the airline companies across the world.

[4] BARNHART, BELOBABA, ODONI, *Operation Research in the Air Transportation Industry*, Transportation Science, Vol 37, No 4, Nov 2003

[5] McGILL and VAN RYZIN, *Revenue Management: Research Overview and Prospects*, Transportation Science, Vol. 33, No. 2, May 1999

2.1.1. Overbooking Models

Overbooking models were aimed to reduce the revenue loss due to no-shows (used to represent 15 to 25% of total bookings at date of departure, down to 10% nowadays). They involve booking more seats than actually available on the aircraft, accounting for the probability of no-shows. Overbooking is a tradeoff (in terms of risks and costs) between denied boarding (airline image) and potential revenue loss from unsold or spoiled seats. These overbooking models include those developed by **Simon**[6], **Vikred**[7] and **Rothstein**[8,9].

2.1.2. Fare Class Mix Models

Some significant progress was then achieved regarding the forecasting part of Revenue Management by **Littlewood**[10], **L'Heureux**[11], and then further explored by **Lee**[12]. The demand forecasts are then fed into a seat allocation optimizer to determine the Booking Limits to be applied to each booking class, using serial nesting described in Chapter 3. The idea is not to allocate seats to partitioned classes, but instead to protect seats for higher classes (all available seats in the entire inventory are available to highest fare class during the entire booking process). The nested seat allocation problem has been solved by **Littlewood**[10] for two classes, and then extended to n classes by **Belobaba**[13,14]. It is based on the notion of Expected Marginal Seat Revenue (known as the **EMSRb**[15] method, which can be defined as "the average fare of the booking class under consideration multiplied by the probability that demand will materialize for this incremental seat"), and uses heuristic decision rules for nested booking classes. It has become the most commonly used SIC model among airlines' RMSs. The use of fare class mix models allowed an increase in revenue by **2 to 4%**[14] with respect to judgmental methods. Optimal formulations for the multiple nested class problem have been published by **Curry**[16], **Brumelle and McGill**[17], and **Wollmer**[18].

---

[6] SIMON, *An Almost Practical Solution to Airline Overbooking*, J. Transp. Econ. Policy 2, 1968
[7] VIKRED, *Airline Overbooking: Some Further Solutions*, J. Transp. Econ. Policy 6, 1972
[8] ROTHENSTEIN, *An Airline Overbooking Model*, Transportation Science 5, 1971
[9] ROTHENSTEIN, *Airline Overbooking: Fresh Approaches Needed*, Transportation Science 9, 1975
[10] LITTLEWOOD, *Forecasting and Control of Passenger Bookings*, AGIFORS, Israel, 1972
[11] L'HEUREUX, *A New Twist in Forecasting Short-Term Passenger Pickup*, AGIFORS, Bowness-On-Windermere, 1986
[12] LEE, *Airline Reservations Forecasting: Probabilistic and Statistical Models for the Booking Process*, PhD Thesis, MIT, Cambridge, MA, 1990
[13] BELOBABA, *Air Travel Demand and Airline Seat Inventory Management*, PhD Thesis, MIT, Cambridge, MA, 1987
[14] BELOBABA, *Application of a Probabilistic Decision Model To Airline Seat Inventory Control*, Operations Research 37, 1989
[15] BELOBABA, 1992
[16] CURRY, *Optimal Airline Seat Allocation with Fare Classes Nested by Origin and Destinations*, Transportation Science 24, 1990
[17] BRUMELLE and McGILL, *Airline Seat Allocation with Multiple Nested Fare Classes*, Operations Research 41, 1993
[18] WOLLMER, An *Airline Seat Management Model for a Single Leg Route When Lower Fare Classes Book First*, Operations Research 40, 1992

2.1.3. OD Control Models

The next step in the development of RMS was to account for network effects and manage the seat inventory by the revenue of the passenger's OD itinerary (instead of according to the fare class request on a flight leg), which became more feasible with the improvements in computer science as the problem became more and more complex. It was very important for airlines which operate large hub and spoke networks. The first OD-Control RMSs used virtual nesting to map itinerary+fare type to a (hidden) virtual class, in order to consider the network value of an ODF instead of just its fare class. They were developed by **Belobaba**[13], **Smith and Penn**[19], **Williamson**[20] and **Vinod**[21,22]. These RMSs gave preference to longer-haul, connecting passengers with higher fares, but did not address the need in certain cases for giving preference to two local passengers whose fares, when added, are higher than the connecting fare. DAVN (Displacement Adjusted Virtual Nesting) addressed this problem based on the Network Revenue Value, defined as the total itinerary fare minus the revenue displacement that might occur on connecting flight legs if the passenger's request for a multiple-leg itinerary is accepted (other than the legs under consideration). There exist a lot of variations on these methods, especially regarding the clustering process.

A simpler model was developed as well, called "Bid-Price Control" by **Smith and Penn**[19], **Simpson**[23] and **Williamson**[24]. In this method, the OD Fare is compared to a bid-price. The bid-price is an approximate displacement cost equal to the sum of dual prices associated with each crossed leg in the itinerary. The dual prices come from a deterministic LP and are used as marginal values for an incremental seat on different legs. If the bid-price is greater than the OD Fare, the request will be denied, if it is lower, the request will be accepted. These methods will be described more extensively later in the thesis.

The benefits they provide have been evaluated using the Passenger Origin-destination Simulator **PODS**[25], by **Belobaba and Wilson**[26], and quantified as an increase in revenue by 4 to 6% relative to conventional leg-based fare class control.

The successive improvements in airline RMS have indeed allowed increasing revenue, accounting for overbooking, fare class mix, and finally network effects. These advances remained nevertheless based on strong assumptions, such as the independence of demand by fare class, which tend to be violated frequently. This translates generally into the inability to recover from revenue loss when competing with a low-cost airline in a set of markets, as explained in the next sections.

---

[19] SMITH and PENN, *Analysis of Alternative Origin-Destination Control Strategies*, AGIFORS, New Seabury, MA, 1988

[20] WILLIAMSON, *Comparison of Optimization techniques for Origin-Destination Seat Inventory Control*, Master's Thesis, MIT, Cambridge, MA, 1988

[21] VINOD, *A set Partitioning Algorithm for Virtual Nesting Indexing Using Dynamic Programming*, Internal Technical Report, SABRE Decision Technologies, 1989

[22] VINOD, *Origin and Destination Yield Management*, in The Handbook of Airline Economics, D. Jenkins (ed.), The Aviation Weekly Group of the McGraw-Hill Companies, New York, New York, 1995

[23] SIMPSON, Using *Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control*, Memorandum M89-1, MIT, Cambridge, MA, 1989

[24] WILLIAMSON, Airline *Network Seat Inventory Control: Methodologies and Revenue Impacts*, PhD Thesis, MIT, Cambridge, MA, 1992

[25] Boeing PODS, developed by Hopperstad, Berge and Filipowski, 1997

[26] BELOBABA and WILSON, *Impact of Yield Management in Competitive Airline Markets*, Journal of Air Transportation Management 3, 1997

2.2. The Low-Cost Carrier Model

Although the low-cost carriers have existed for decades (Southwest Airlines was founded in 1971), it is only since the beginning of the new century that this kind of airline companies has significantly developed in the US, and also in Europe (Easy Jet, Ryanair), Asia (Air Asia, Tiger) and Oceania (Virgin Blue, Jet Star). Low-cost carriers carry now more than 25% of the domestic passengers in the US.

The effects of the LCC emergence have dramatically changed the airline industry landscape in the world. The main characteristic of these airlines – apart from their very low costs, is their simple fare/restriction product structure, that the legacy carriers are forced to match to avoid losing their market share. But this causes significant drops in revenues due to the fact that the legacy carriers have higher operating costs and their RMSs are based on the restrictions and advance purchase requirements features. As we will discuss later, the legacy carriers face a spiral down effect in their revenues in the attacked markets when restrictions are removed.

Weber and Thiel[27] have presented an analysis of the "Methodological Issues in Low-Cost Carriers Revenue Management", from which I summarize the main points of the comparison between LCC and legacy carriers:

Low-Cost Carriers	Legacy Carriers
✓Single product: no segmentation	✓Multiple product; segmented
✓1 way fares only	✓Mainly return fares
✓No connecting traffic	✓Mainly network traffic
✓No refund of cancelled tickets	✓Refundable ticket depending on fare
✓Single product offered at different fares	✓Fare differentiation according to product differentiation
✓Booking (bkg) control via fares offered	✓Bkg control via booking class available
✓Only 1 fare available at a time	✓Several fare products available at a time
✓Demand supposedly price-sensitive	✓Demand supposedly product-sensitive
✓Time series based forecast does not reflect price elasticity	✓Time series based forecast dominates Revenue Management
✓RM mainly manual or semi-automatic	✓ RM control by fully automatic systems

Of course, all low-cost carriers do not fit exactly in the first column, and there is a very large range of models between the two extremes described above.

But from this, we can point out the number of non-negligible differences between both models. Nevertheless, there are also some points which apply to both low-cost and legacy carriers:

- ✓ The average fare increases as we approach the departure date
- ✓ The low amount of revenue obtained from the lowest fare classes

[27] WEBER and THIEL, *Methodological Issues in Low-Cost Carriers Revenue Management*, Presentation to AGIFORS, Auckland, New Zealand, Jan 2004

Dunleavy and Westerman[28] provide also a good analysis of the low-cost carrier model compared with the legacy carrier model. They point out the complexity problem the legacy carriers' RM departments have to face regarding the finer and finer segmentation of demand (caused by pricing strategies); whereas in the case of low-cost carriers, a simple single product is offered at a unique price which changes depending only on how close the day-of-departure is to the day-of-booking. Using their advantages – lower costs since they started operations, and avoidance of competition with other LCCs – the low-cost carrier model provides sustainable profitability to these airlines – at least in a short term.

The point is that both models have to co-exist on the same legs, and inevitably, each of them has to incorporate a significant part of the competitor's model in its own model. The low-cost carriers may need some less simple RMS, as well as legacy carriers need to extend their RMS by adding some features allowing them to address the presence of LCCs in some of their markets. We will, in this thesis, focus on the second part.

Along with the emergence of low-cost carriers, a new segregation of demand has arisen. **Boyd and Kallesen**[29] provide an analysis of the differences between priceable and yieldable demand, also respectively called price-oriented and product-oriented demand. Yieldable demand refers to customers who buy a ticket with respect to the product it represents (with restrictions and advance purchase requirements) and corresponds to the traditional assumptions of Revenue Management. Priceable demand corresponds to customers who buy the lowest available fare, irrespective of restrictions and advance purchase requirements, and it generally refers to markets in which a low-cost carrier entered. In most cases, both types of demand coexist in the same leg, but they cannot be addressed the same way, providing new challenges in modeling, forecasting and optimization for Revenue Management Systems.

2.3. Traditional Revenue Management in Less Restricted Environments: the "Spiral Down" Effect

The legacy carriers have responded to the entrance of LCCs in some of their legs with major changes in their sets of fare products (in the specified legs). They tend to partially or fully match the LCC in terms of fares products:

- ✓ Lower fares
- ✓ Removal of certain restrictions (Saturday night stay, non-refundability and/or change fees)
- ✓ Reduced advance purchase requirements

These changes might be quite complex; they might be applied to all fare classes in a market, or only to a reduced number of classes and they might be applied to connecting and local fares or only to local fares.

[28] DUNLEAVY and WESTERMAN, *Future of Airline Revenue Management*, Journal of Revenue and Pricing Management, Vol 3, No 4, Jan 2005
[29] BOYD and KALLESEN, *The Science of Revenue Management when Passengers Purchase the Lowest Available Fare*, Journal of Revenue and Pricing Management, Vol 3, No 2, Jul 2004

In these new kinds of hybrid networks, the fundamental assumptions of traditional RMS are in jeopardy:

- ✓ The independence of demand by fare class is no longer valid. Customers are offered undifferentiated products, so they buy the lowest available fare. Thus demand for highest classes do not materialize, and it becomes impossible to establish forecasts for highest fare classes based on historical data.
- ✓ The customer segmentation rules are violated as well. The differences between "leisure" and "business" passengers are eroding, business passengers becoming more sensitive to price for instance.

In these conditions, it happens that the traditional RMSs perform quite poorly, as far as they are by definition based on the fare product structure (restrictions + advance purchase requirements). The combination of the violation of the assumptions and the poor performance of traditional RMSs leads to the "buy-down" and "spiral-down" effects – presented below – which translate into severe revenue losses.

Both phenomena have been described based on analysis done with the Passenger Origin-Destination Simulator PODS by **Cusano**[30], but also by **Ozdaryal and Saranathan**[31].

The buy-down Phenomenon consists of a series of simultaneous processes which, combined, lead to a decrease in the average fare paid per passenger:

- ✓ The removal (total or partial) of the restrictions and advance purchase requirements leads to the violation of the customer segmentation principles
- ✓ The assumption of demand independence between fare classes is invalidated
- ✓ The excess in capacity caused even more losses when lower class demand started to materialize close to departure date (violation of the customers' behaviors assumptions)

The Spiral-down Phenomenon is the direct consequence of the buy-down phenomenon. As said, the removal of fare product differentiation led to a decrease in yield. Legacy airlines tried without success to replace these fare rules by new inventory strategies, and were stuck into a cycle leading to lower and lower revenues as shown on *Figure 1*[31].

[30] CUSANO, *Airline Revenue Management Under Alternative Fare Structures*, Master's Thesis, MIT, Cambridge, MA, 2003

[31] OZDARYAL and SARANATHAN, *Revenue Management in Broken Fare Fence Environment*, Presentation to AGIFORS, Auckland, New Zealand, Jan 2004

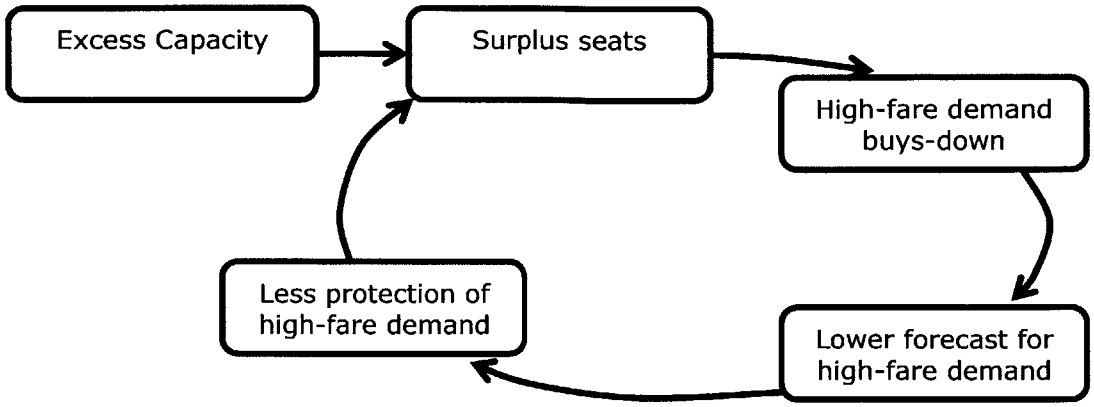

Figure 1: Spiral-Down Effect

When a traditional RMS deals with an unrestricted environment, the excess of capacity (caused for instance by the Sept. 11, 2001 events) leads high-fare passengers to buy-down. This consequently encourages lower forecasts for high-fare demand (given that high-fare demand did not materialize), reducing the protection for high-fare passengers. Finally, this change in booking limits leads to an increase in the number of seats available for lower classes. The cycle then spirals down.

The customer, who is given the choice between buying a ticket to the legacy carrier or to the low-cost carrier, will for sure choose the LCC if the legacy carrier does not match with its competitor. So the legacy carrier has to match, and then, as described earlier, the spiral down effect feeds itself when the airline removes the fences between fare classes. This results in lower revenues for the "attacked" airline.

Cooper and Homen-de-Mello[3] have modeled mathematically the spiral down effect occurring when using traditional forecasting methods. The purpose of their paper is to study the (more general) effect of a violation of the fundamental assumptions taken for granted in the traditional revenue management algorithms. They define, for the model they have built, precisely under what conditions and how the spiral down effect occurs.

Since 2001, the revenue management world has focused on the spiral-down problem to try to come up with alternatives, but at this point, no one has found the optimal solution in helping the legacy carriers with recovering significantly from revenue loss due to competition with LCC in a set of markets. This thesis will present two complementary alternatives to traditional revenue management methods designed to address the problems raised above.

3. APPROACH

3.1. PODS Overview

This thesis makes use of PODS, the **Passenger Origin-Destination Simulator**[25], developed at Boeing by Hopperstad. It simulates realistically the environment of competing airline companies. PODS evolved from the **Decision Window Model**[32] by incorporating major simulation capabilities in addition to schedules, airline image and aircraft types, such as the Revenue Management systems, the fares offered on each flight and path, the restrictions associated with each fare, and the possibility for the passenger to choose between several competing airlines.

PODS simulates for a single day's departures the booking process on a network of flights for one or more airlines which can use different Revenue Management methods. The simulator produces various output files which allow the researcher to analyze the impact of the different revenue management and forecasting methods and/or different fare structures in various environments.

Since PODS has been described several times by **Zickus, Gorin and Carrier**[33], the description of the simulator in this thesis will be an overview and the reader should refer to these materials for more detailed descriptions.

3.1.1. The Simulator

PODS will be used to simulate several airlines operating a network of flights in a competitive environment in order to measure the performance of different revenue management methods. The simulator is composed of 4 major elements linked as shown on *Figure 2*.

[32] *Decision Window Path Preference Model (DWM),* The Boeing Company, Feb 1994
[33] ZICKUS, *Forecasting for Airline Network revenue Management; Revenue and Competitive Impacts*, May 1998
GORIN, *Airline Revenue Management: sell-up and Forecasting Algorithms*, June 2000
CARRIER, *Modeling Airline Passenger Choice: Passenger Preference for Schedule in the Passenger Origin-Destination Simulator (PODS)*, May 2003

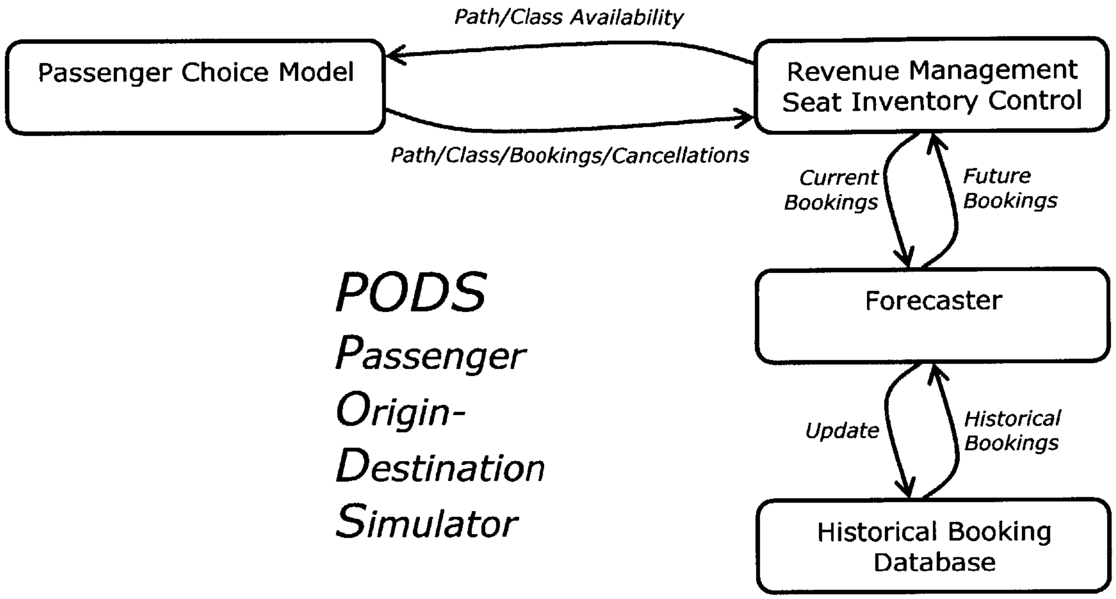

Figure 2: PODS Flow Chart

Source: **Hopperstad**[25]

The simulation flow begins at the bottom of *Figure 2*. The booking data from previous flights is collected into the historical database, which, with the current booking status (from the Revenue Management optimizer), is used by the forecaster to estimate the future demand for a given flight. These expected bookings are then fed into the Revenue Management optimizer, which computes the seat protections and availability based on current path and class bookings as well as cancellations. Finally, the availability data is fed into the passenger choice model, which generates future passengers' characteristics and assigns them to available path-fare combinations according to their decision window, budget and preferences. Then all the information is input back into the model again to provide historical data to the Revenue Management optimizer for the next optimizations and future flight departures

In the next paragraphs, only the relevant parameters, with respect to the subject of this thesis, will be described. For a more complete description of the PODS Model, its inputs, fares, restrictions, outputs, various networks, the reader should refer to the works of **Zickus, Gorin and Carrier**[33] which contain complementary and extensive explanations of the different aspects of the simulator. The following is just a brief summary.

PODS is a simulator of the behavior of the passengers, the way they choose between paths, airline companies, and fare products given what the airline companies decided to make available in terms of these fare products. There are therefore two distinguishable components in PODS as shown in *Figure 3*.

PODS

How Passengers choose flights and fares
Based on the Passenger Choice Model

How Airlines set their availabilities of paths and fare classes
Based on historical data and forecasts

Figure 3: The Two Main Components of PODS

One particular run of PODS is comprised of 5 independent trials, each of them being the repetition (600 times generally) of a single departure day (for instance Mondays) called a sample, over and over without cycles or seasonality. The first 200 samples of each trial are burnt to avoid transient effects. In a given trial, the samples are not independent, since they are used for historical purposes. Finally, the results provided are averaged over the 5 trials.

3.1.2. The Network Parameters

At the network level, the inputs include:

- ✓ The number of Airlines
- ✓ The number of Origin-Destination Markets
- ✓ The number of flights per day
- ✓ The departure time and capacity of each flight
- ✓ The number of fare classes and associated fares, advance purchase requirements and restrictions (like non-refundability, change fee, Saturday night stay)

3.1.3. The Passengers

At the passenger level, the inputs include:

- ✓ The demand by day, OD and Passenger Type (Business, Leisure) are given in terms of means and standard deviations that vary day after day
- ✓ The arrival curves by passenger types, which typically look like those of *Figure 4*. It translates the fact that leisure passengers book early while the majority of business passengers book during the last days before departure date.

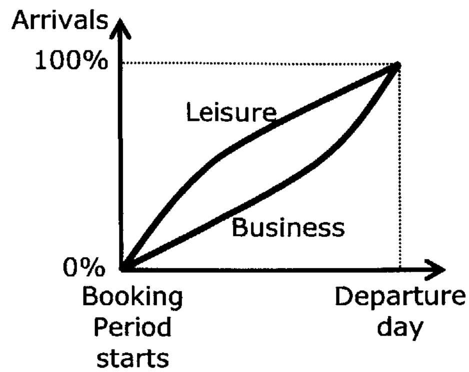

Figure 4: Arrival Curves by Passenger Type

The Passenger Choice Model.
This model is based on and has evolved from the Boeing **Decision Window Model**[32] (DWM), itself using as input the maximum willingness-to-pay (WTP) curves given by passenger types as shown in *Figure 5.*

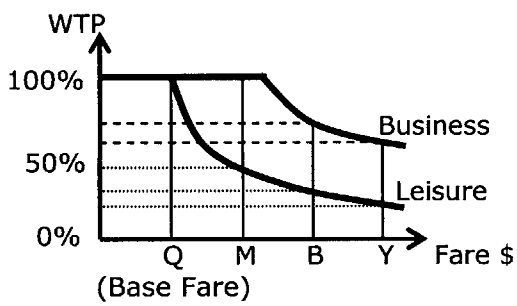

For instance, in this 4-fare-classes environment, we notice that only 50% of the leisure passengers are willing to pay for the M fare whereas still 100% of the business passengers are willing to pay for it.

Figure 5: Willingness-To-Pay Curves by Passenger Type

To account for the cost they represent for a given passenger type, the OD fares are also adjusted with the disutility costs associated with the different restrictions R_i:

$$\text{Total Generalized Cost} = \text{OD Fare} + \Sigma_{R_i}(\text{Evaluated Cost of } R_i)$$

These features are used and organized in the process shown in *Figure 6.*

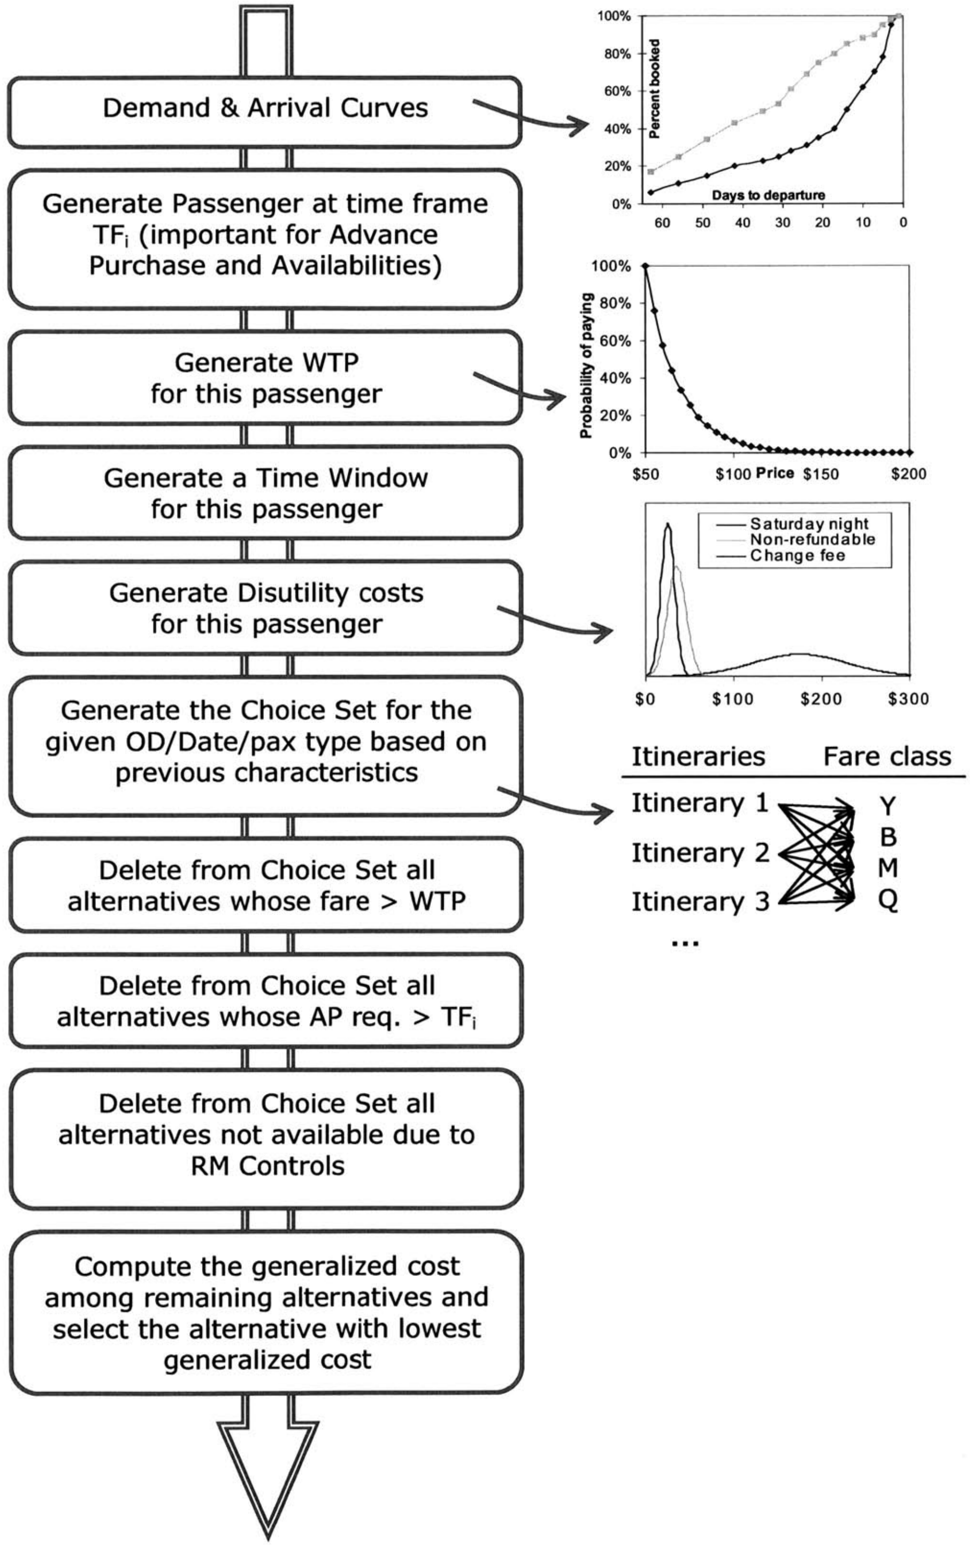

Figure 6: Passenger Choice Model – Process Overview

3.1.4. Forecasting and Detruncation

As mentioned during the introduction and the literature review, the main assumption on which the traditional revenue management methods are based is the independence of demand by fare class. The forecasts by fare-class are re-computed at every Time-Frame (TF) during the booking process, that is at arbitrarily fixed shrinking intervals until the day-of-departure.

Forecasting (by time-frame/itinerary/fare class)
The forecasts in PODS are computed at each time frame, and are based on historical data (26 previous samples or "Mondays"). The goal is to estimate the unconstrained demand, and the first step for doing this is to test if we need to detruncate the historical data or not, as explained in *Figure 7*.

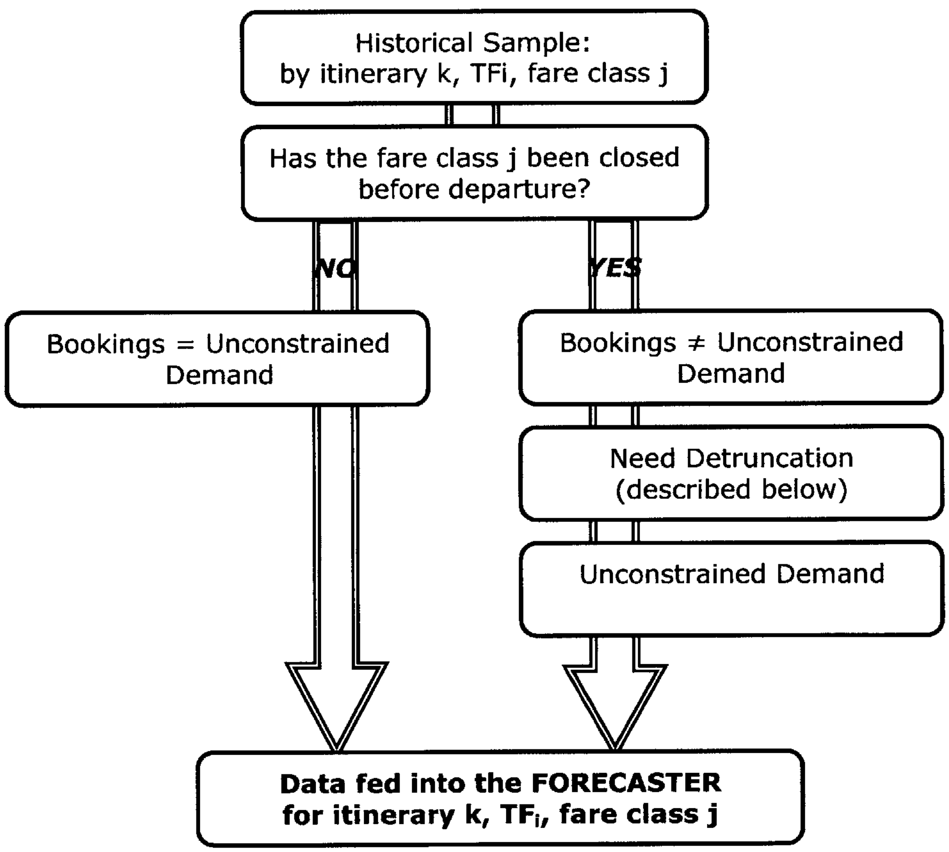

Figure 7: Detruncation/Forecasting Process

Once the unconstrained demand is estimated from the previous samples, several forecasting methods can be used:

- ✓ *Pick-Up Moving Average* (PUMA), based on the "Bookings-To-Come" from TF_i: average over the previous samples of occurrences of additional bookings between TF_i and the departure date

- ✓ *PUMA with exponential smoothing*: allows putting more weight on the more recent samples with respect to older ones
- ✓ *Booking Regression*: estimates from the previous samples the parameters a and b so that the number of booked seats (BKD) is given by the equation $BKD(TF_0) = a + b*BKD(TF_i)$

Detruncation

Detruncation is needed when fare classes closed down before departure date, in order to estimate the unconstrained demand and feed the forecaster. Without detruncating, the forecast would not be accurate and revenues would not be optimized.

Two different methods can be used to detruncate. The first method is called *Booking Curve* Detruncation (Probability-based) and is presented in *Figure 8*.

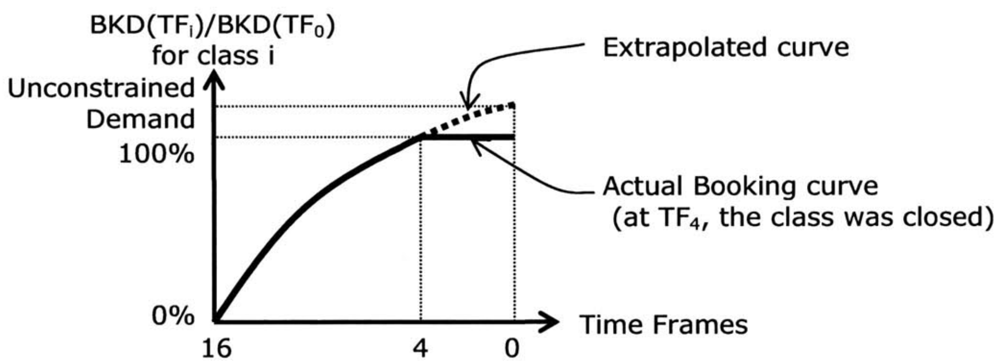

Figure 8: Booking Curve Detruncation Process

The second method is *Projection* Detruncation (Spill-based). Given the fact that the class was closed down before the departure date, we only observe a part of the unconstrained demand (see *Figure 9*). The method is based on estimating the ratio of opened class over the unconstrained demand, and iterating until convergence. It is also governed in PODS by a parameter tau (τ).

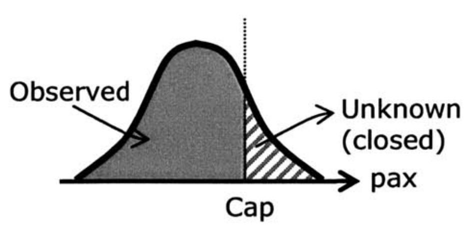

Figure 9: Projection Detruncation Process

For more details, refer to **Gorin**[33], in which is described precisely the mechanisms involved as well as the range and influence of the tau parameter.

3.1.5. Revenue Management - Seat Inventory Control

The goal of the revenue management optimization is to compute the booking limits for the different classes of each flight given the demand forecasts for each itinerary. The family of Revenue Management methods is very broad, but we can regroup the methods into sub-families which share some common features, from the least to the most sophisticated method.

Base Line
Used as reference to measure the performance and compare revenue management methods, the **First Come First Served** (FCFS) booking method is the most basic way to fill the aircraft. There is no booking limit in this method; the passengers buy seats as soon as they arrive, and as long as the requested class is not filled or closed due to advance purchase requirements.

Threshold Algorithms
The Threshold algorithms are the first step in managing booking limits. A threshold (between 0% and 100%) is associated with each fare class and as soon as the ratio of "total bookings over capacity" reaches one of the thresholds, the corresponding class is closed down. In PODS, two different threshold methods are implemented: Fixed Threshold (FT) for which the threshold values are specified (typically 1.00, 0.85, 0.60, 0.40 for the 4-fare-class case) in the input file and remain the same for the entire booking period, and Adaptive (or Accordion) Threshold (AT), for which a load factor target is set in the input file - generally at 80% - and the threshold values are computed and optimized at each time frame according to the actual bookings recorded until this point.

Fare Class Yield Management (FCYM)
It is a leg-based RM Method, generally used with Pick-Up Forecasting and Probabilistic Detruncation. The optimizer deals with a nested inventory, based on the Expected Marginal Seat Revenue (**EMSRb**[11]), the revenue we can expect from the next seat available for a given class (= OD Fare * probability to sell it).

Again, the optimization is based on the assumption that the demand for each fare class is described by an independent distribution. Given the standard fare structure and demand for a 100-seat flight represented in *Table 1*, the EMSRb Booking Limit Optimization is shown for this example on *Figure 10*.

Class	Fare	Avge	stdev	*BL*
Y	$400	20	5	*100*
B	$250	30	8	*82*
M	$160	40	10	*46*
Q	$100	40	10	*18*

Table 1: Example of Fare Class Structure in the Process of Optimizing the Seat Inventory

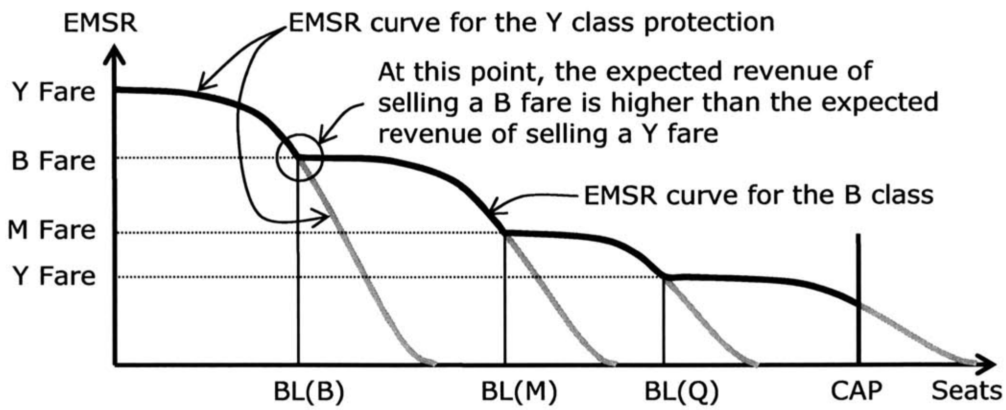

Figure 10: EMSR Curve and Booking Limits

The booking limit for fare class i represents the total number of bookings at which the fare class is closed down, to protect seats for higher fare classes (<i). For instance, in the case described above, if BL(M) bookings have been recorded, only B and Y fares will be available for the next requests. Nested limits are shown on *Figure 11*.

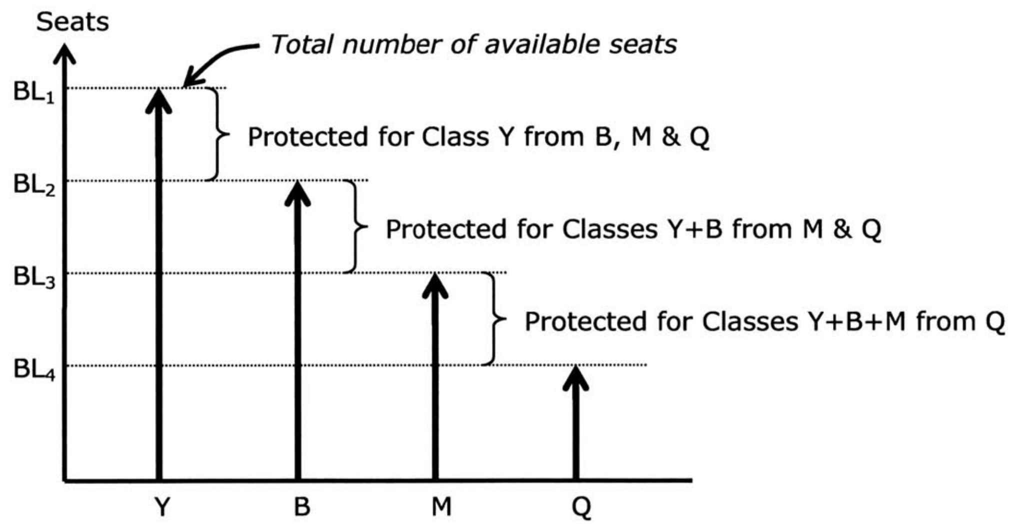

Figure 11: The Nested Booking Limits

Greedy Virtual Nesting (GVN)

GVN is the first step in the process of accounting for connecting passengers, even if it is still a leg-based optimization. If we deal with a single leg without any connections, the result in terms of Booking Limits will be the same as the one obtained in the previous case. Let's consider now the case when we have two legs presented on *Figure 12*.

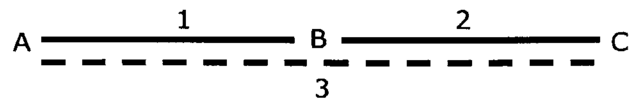

Figure 12: Simple 2-legs OD Market

GVN will create, as its name indicates, n virtual classes for each LEG, in which it will order both local and connecting fares. Then, it applies the EMSR optimization for each LEG to the virtual classes as shown on *Figure 13*.

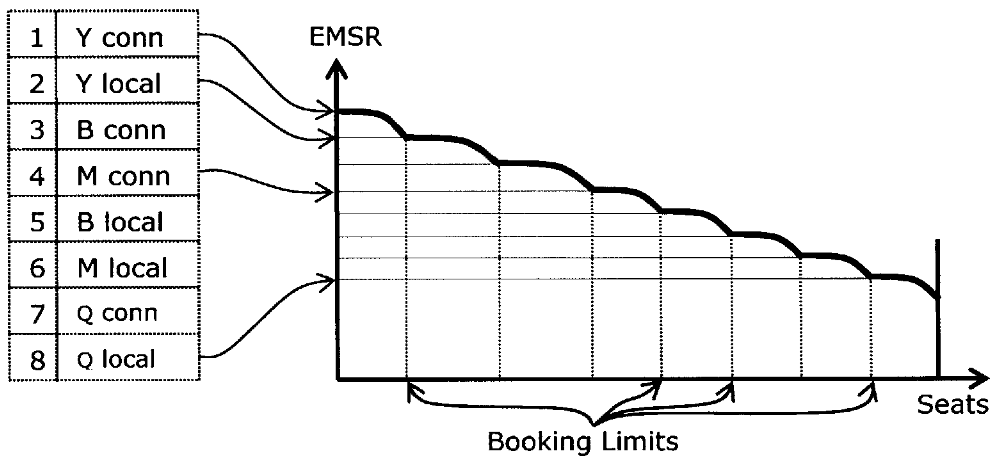

Figure 13: EMSR Optimization in the GVN Case

Consequently, a connecting booking will be accepted if ALL the virtual classes in which the fare is bucketed in each of its legs are still open at the time of the request. Given the relative high fare of the connecting itinerary with respect to locals individually, preference will be given to connecting passengers in most of cases. The problem is that it may cause the denial of two local passengers on both legs if the itinerary involves two legs, and often, it happens that the sum of the two local fares is greater than the connecting fare for a same fare class.

Displacement Adjusted Virtual Nesting (DAVN)
The goal of DAVN is to answer the problem raised at the end of the previous paragraph. This time, the method requires a path/class forecasting and detruncation. The principle is to apply a penalty to the connecting fares that accounts for the potential displacement of a local passenger. On the leg i that belongs to the connecting itinerary, the passenger's total fare will be replaced in the bucketing by the fare minus the sum of the displacement costs associated with all other legs involved in the itinerary.

$$Bucketed\ Fare_{Legi} = ODFare - \sum_{\substack{j \in Itinerary \\ j \neq i}} d_j$$

Where d_j represents the displacement cost associated with leg j.

The displacement costs are computed as the dual solution of a deterministic linear program solved over the network, in which we look for maximizing total revenue subject to capacity and forecast constraints. This LP problem was described by **Williamson**[34].

Thus, for each given leg, the ranking in the bucketing process is based on the network value of each booking, and gives less preference to connecting passengers than GVN does. The displacement costs as well as the size of the buckets are regularly re-optimized during the booking process.

Example: if we use the small network described above in the GVN sub-section (2 legs, 3 markets), given the OD forecasts (considered being deterministic), the LP is formulated as follows:

$$MAX \ \left(p_1^1 x_1^1 + p_1^2 x_1^2 + p_2^1 x_2^1 + p_2^2 x_2^2 + p_3^1 x_3^1 + p_3^2 x_3^2\right) \ = \ MAX \ \left(\sum_{\substack{OD \\ i}} \sum_{\substack{Class \\ j}} p_i^j x_i^j \right)$$

Subject to

$$x_i^j \ \langle \ f_i^j \qquad\qquad \forall \ OD \ i, \ class \ j$$

$$\sum_{OD \ i} \ \sum_{class \ j} x_i^j \delta_i^k \ \langle \ C_k \qquad\qquad \forall \ leg \ k$$

Where:

i represents an OD market
j represents a fare class
k represents a leg
p_i^j *represents the fare of fare class j in the OD i*
x_i^j *represents the number of passengers flying in fare class j in the OD i*
f_i^j *represents the mean of the forecast for fare class j in the OD i*
δ_i^k *= 1 if leg k is traversed by the passenger's itinerary in the OD i*
= 0 otherwise
C_k *represents the aircraft capacity on leg k*

The solution of this LP provides optimal values for the set $\left(x_i^j\right)_{i,j}$*, and more interestingly the dual solution* $\left(d_k\right)_k$ *represents, for each leg k, the marginal revenue obtained by adding an extra seat to the capacity on leg k. These values are used as the displacement costs in the bucketing process of the connecting passengers described above.*

[34] WILLIAMSON, Airline Network Seat Inventory Control: Methodologies and Revenue Impacts, June 1992

Heuristic Bid-Price (HBP)

HBP is based on the GVN structure, and tries to approach the ideas developed in DAVN while still using a leg-base forecast instead of the path/class forecast needed for DAVN. It is both an improved GVN and an approximation of DAVN. The HBP presented below was developed by **Belobaba**[35].

Let's consider the network previously described (2 legs, 3 markets). For each leg, HBP compute the EMSR-optimized Booking Limits as shown on *Figure 14*.

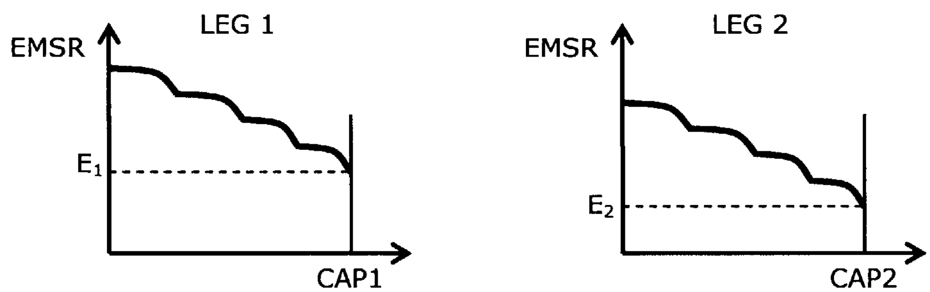

Figure 14: EMSR Curves in the case of HBP

E_1 and E_2 are, for each leg, the minimum value at which we accept a booking (whatever the fare class). In no case we want to accept a booking for a value less than E_1 in leg 1 and less than E_2 in leg 2. (It should be noticed that the higher the demand, the higher E_i)

Therefore, we can infer that some combination of these two numbers E1 and E2 is a good indication for a minimum price at which we would accept one connecting booking. In HBP, this minimum price is called the Bid Price (BP), and is defined by:

$$BP_{OD3} = \max(E_1,E_2) + d*\min(E_1,E_2)$$

Where d is the bid-price coefficient and can roughly be interpreted as the probability to displace two connecting local passengers (e.g. for instance 0.5*0.5=0.25 if we assume roughly that the mix is 50% local, 50% connecting).

Probabilistic Bid-Price (ProBP)

This is the – theoretically – most advanced RM Method at this point, using a forecast based on path and fare class. It is a probabilistic network optimization algorithm (the drawback of DAVN was to be based on deterministic assumptions). On the other hand, it is the reason why it is quite complex and more difficult to implement. For a complete and detailed description of ProBP, refer to **Bratu**[36].

ProBP is an iterative algorithm. The first step proposes to draw a complete EMSR graph for both legs of the connecting itinerary, taking into account all possible other

[35] BELOBABA, The Evolution of Airline Yield Management: Fare Class to Origin-Destination Seat Inventory Control, Handbook of Airline Marketing, 1st Ed., 1998, Chapter 23, pp. 285-302

[36] BRATU, Network Value Concept in Airline Revenue Management, June 1998

connections. From the graphs, we compute the values E_1 ($=E^1_1$) and E_2 ($=E^1_2$) defined in the HBP paragraph. This is the end of the first iteration.
Next, the highest OD fares (let Y^1_1 and Y^1_2 be these highest fares) are pro-rated:

$$Y_1^2 = Y_1^1 \times \frac{E_1}{E_1 + E_2} \quad \text{and} \quad Y_2^2 = Y_2^1 \times \frac{E_2}{E_1 + E_2}$$

And then the EMSR optimization is applied once again, providing E^2_1, E^2_2 ...
It happens that the sequences of E^i_1 and E^i_2 converge very quickly and it has been derived (See **Bratu**[19]) that is converges toward the optimal solution, (E^f_1 and E^f_2). The Bid-Price is then defined by $BP = E^f_1 + E^f_2$. Again, a connecting booking will be accepted if its fare is greater than BP.

Some other RM methods exist which combine some features of several of the methods described above.

3.2. The Specifics of PODS Related to this Thesis

3.2.1. Environments

Single Market
The single market case (*Figure 15*) in PODS consists of 2 airline companies, competing in one Origin-Destination Market, with 3 flights per day in only one direction (same direction for both). Both of them offer 8 fare classes with associated products.

Figure 15: Single Market

Network D with New Entrant
The Network D is more complicated than the single market, since it involves 482 OD Markets and 3 airline companies. It is based on the networks of 2 legacy carriers which operate a hub (Hub 1 in MSP, Hub 2 in DFW). Each airline operates 41 legs, between 20 spokes on the West coast and its hub, 20 other legs between its hub and 20 spokes on the East coast, and finally 1 leg between both hubs. The legs are organized in banks, and there are 3 of them during the day, corresponding to 3 flows from the West coast to the East coast.

The "New Entrant" (e.g. LCC) is operating non-stop flights only between Hub 1 and the 10 largest markets on the East coast (with three flows per day as well). It is therefore in direct competition with the legacy carrier operating Hub 1 in 10 OD Markets. Obviously, the airline operating Hub 2 (called AL2) will almost not be affected by this competition given that it only operates connecting flights in the 10 "New Entrant Markets" (NEM) and was already not competitive against AL1 in these markets. The schedule of the new entrant's flights are the same as the legacy carrier's in order to avoid any schedule effects which are not the object of this analysis.

PODS allows to use, for both legacy carriers, a different fare product structure in these 10 markets.

The 482 Origin-destination Markets are mixed between

- ✓ Non-stop flights from 20 western spoke cities to hubs H_1 and H_2
- ✓ Non-stop flights from hubs H_1 and H_2 to 20 eastern spoke cities
- ✓ 82 non-stop "local" markets to/from/between hubs operated by the legacy carriers
- ✓ 10 non-stop "local" markets between hub H_1 and10 eastern operated by the LCC
- ✓ 400 connecting O-D markets

Figure 16, *Figure 17* and *Figure 18* represent the Airlines' networks.

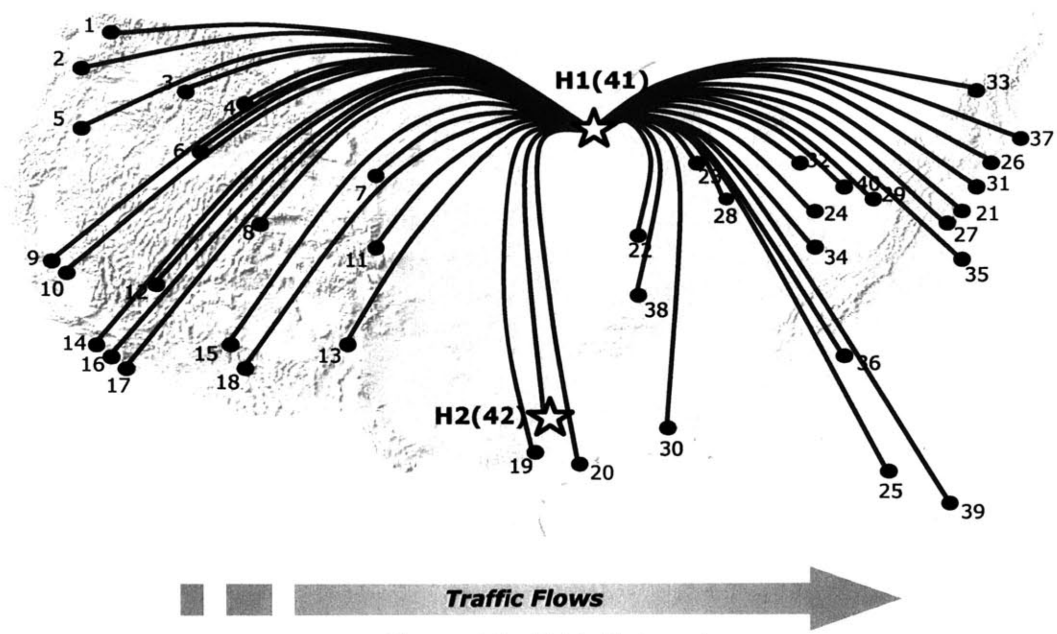

Figure 16: AL1's Network

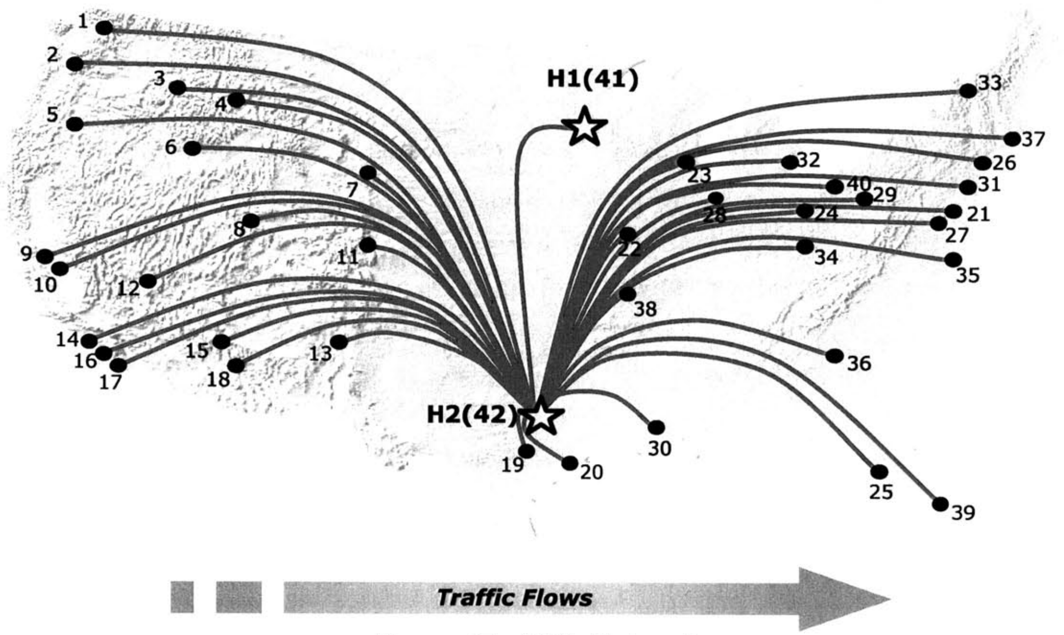

Figure 17: AL2's Network

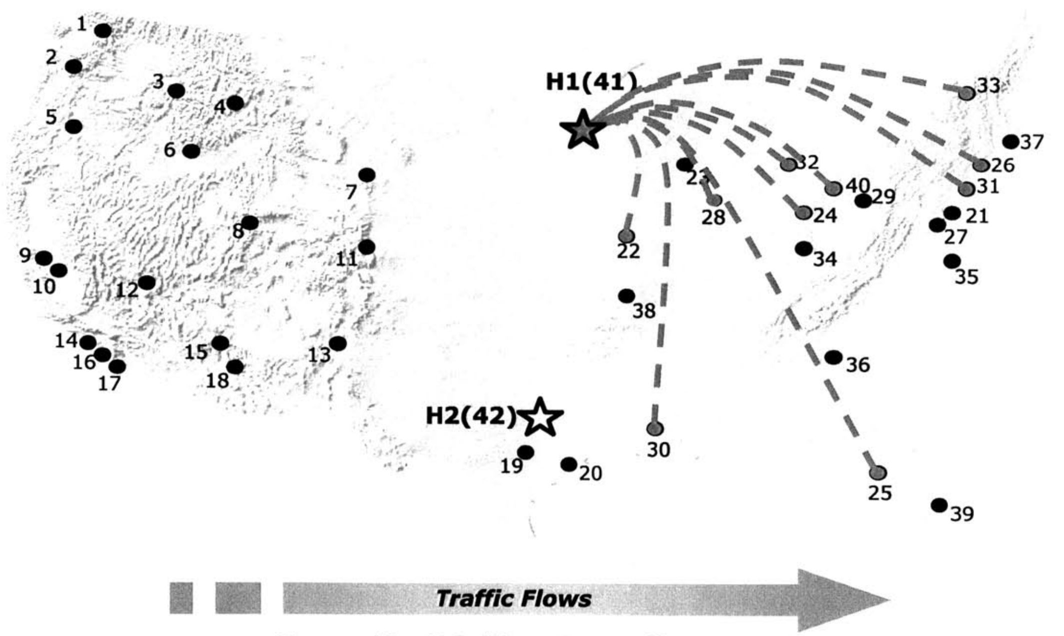

Figure 18: AL3 (New Entrant)'s Network

The impacts of the New Entrant on AL1's network are shown on *Table 2*.

IMPACT		Total	Affected	Percent
Total	Legs	126	30	24%
	Markets	482	10	2%
Local Markets		42	10	24%
Network Level	Passengers	7,200	878	12%
	RPMs	9,524,153	595,306	6%
Leg-level	Passengers	10,629	2,740	26%
	RPMs	10,371,190	1,856,086	18%

Table 2: Impact of the New Entrant on AL1's network (Source: Gorin[44])

3.2.2. Fare Products

The fare products in PODS are comprised of fares by fare class, associated with restrictions and advance purchase requirements. The restrictions are defined by their disutility costs. R1 is the strongest one (Saturday night stay requirement for instance), R2 and R3 have a lower cost (change fee or non-refundability for instance). The Advance Purchase requirements (AP) are defined by days before the date-of-departure, and generally they are referred in terms of Time-Frames (TF). The re-optimization time in PODS are indeed organized in 16 Time-Frames, corresponding to days to departure date as shown in *Table 3*.

Time Frame		1		2		3		4		5		6		7		8		9		10		11		12		13		14		15		16	
Days to departure	63		56		49		42		35		31		28		24		21		17		14		10		7		5		3		1		0

Table 3 Time-Frames Definitions

Fully Restricted Fare Products:
The *Table 4* presents an example of fare product, with 4 fare classes (Network D case) in a fully restricted environment.

Class	Fare	Advance Purchase TF	Advance Purchase Days	Restrictions R_1	Restrictions R_2	Restrictions R_3
Y	\$ 400	16	0	0	0	0
B	\$ 200	12	7	1	0	0
M	\$ 150	10	14	1	1	0
Q	\$ 100	8	21	1	1	1

Table 4: Traditional 4 Fare Classes Structure

The Advance Purchase (AP) requirements are given here in terms of days to departure date or in terms of number of Time Frames during which the fare product is available, starting from the beginning of the booking period (TF_1).

Fully Unrestricted/Undifferentiated Fare Products:
Table 5 and *Table 6* are two examples of fare products, with either 8 (single market case - *Table 5*) or 4 fare classes (Network D case, in the so-called 10 NEM - *Table 6*) in a fully unrestricted environment, without Advance Purchase requirements or Restrictions.

Class	Fare	Advance Purchase TF	Advance Purchase Days	Restrictions R_1	Restrictions R_2	Restrictions R_3
Y	$ 200	16	0	0	0	0
B	$ 160	16	0	0	0	0
M	$ 125	16	0	0	0	0
Q	$ 90	16	0	0	0	0

Table 5: Unrestricted 4 Fare Classes Structure

With respect to the restricted environment, the fare ratios have been compressed (and the base fare is reduced by $10), restrictions have disappeared and all fare products are available during the entire booking period (e.g. no more advance purchase requirement) unless closed by the Revenue Management system.

Class	Fare	Advance Purchase TF	Advance Purchase Days	Restrictions R_1	Restrictions R_2	Restrictions R_3
1	$ 625	16	0	0	0	0
2	$ 515	16	0	0	0	0
3	$ 405	16	0	0	0	0
4	$ 295	16	0	0	0	0
5	$ 240	16	0	0	0	0
6	$ 185	16	0	0	0	0
7	$ 155	16	0	0	0	0
8	$ 125	16	0	0	0	0

Table 6: Unrestricted 8 Fare Classes Structure

In the single market case "unrestricted", we usually use the 8-fare-class structure, with, like in the case of the 4-fare classes, the availability of all undifferentiated 8 fare products during the entire booking period.

3.3. Sell-up: Introducing Alternatives to Address the Unrestricted Environments' "Spiral Down" Problem

3.3.1. Traditional Sell-up Utilization

Sell-up happens when a passenger is denied booking on a particular flight and decides to pay more for the same flight and accepts the next higher fare available. In the end, if the passenger travels with the same airline, he (or she) can be either recaptured on the same flight in a higher fare class, or on another flight in the same fare class. In both cases, this represents an increase in revenue that we want to account for in the seat inventory control process.

The sell-up model, presented by **Belobaba-Weatherford**[37] and its benefits are extensively detailed by **Gorin**[33] and **Bohutinsky**[38], but as the methods presented further are significantly based on this important notion, this is a short summary about sell-up, willingness-to-pay and the way they can be applied into PODS.

The Maximum Willingness-To-Pay (WTP) represents the maximum price a passenger is willing to pay for a flight, and is used by the Revenue Management system in the estimation of the probability that a passenger will agree to sell-up. From this point of view, one of the main means used in Revenue Management is to get each passenger to pay his (or her) maximum price. Legacy airlines have tried to differentiate customers according to their maximum WTP, generally into 2 passenger types: business (time sensitive, not price sensitive, time constraint) and leisure (price sensitive, less time constraint and look for interesting deals). Consequently, the disutility costs associated with each restriction is different according to the passenger type. The differentiation done between business and leisure passengers is another important assumption which fades nowadays.

Therefore, accounting for sell-up as been considered as a means to increase profit and airlines have tried to integrate it into the revenue management algorithms.

Bohutinsky's Master's Thesis has shown that:

- ✓ sell-up depends on demand relative to capacity (the higher the demand, the higher the sell-up probability)
- ✓ sell-up rate is higher in higher classes
- ✓ in a competitive market, the competitor can recapture the denied passenger, especially as travel websites are developed which allow easy comparison between the products offered by each airline

Belobaba and Weatherford developed a sell-up modification to be used in the EMSRb algorithm, based on a decision tree which allows for adjusting the Booking Limits in order to account for sell-up probabilities.

[37] BELOBABA-WEATHERFORD, *Comparing Decision Rules that Incorporate Customer Diversion in Perishable Asset Revenue Management Situations*, Decision Science 27, 343, 1996
[38] BOHUTINSKY, The Sell-up Potential of Airline Demand, Masters Thesis, MIT, Cambridge, MA, 1990

In the 2-fare-class case, the idea is as follows. The new booking limit (which intuitively increases protection for higher fare classes) simply accounts for the probability of sell-up:

$$EMSRb\left(S_2^1 + V_2^1\right)\times\left(1-p\right)+p\times f_1 = f_2$$ [13]

Where:

S_2^1 represents the protection for class 1 from class 2

V_2^1 represents the additional protection obtained through the heuristic

p represents the probability of sell-up from class 2 to 1

f_i represents the fare of class i

It is illustrated on the following *Figure 19*, where f_2 appears to be the weighted mean of f_1 (affected of weight *p*) and $EMSRb(S_2^1+V_2^1)$ (affected of weight *1-p*):

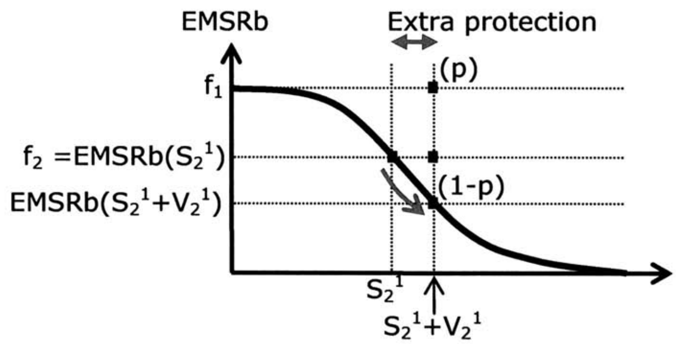

Figure 19: Additional Protection from Sell-up Accounting in the EMSR Optimization

The method had both benefits and drawbacks. Form Gorin and Bohutinsky, it appears that this simple algorithm allows for an increase in revenue by 2% with respect to traditional EMSRb algorithm. It depends strongly on the fare ratio between consecutive fare classes; sell-up rate and demand distribution of the higher fare class, but there were at this point no estimation of the sell-up rate, which is critical to the algorithm. As we will see, the algorithms we will refer to in this thesis are based on sell-up and elasticity and the estimation part is one of the critical and difficult ones.

3.3.2. New Methods to Address Undifferentiated Environments

To respond to the introduction of fare products with almost no restriction nor advance purchase requirements, legacy carriers typically match their competitors. But in this different environment, the fundamental assumption (independence of demand by fare class) on which the traditional revenue management methods are based is largely violated. At this point, airlines must develop new methods adapted to this new environment, among which:

- ✓ Basic Load Factor based control: Threshold Revenue Management methods
- ✓ Modified version of leg-based dynamic programming

- ✓ Q-Forecasting
- ✓ Fare Adjustment methods

This thesis will focus on the last two alternatives, Q-Forecasting and Fare Adjustment methods. Dynamic Programming will not be used in this thesis. The threshold revenue management methods – which are used by the "new entrant" in PODS - are simple enough to be considered as self-explaining (see sub-section 3.1.5 for a more detailed description).

3.4. Q-Forecasting

3.4.1. Overview of Processes and Concepts

Legacy carriers respond to the entrance of a low-cost carrier in a given set of markets with significant change in their fare structure in these markets to match the new entrant fare products, which tend to be less differentiated. Q forecasting is designed to answer to the fact that we cannot rely anymore on the "independence between fare classes demand" assumption, which is especially true in an undifferentiated environment. The lower differentiation there is between fare products, the more interdependence between fare classes. This leads to the concept of "spiral down": as samples go, demand is underestimated and revenues drop:

- ✓ under undifferentiated fare structure, customers buy the lowest fare available
- ✓ associated revenues decrease
- ✓ this "buy-down" phenomenon is recorded as lower fare demand and thus the forecaster predicts less higher class demand
- ✓ the optimizer decreases protection for higher fare classes (since that demand has shifted to lower classes)
- ✓ this allows even more buy-down, and the spiral-down continues

Current Revenue Management methods are not adapted to the undifferentiated environment since that they are based on the segmentation of demand achieved via advance purchase requirements and restrictions and since that they assume independence of demand by fare classes.

Belobaba and Hopperstad[39] have developed modified forecasting and optimization approaches that do not rely on independent class demand. The basic idea is therefore to forecast only total demand at lowest class and to account for the passengers' willingness to pay higher fares. The forecaster will predict an expected number of "Q-equivalent passengers", meaning the demand for the lowest class (in case only this class is available). The aim is still to use a fare class structure to set the booking limits, but demand and bookings are transformed to Q-equivalent bookings and re-partitioned from total equivalent demand at lowest fare to demand

[39] BELOBABA and HOPPERSTAD, *Algorithms for Revenue Management in Unrestricted Fare Markets*, AGIFORS, Auckland, New Zealand, Jan 2004

by fare class. The transition between fare class demand and Q-equivalent demand is based on the Willingness-To-Pay, and will be described later. *Figure 20* presents an overview of the process.

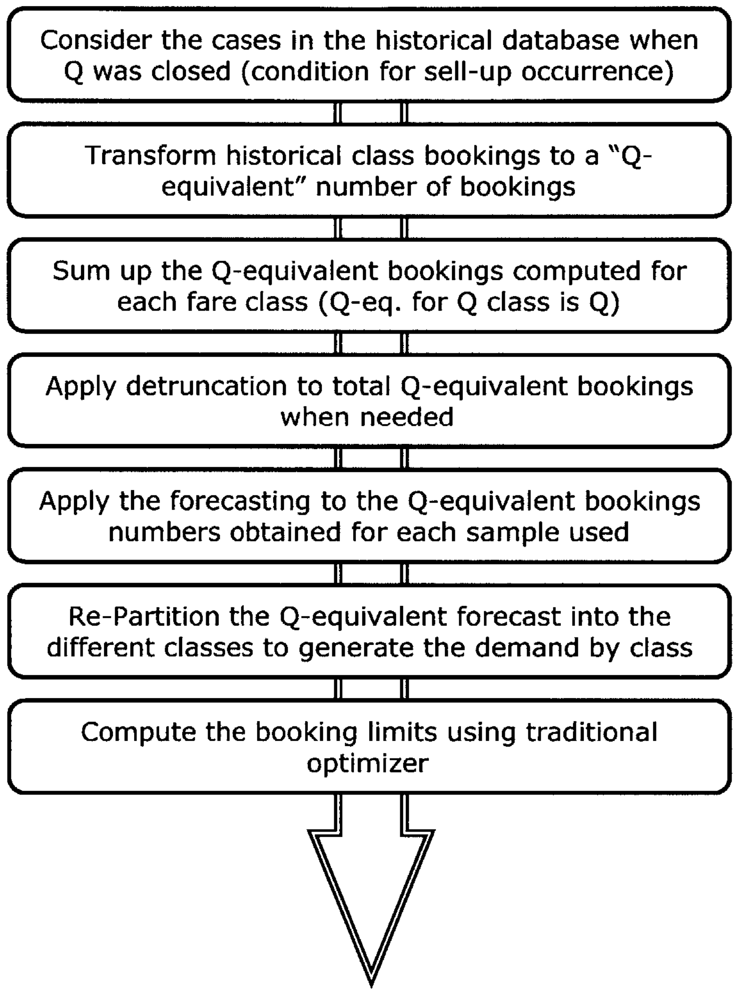

Figure 20: Q-Forecasting Process Overview

3.4.2. The Sell-up Formulation

PODS makes use of an exponential distribution for the probability of sell-up. The probability $p\,\mathrm{sup}_{Q\rightarrow f}$ that "a random passenger would sell-up to class f, given he would have chosen the lower class (called here class Q) and that no other class is available"[40] is:

[40] HOPPERSTAD, 2004

$$p\sup_{Q\to f} = e^{-\left(\frac{fare_f}{fare_Q}-1\right).\sup con_{tf}}$$

Where:

supcon$_{tf}$ is the sell-up constant for time frame *tf*

fare$_i$ represents the fare of class *i*

The airline in PODS has to specify as input a value for *supcon*$_{tf}$ in terms of the fare ratio *Frat5*$_{tf}$ (relative to lowest fare *Q*) at which 50% of the passengers are expected to sell-up for each Time-Frame. The values *supcon*$_{tf}$ and *Frat5*$_{tf}$ are linked by the following formula.

$$\sup con_{tf} = \frac{-\ln\left(\frac{1}{2}\right)}{Frat5_{tf}-1}$$

Typical values for Frat5 exhibit an S-shape reflecting the change in the mix business/leisure across time frames as shown on *Figure 21*.

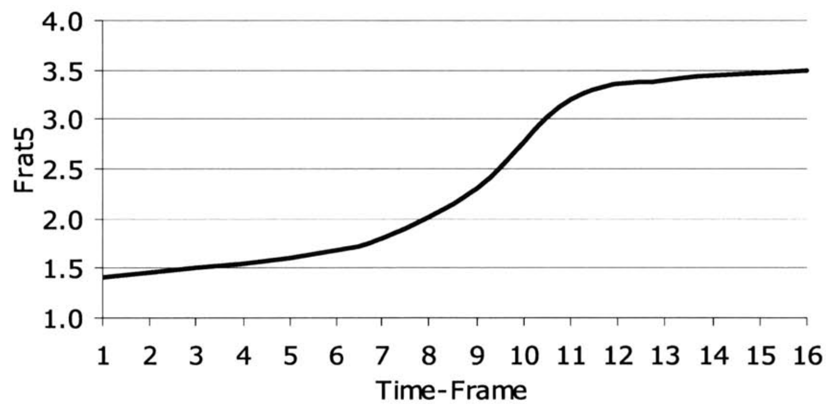

Figure 21: Typical Frat5 Curve

The Frat5 are indeed expected to increase during the booking process for a given flight. Early in the booking process, 50% of the customers (leisure passengers in majority) would only be willing to sell-up from a Q fare of $100 to an M fare of $140. This implies a Frat5 value of 1.4. On the contrary, at the end of the booking process, it might be that 50% of the customers (business passengers in majority) would be willing to sell-up from a Q fare of $100 to a Y fare of $350. This implies a Frat5 value of 3.5.

On *Figure 22* are shown the graphs of *psup*$_{Q\text{->}f}$ as a function of the fare ratio for 3 different values of the *sell-up constant supcon*$_{tf}$, computed from the previous Frat5 graph respectively at Time-Frames 1, 5 and 7, using the above formula.

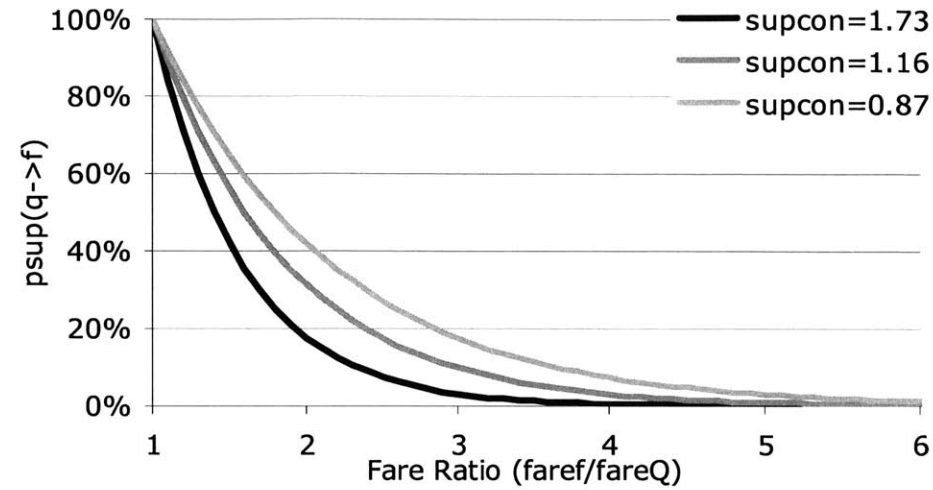

Figure 22: Probability of Sell-up Curves as a function of the fare ratio

As we can notice, the negative exponential tends faster toward zero as the $supcon_{tf}$ is higher.

The sell-up probabilities by fare class (assuming the fares described in *Table 5* and the above Frat5) as a function of Time-Frame are presented in *Figure 23*.

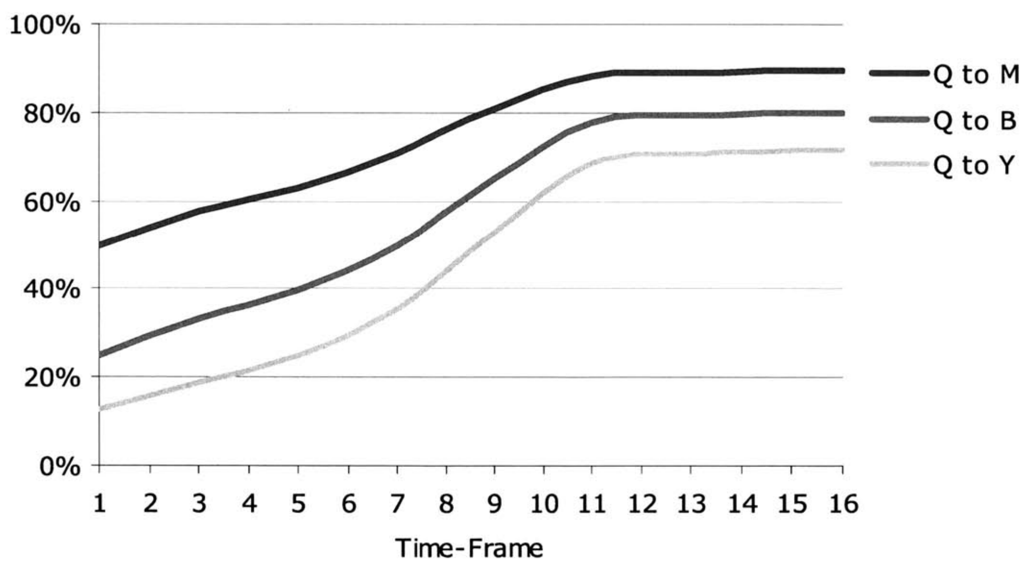

Figure 23: Probability of Sell-up as a Function of Time-Frame

The classes are ranked from the highest to the lowest class as follow: Y, B, M and Q. As expected, at any given TF, the probability of sell-up from the lowest class (Q) to class i decreases when i decreases (e.g. up to the 1st class = Y class).

Modification to the forecaster:
As suggested above, the standard forecasting and detruncation are used, but only applied to the total Q-equivalent demand. For each time-frame, all historical bookings hbk_f for all classes f are converted to Q-equivalent bookings given estimated sell-up rates $psup_{Q\to f}$ as inputs according to the straight forward formula:

$$hbk_Q = \sum_f \frac{hbk_f}{p\sup_{Q\to f}}$$

The *Table 7* is an example of transformation of historical bookings records to Q-equivalent bookings:

Class f	Bookings hbk_f	Sell-up $p\sup_{Q\to f}$	$\frac{hbk_f}{p\sup_{Q\to f}}$
Y	5	15%	33
B	10	40%	25
M	20	70%	28
Q	30	100%	30
		hbk_Q	**116**

Table 7: Scale-Down Computation for a 4-fare classes structure

Then the standard forecasting (exponential smoothing, linear regression, pick-up…) and detruncation (booking curve, projection…) are applied to the hbk_q, providing the mean ($fc\mu_Q$) and the standard deviation ($fc\sigma_Q$) of the forecasted Q-equivalent demand.

Notice (the terminology will be used again later in the thesis) that for detruncation in PODS, a path is marked as *closed* only if the highest class (generally Y) is closed.

Finally comes the optimization process. In order to be able to use the same optimizer as in a traditional process, we need to convert the forecasted Q-equivalent into an n-fare-classes structure. This is done by reversing the previous process. Under the assumption that passengers arrive in reverse order relative to their WTP, the approach is to re-partition the Q-equivalent demand into class demand based on the probability that a random passenger will sell-up to f but not to f-1. In PODS, the corresponding revised mean and standard deviation of demand for class f are computed from:

$$fc\mu'_f = fc\mu_Q \cdot \left(p\sup_{Q\to f} - p\sup_{Q\to f-1}\right)$$

$$fc\sigma'_f = fc\sigma_Q\sqrt{p\sup_{Q\to f} - p\sup_{Q\to f-1}} = \frac{fc\sigma_Q}{\sqrt{fc\mu_Q}}\cdot\sqrt{fc\mu'_f}$$

The new mean for class *f* demand will be proportional to the Q-equivalent forecast and to the probability of sell-up to this specific class. The standard deviation can be written as a function of the repartitioned demand for the same class, and the Q-equivalent demand forecast mean.

Then, for instance, the standard EMSRb optimizer provides the booking limits corresponding to these values.

Table 8 shows the re-partitioning based on the values used in the previous example (*Table 7*) for a specific departure, assuming the forecasted demand obtained across time frames was 125.

Class f	Sell-up Proba $p\sup_{Q\to f}$	125*($p\sup_{Q\to f}$ - $p\sup_{Q\to f-1}$)
Y	14%	125*(0.14-0.00) = 17.50
B	42%	125*(0.42-0.14) = 35.00
M	75%	125*(0.75-0.42) = 41.25
Q	100%	125*(1.00-0.75) = 31.25
	tot	125

Table 8: Re-partitioning Computation for a 4 Fare Classes Structure

3.4.3. Sell-up Parameters in PODS

The implementing in PODS is an expanded version of the theory presented above. In PODS, 2 sets of Frat5 are used for each airline:

- ✓ Time-Frame Frat5 (TF-Frat5), which are used to transform historical booking records (occurring during the Time Frame) to Q-equivalent bookings
- ✓ Bookings-To-Come Frat5 (BTC-Frat5), which are used to re-convert Q-equivalent forecasts to partitioned class forecasts (and therefore need to be "Bookings-To-Come" values)

The difference is that the BTC-Frat5 value for TF_i refers to the remaining booking period, whereas the TF-Frat5 value for TF_i is relative to the current Time-Frame. The BTC-Frat5 for TF_i can be derived (in the input sets used later in the simulations) from the TF-Frat5 set, for instance as the average of the TF-Frat5 for TF_i to TF_{16}.

As well, the use of Q-forecasting can be turned on or off for each airline in a given set of markets: when used, it is turned on in the markets with undifferentiated fare product structure.

In order to conduct sensitivity analyses relative to the input Frat5, we have arbitrarily come up with several sets of TF-Frat5 respecting the "S" shape, but with decreasing aggressiveness (Set "A" is more aggressive than set "B", meaning that the Booking Limits corresponding to the set "A" will protect more seats for the higher classes). Sets "A" to "E" are shown on *Figure 24*. The derived BTC-Frat5 sets are presented on *Figure 25*.

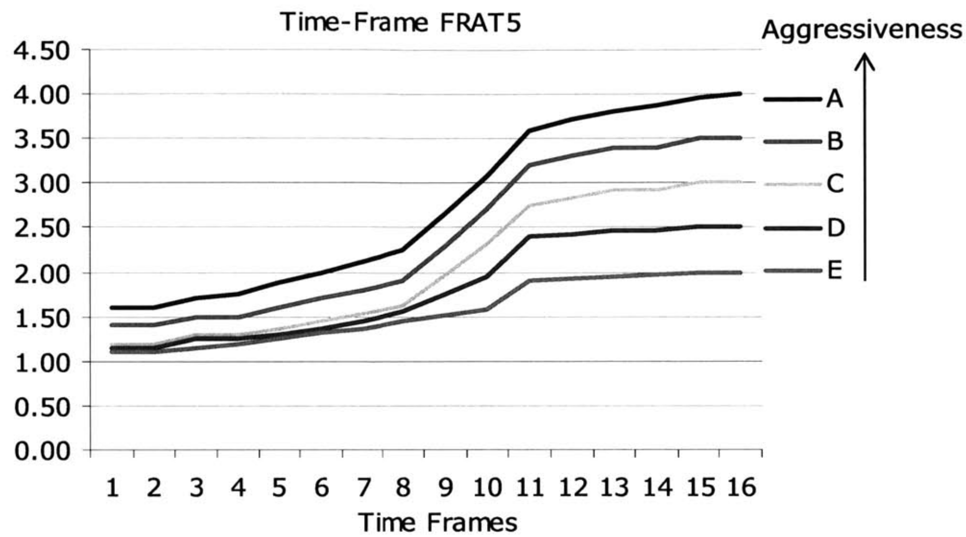

Figure 24: Time-Frame Frat5 as a Function of Time-Frames

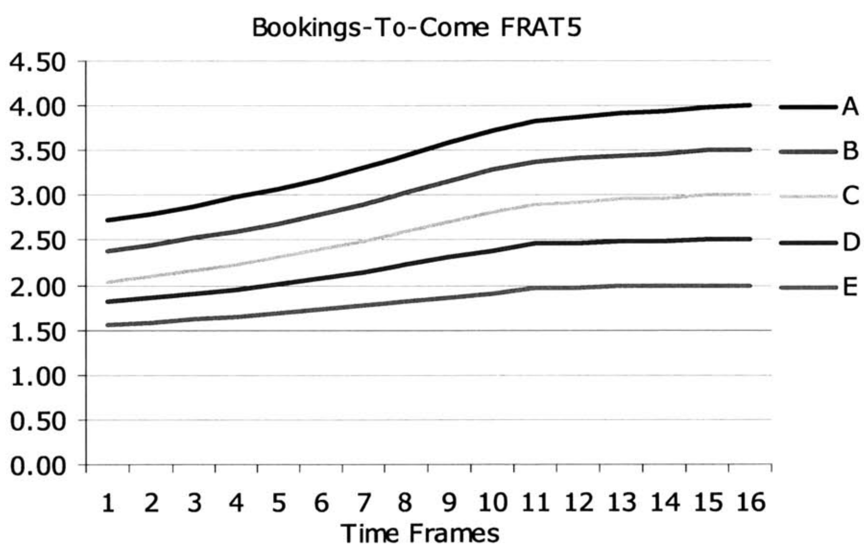

Figure 25: Bookings-To-Come Frat5 as a Function of Time-Frames

Several other input-Frat5 sets will be described and presented further in the thesis, but these ones will be extensively used. The next step, described in the next sub-section, will be to manage to correctly estimate the Frat5 values dynamically from the historical booking records in order to make the Q-Forecasting process independent from the input and more feasible for the airlines (the Frat5 values would theoretically be different every day, and for every given itinerary).

3.4.4. The Sell-up Estimator

As previously introduced, the sell-up estimator is designed to provide inputs dynamically to the Q-Forecasting method. Q-Forecasting with input Frat5 as described in the previous paragraph would require knowing and changing the frat5 values of each flight at each time-frame, which is unrealistic. The goal is therefore to add to the method a feature which estimates dynamically, as part of the forecasting process, based on historical booking records, the values of the Frat5.

The method described below has been introduced by **Hopperstad**[41]. The process is presented on *Figure 26.* Each step shown on this figure will be described in detail below.

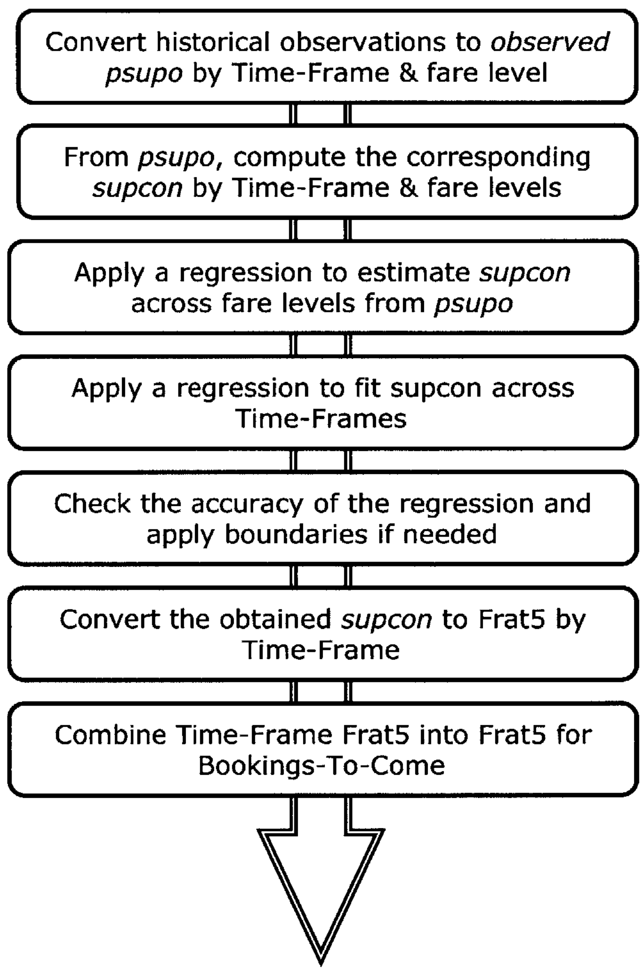

Figure 26: Sell-up Estimation Process

In the estimation process, a scale of fare ratios is defined. For each relevant path, the fare ratio fr_i are defined for class i in terms of the ratio of the class i fare to the Q

[41] HOPPERSTAD, 2004

fare (e.g. the lowest fare). The objective is then to estimate the probabilities of sell-up *psup* for each level by time frame, apply a series of consecutive regressions and to report them as Frat5 values.

The computation of the observed probability of sell-up, *psupo*, is part of the forecast prediction method. The estimation process is applied *only* to cases when the *Q class was closed*. For each Time-Frame, and each level corresponding to unclosed flights, the estimator records:

- ✓ The number Bookings which occurred during the Time-Frame *tf* for the lowest class available (not Q, by assumption)
- ✓ The Q-equivalent forecast for the considered Time-Frame. Notice that Time-Frame forecasts (*fbtf*) are obtained from Bookings-To-Come forecasts (*fbtc*) using estimated booking curves (*pbook*)

$$fbtf = \frac{fbtc \cdot \left(pbook_{tf} - pbook_{tf-1}\right)}{1 - pbook_{tf-1}}$$

As a reminder, $pbook_{tf}$ represents the ratio of the number of bookings recorded until Time-Frame *tf* over the total number of bookings at day-of-departure. ($pbook_{tf}$-$pbook_{tf-1}$) represents the ratio of expected bookings during tf over the total number of bookings, and the quotient $(pbook_{tf} - pbook_{tf-1})/(1 - pbook_{tf-1})$ represents the proportion of remaining bookings which occur during Time-Frame *tf.*

The observed probability of sell-up, *psupo*, is defined for each Time-Frame and each fare level, as:

$$p\sup o_{f,tf} = \frac{bookings_{f,tf}}{Q_{eq}\,forecast_{tf}}$$

Example:
Suppose we have 4 fare classes, Y, B, M and Q. For a given observation (Time-Frame *tf*), Q is closed but not M. During Time-Frame *tf*, 10 bookings have been recorded. The Q equivalent forecast of the same Time-Frame was 30. Therefore, the observed probability of sell-up for *tf* was: $psupo_{tf}$=10/30=33%.

The idea of the forecast prediction method is that the sum of the Q-equivalent forecasts is obtained owing to previous estimates of sell-up in previous samples. The consecutive recalculations of *psupo* constitute a series of corrections and convergence is achieved quickly.

At this point, we have a series of $psupo_{f,tf}$. The next step is to estimate a corresponding value of $supcon_{tf}$ which do not depend on the fare levels. Knowing the assumed relation between $psup_{f,tf}$ and $supcon_{tf}$:

$$p\sup_{f,tf}\left(fr_f\right) = e^{-\sup con_{tf}\left(fr_f - 1\right)}$$

The Time-Frame *supcon* fit is done by weighted least squares fit (for a given Time-Frame) using levels for which there were forecasts and bookings. The method consists of picking $supcon_{tf}$ and a scaling constant *c* such that:

$$\sum_f \left(p\sup o_{f,tf} - c \cdot e^{-\sup con_{tf}(fr_f - 1)}\right)^2 \text{ is minimized}$$

Where: $psupo_f$ = observed sell-up probability, class f
fr_f = fare ratio, class f

The obtained $supcon_{tf}$ is now independent from the fare level. The following step deals with checking the obtained $supcon_{tf}$ values across Time-Frames by applying a linear regression fit. We pick an intercept b and a slope a such that:

$$\sum_{tf} \left(b + a \cdot tf - \sup con_{tf}\right)^2 \text{ is minimized}$$

Then a correction is added to prevent negative values:

$$\sup con'_{tf} = \max\left[0.2, b + a \cdot tf\right]$$

From this value of the $supcon'_{tf}$, the Time-Frame Frat5 values can be derived using the following formula.

$$\sup con_{tf} = \frac{-\ln\left(\frac{1}{2}\right)}{Frat5_{tf} - 1}$$

Finally, using again the estimated booking curves, we compute the Bookings-To-Come values of the $psup'_{f,tf}$ as a weighted average across Time-Frames of the remaining Time-Frame values of the $psup'_{f,tf}$. The weights correspond to the proportions of bookings expected during each Time-Frame.

$$p\sup'_{f,tf,BTC} = \frac{\sum_{n=tf}^{nft} p\sup'_{f,tf,BITC}\left(pbook_n - pbook_{n-1}\right)}{1 - pbook_{tf-1}}$$

From these probabilities of sell-up, the Bookings-To-Come Frat5 are then computed.

3.5. Fare Adjustment Methods

The Fare Adjustment methods have been developed by **Fiig and Isler**[42] in 2004 and then implemented in PODS by Hopperstad. The main point of the methods is to define an alternative capable to address the problem of the coexistence of different fare structures in a single leg. Typically, one of them is a traditional, restricted fare structure, whereas the other one is an unrestricted, low-cost-carrier-type fare structure.

For instance, in PODS' Network D, the New Entrant airline (dedicated to be the low-cost carrier) enters 10 non-stop markets. In these local markets, the legacy carrier with which the New Entrant is competing has to match the fare structure for the local passengers. On the contrary, the connecting passengers (for the network carrier) flying on this particular leg are still in a traditional, restricted fare product environment. Both structures co-exist on the same leg.

There are 2 distinct fare adjustment methods proposed for DAVN, so called MR (for Marginal Revenue) and KI (for Karl Isler). The only difference is that MR is a continuous method whereas KI is a discrete method. Apart from this, both are based on the same concepts and will be presented simultaneously.

In this section, we will focus first on explaining the conflict resulting from the presence of 2 different fare structures in the same leg for a given airline. Then the overview will introduce the main mechanisms of the methods, which will be described in the next section. Finally, the differences in terms of sell-up definition between Fare Adjustment methods (FA) and Q-Forecasting methods (QF) will be explained.

3.5.1. Overview of the Process

To understand the concept, it is important to understand the problem it addresses. The co-existence of several distinct fare structures in the same leg. Let's consider again the example of PODS' Network D with the New Entrant Airline (LCC) in a set of 10 local markets.

Under the restricted fare structure, demand for each fare class is assumed to be independent and sell-up potential is not taken into account by the forecaster. Thus, demand is segmented thanks to restrictions and advance purchase requirements. On the contrary, under the unrestricted fare structure, revenues are maximized by incorporating sell-up behavior into forecasting and optimization. In this case, higher-class bookings depend entirely on closing down lower classes.

Therefore, in DAVN for instance, it may happen that in the same virtual bucket, 2 OD Fares co-exist, each of them belonging to a different market and fare structure. Then, the closure down of the bucket, which automatically make both fares unavailable, may be optimal under the restricted fare structure and not at all under

[42] FIIG and ISLER, *SAS O&D Low-Cost Project*, PODS, Minneapolis-St Paul, 2004

the unrestricted fare structure (respectively, may the closure down of the same bucket may be optimal under the unrestricted fare structure and not at all under the restricted fare structure). Thus, revenues will not be maximized. On *Figure 27* is described such a conflict.

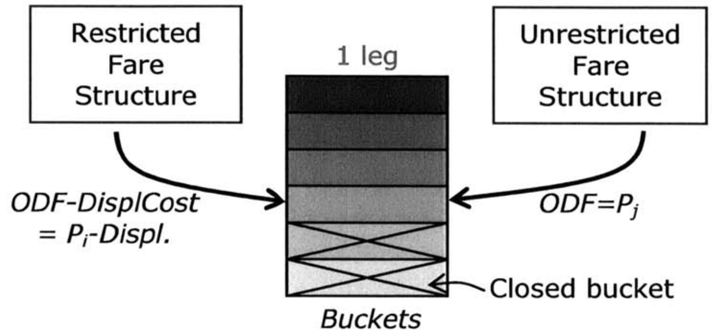

Figure 27: Co-existence of distinct fare structures on the same leg

As a reminder, DAVN uses virtual buckets in which it orders the network revenue contributions, which are equal to the OD-Fare minus the displacement cost (cf Section 3.1.5). In the case considered here, the unrestricted fare structure applies to local passengers, that is why there is no displacement cost subtracted from the OD Fare P_j on the right hand side. If now the third bucket (starting from the bottom) is closed down, both P_j and P_i are made unavailable, which might, as suggested above, not be the optimal strategy for both structures.

The principle of the fare adjustment methods proposed by the Fare Adjustment methods is, instead of optimizing the quantities "ODFare - Displacement Cost", to optimize the quantities "Marginal Revenue - Displacement Cost". Given the specific demand curve used in PODS (negative exponential), it also means (by hazard) to subtract to the quantity "ODF - Displacement Cost" another amount, so-called Price-Elasticity Cost (PEcost) which, in a sense, accounts for the risk of buy-down.

The effect of this method is to decrease the displacement-adjusted fares of the unrestricted fare structures and send the corresponding fares to lower buckets. The PEcost, which is not applied to the restricted fare structures, will allow decoupling the fares which previously were in the same buckets and allow managing both structures more independently, as shown on the *Figure 28*.

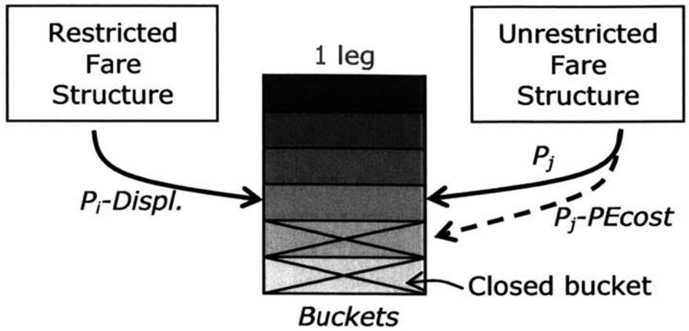

Figure 28: Decoupling the Fare Structures

In this example, we replace in the bucketing process the fare P_j from the unrestricted fare structure by the adjusted fare P_j-PEcost. The fare is sent to a lower bucket which is now closed at this point of the booking process, whereas the fare P_i is still available.

The generalized formula for bucketing (for all legs and all kinds of markets) in the PODS version of DAVN is now:

OD Fare – Displacement Cost - PEcost

The insertion of the MR process into DAVN is summarized on *Figure 29*.

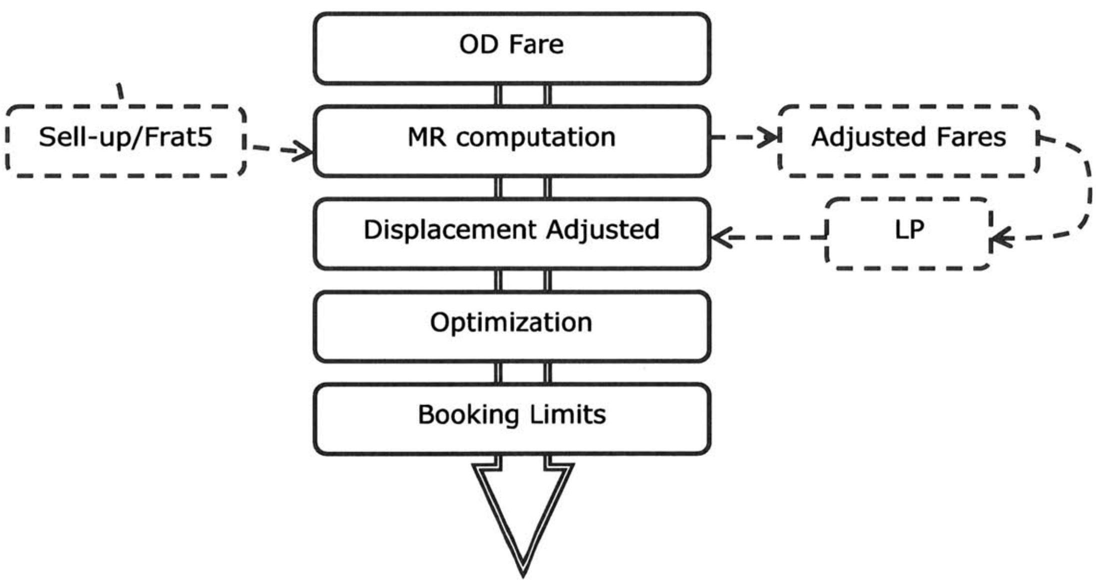

Figure 29: Insertion of MR into the DAVN Process

The LP which is solved to compute the displacement costs only sees – in the "New Entrant" markets – the adjusted fares equal to the ODF minus the PEcost (it does not "see" at all the ODF, like in the traditional DAVN). MR uses as secondary input the sell-up data from the Frat5 sets.

3.5.2. Formulation

Given the following notations:

Q=Demand
P=Price
TR=Total Revenue,

The Total Revenue is computed as follow:

$$TR(Q) = P(Q) \times Q$$

And then, depending on the formulation chosen, continuous (Fiig) or discrete (Isler), the Marginal Revenue is given by:

$$MR_{Fiig} = \frac{\partial TR}{\partial Q} \qquad \text{and} \qquad MR_{Isler} = \frac{TR(Q+\Delta Q) - TR(Q)}{\Delta Q}$$

3.5.3. PODS parameters

In PODS, we assume a negative exponential form for the demand curve. Therefore, the Price Elasticity follows the formula:

$$PE = e^{\ln\left(\frac{1}{2}\right) \times \frac{\frac{p}{p_0} - 1}{Frat5 - 1}}$$

And the PEcost:

$$PE\cos t = p_0 \cdot \frac{Frat5 - 1}{-\ln\left(\frac{1}{2}\right)}$$

On the *Figure 30* are represented the Marginal Revenue curves as a function of the demand Q, assuming a negative exponential form. The PEcost can be shown as the difference between a given fare and the Marginal Revenue associated with the corresponding demand value. Given the configuration, the exhibited Discrete PEcost is in this case smaller than the Continuous PEcost.

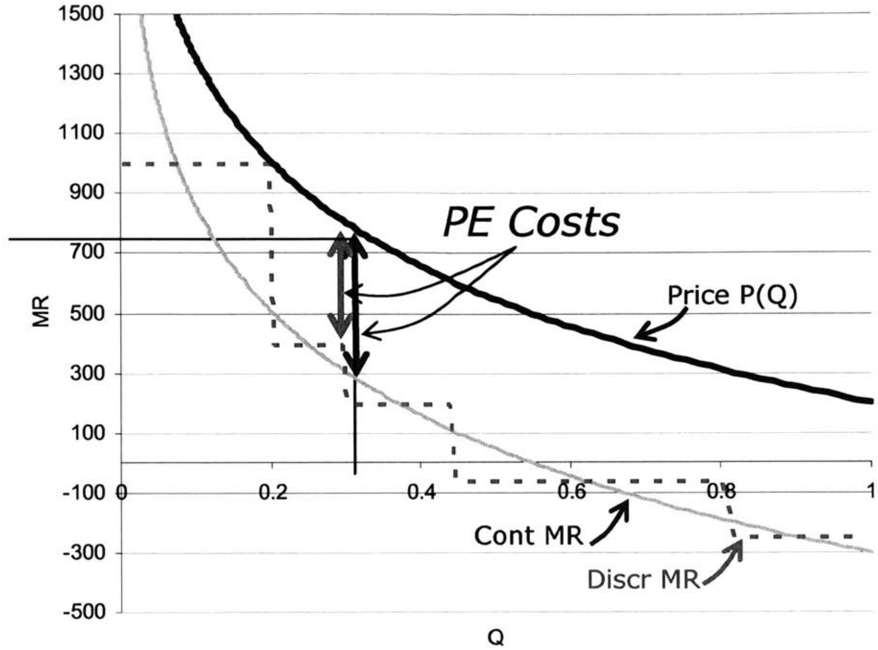

Figure 30: Representation of the MR curves as a function of Q

From the previous graph, we can notice that we do not want to sell any ticket at a fare lower than the one corresponding to a Marginal Revenue equal to zero (here for Q=0.55 and Price=\$500). Like in the case of Q-Forecasting, the sell-up rates are given for the Fare Adjustment methods in terms of Frat5. The sets of Frat5 used for these methods are different from the Q-Forecasting ones for the reasons explained below.

3.5.4. Differences with Q-Forecasting

Sell-up is not used in the same way by both methods:

- ✓ Frat5 for Q-Forecasting are *conditional*, they represent buy-up rates given that lower classes are closed
- ✓ Frat5 for Fare Adjustment methods are *unconditional*, they are applied to all booking requests regardless of the open/closure status of the different classes

Let's consider the following example by **Fiig**[43]: assume we have two airlines A and B in competition, and two classes, Y and Q, with respective fares f_Y and f_Q. Let 0/1 denote if a booking class is closed/open. The system can be represented by the states (we only consider the cases when Y is open) represented on *Figure 31*.

[43] FIIG, *Memo Price Elasticity*, SAS internal document, 2005

$$
\begin{array}{c} \\ Y \\ Q \end{array}
\begin{array}{c} A \quad B \\ \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \end{array}
\quad
\begin{array}{c} A \quad B \\ \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \end{array}
\quad
\begin{array}{c} A \quad B \\ \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} \end{array}
\quad
\begin{array}{c} A \quad B \\ \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \end{array}
$$

Figure 31: States of the system

In the real world, airline A knows only the parameters of the first column in each matrix and the sell-up rate from Q to Y when Q is closed. In this example, we will assume that the airline also knows the availability of its competitor (e.g. the second column). In addition, as shown on *Figure 32*, for each state of the system, we know the probability of sell-up *psup* and the frequency at which the state occurs (given the RM controls of A and B).

$$
\begin{array}{cccc}
 & A \quad B & psup & frequency \\
\begin{array}{c} Y \\ Q \end{array} & \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} & 0 & 0.05 \\
\begin{array}{c} Y \\ Q \end{array} & \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} & 0 & 0.10 \\
\begin{array}{c} Y \\ Q \end{array} & \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} & q & 0.05 \\
\begin{array}{c} Y \\ Q \end{array} & \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} & 0 & 0.08
\end{array}
$$

Figure 32: Sell-up data

There is only one case where sell-up can occur, that is when no more Q is available, neither in Airline A's inventory nor in Airline B's inventory. In all other cases, passengers carry on buying Q seats as long as there are some available.

Under these assumptions and notations, the price elasticity for Q-Forecasting is given by:

$$
p\left(Y_A \mid Q_A closed\ history\right) = \left(\frac{\sum p\sup_i \times freq_i when\ Q_A closed}{\sum freq_i when\ Q_A closed} \right) = \frac{q \times 0.05}{0.05 + 0.10} = \frac{q}{3}
$$

On contrary, the price elasticity which will be used in the fare adjustment methods is estimated from the Marginal Revenue (discrete formulation here):

$$
MR = \frac{TR(Q\ open) - TR(Q\ closed)}{PAX(Q\ open) - PAX(Q\ closed)},
$$

where TR is the Total Revenue and PAX the number of Passengers in a given class.

If Q_A is open, the Total Revenue can be written as:

$$TR(Q_A open) = PAX(Q_A open) \cdot f_{Q_A},$$

where PAX($Q_A$open) represents the number of passengers who buy a seat to A when Q_A is open.

If Q_A is closed, the Total Revenue is:

$$TR(Q_A closed) = PAX(Q_A closed) \cdot f_{Y_A},$$

where PAX($Q_A$closed) represents the number of passengers who buy a seat to A when Q_A is closed (so they buy a Y class ticket). It can be written as:

$$PAX(Q_A closed) = PAX(Q_A open) \times \Pr(Y \mid Q_A closed\ \ future)$$

We add the notation: $p\sup = \Pr(Y \mid Q_A closed\ \ future)$

So when Q_A is closed,

$$TR(Q_A closed) = PAX(Q_A open) \times p\sup \cdot f_{Y_A}$$

Finally, by inserting these formulas into the MR equation, we obtain:

$$MR = \frac{f_{Q_A} - p\sup \times f_{Y_A}}{1 - p\sup}$$

The sell-up rate here is not equal to the historical sell-up rate computed in the Q-Forecasting case. The computation of the sell-up rate then depends on the degree of knowledge of the competitor's availability.

If we assume that the competitor RMS is completely independent of A's, the historical frequencies for the 4 states will be the future frequencies as well. If we apply the method to the example considered above, we obtain the results presented on *Figure 33*.

	HISTORY			*Q-CLOSED*			*FUTURE*		
	A B	*Psup*	*Freq*		*Psup*	*Freq*		*Psup*	*Freq*
$\begin{array}{c}Y\\Q\end{array}$	$\begin{pmatrix}1&1\\1&0\end{pmatrix}$	0	0.05	$\begin{pmatrix}1&1\\1&0\end{pmatrix} \Rightarrow \begin{pmatrix}1&1\\0&0\end{pmatrix}$	q	0.05	$\begin{pmatrix}1&1\\0&0\end{pmatrix}$	q	0.10
$\begin{array}{c}Y\\Q\end{array}$	$\begin{pmatrix}1&1\\0&1\end{pmatrix}$	0	0.10	$\begin{pmatrix}1&1\\0&1\end{pmatrix} \Rightarrow \begin{pmatrix}1&1\\0&1\end{pmatrix}$	0	0.10	$\begin{pmatrix}1&1\\0&1\end{pmatrix}$	0	0.90
$\begin{array}{c}Y\\Q\end{array}$	$\begin{pmatrix}1&1\\0&0\end{pmatrix}$	q	0.05	$\begin{pmatrix}1&1\\0&0\end{pmatrix} \Rightarrow \begin{pmatrix}1&1\\0&0\end{pmatrix}$	q	0.05			
$\begin{array}{c}Y\\Q\end{array}$	$\begin{pmatrix}1&1\\1&1\end{pmatrix}$	0	0.08	$\begin{pmatrix}1&1\\1&1\end{pmatrix} \Rightarrow \begin{pmatrix}1&1\\0&1\end{pmatrix}$	0	0.80			

Figure 33: Sell-up rate when the RMS are independent

In this case, the future sell-up rate is equal to q*0.1, which is very different from the historical sell-up rate computed in the case of Q-Forecasting (it was equal to q/3).

If now we assume that both RMS are very highly correlated (the other extreme case), it means that the competitor B mimics the RMS of Airline A. Consequently, both airlines will have the same availability. Applying the concept to the same example, the results are shown on *Figure 34*.

	HISTORY			*Q-CLOSED (and B follows)*			*FUTURE*		
	A B	*Psup*	*Freq*		*Psup*	*Freq*		*Psup*	*Freq*
Y Q	$\begin{pmatrix}1 & 1\\ 1 & 0\end{pmatrix}$	0	0.05	$\begin{pmatrix}1 & 1\\ 1 & 0\end{pmatrix} \Rightarrow \begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix}$	q	0.05	$\begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix}$	q	0.90
Y Q	$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}$	0	0.10	$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix} \Rightarrow \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}$	0	0.10	$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}$	0	0.10
Y Q	$\begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix}$	q	0.05	$\begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix} \Rightarrow \begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix}$	q	0.05			
Y Q	$\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}$	0	0.08	$\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix} \Rightarrow \begin{pmatrix}1 & 1\\ 0 & 0\end{pmatrix}$	q	0.80			

Figure 34: Sell-up rate when the RMS are highly correlated

In this case, the future sell-up rate will be equal to q*0.9, which is again very different from the historical sell-up rate computed previously.

All degrees of correlation between both airlines are possible, and can be modeled easily. This concludes the section aimed to show that sell-up rates can be different when we deal with Q-Forecasting from when we deal with Fare Adjustment methods. In his Memo, **Fiig**[43] shows as well the correspondence between the discrete and continuous price elasticity.

Both processes – QF and FA – are not totally independent, but are complementary, and we will see in the next section that the airlines benefit from using both simultaneously with DAVN. The insertion of both processes is schematized on *Figure 35*.

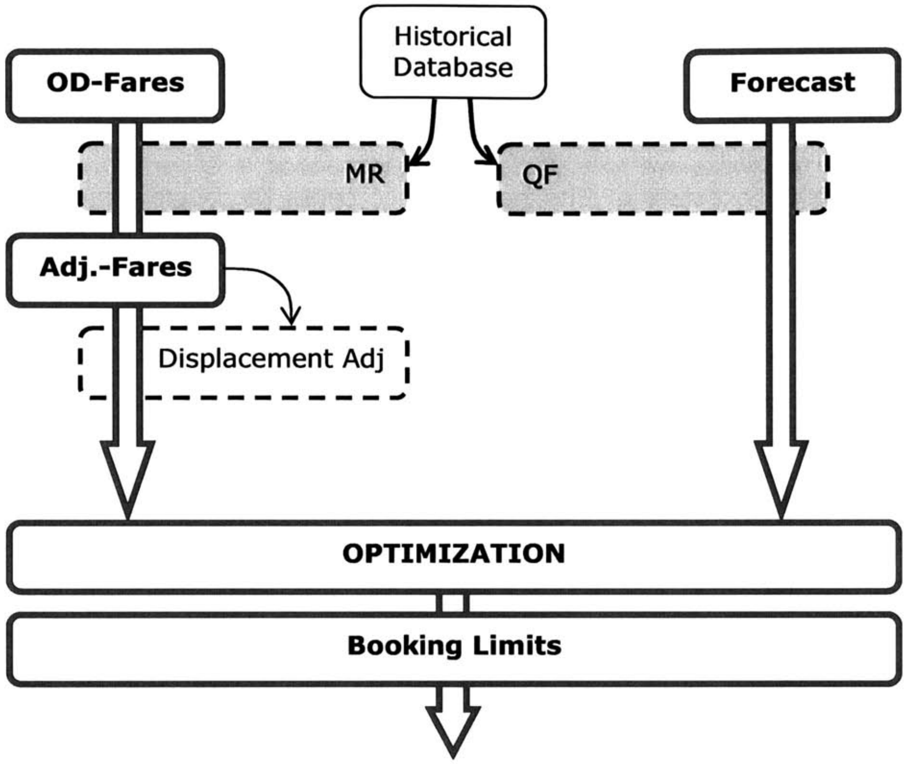

Figure 35: Insertion of MR and QF into the DAVN process

The historical database is used for the forecasting and sell-up estimations, optimally in both Q-Forecasting and Fare Adjustment methods. On the Forecasting side, Q-Forecasting is inserted at the very beginning of the process. On the Fare Adjustment side, the adjustment occurs as well very early in the process, before subtracting the displacement cost since that the LP solved to compute them uses the adjusted fares. Then, the "MR+Displacement adjusted" fares are sent into the optimizer with the forecasts to compute the booking limits.

4. Simulation Results

This chapter is divided in four sections. The first one presents quickly the spiral-down effect in the networks we will use. Then, the section 4.2 summarizes the results obtained with Q-Forecasting. The next one presents the performance of the Fare-Adjustment methods, and the last section details the benefits obtained from the combination of both methods.

4.1. Traditional RM in Unrestricted Environments: Spiral-Down Effect

4.1.1. Performance in the Single Market Case

In order to simulate a sudden change ("shock") in restrictions - from a fully restricted fare product to a fully unrestricted fare product - the burn-in process of PODS (first 200 samples of each trial) was modified. In a specific version, we are allowed to change the fare product starting from the 201st sample so that the new fare product is used for the 400 last samples, which account for the reported results.

This has been tested in the single market case:

- ✓ 2 carriers
- ✓ 3 flights per day
- ✓ Various Demand Factors (DF) corresponding to the following Average Load Factors
  - DF=0.8 means 65% ALF for EMSRb without shock
  - DF=0.9 means 73% ALF for EMSRb without shock
  - DF=1.0 means 77% ALF for EMSRb without shock

We are in the case of 4 fare classes, following the example of *Table 9*.

Class	Y	B	M	Q
Fare	$625	$295	$185	$125

Table 9: 4-Fare-class Price Structure

Both airlines in these cases will use standard advance purchase requirements, and the fare products, before and after the shock, are structured as shown respectively in *Table 10* and *Table 11.*

TF	R1	R2	R3
16	0	0	0
12	1	0	0
10	1	1	0
8	1	1	1

Table 10: Restricted Fare Product in the 4-Fare-Class Case

TF	R1	R2	R3
16	0	0	0
12	0	0	0
10	0	0	0
8	0	0	0

Table 11: Unrestricted Fare Product in the 4-Fare-Class Case

In all cases below, Airline 1 and Airline 2 undergo the shock in their fare products. Given the symmetry in their results, only the performance of Airline 1 will be reported here.

On *Figure 36* is represented Airline 1's revenue (USD) per day (3 flights in the single market case) when the shock occurs, computed for each of the 600 samples of a particular trial. The black curve is a moving average over 30 departure days.

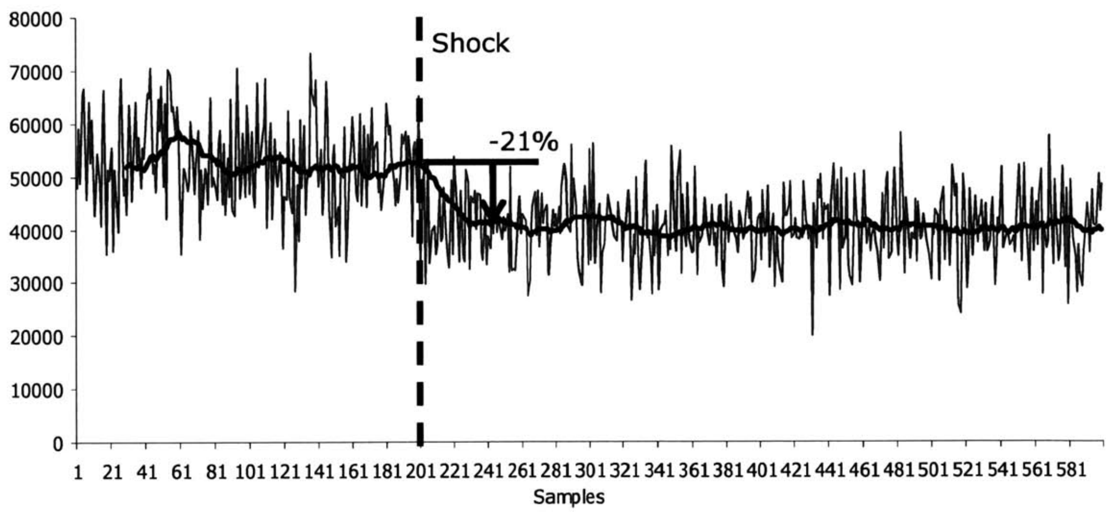

Figure 36: Trend in Airline 1's Revenue when a shock occurs at sample 200

As one can notice, the shock induces a drop in revenue by 21% on average, and the airline does not recover from this decrease in the next samples. On *Figure 37* is shown the revenue of Airline 1 without and with shock, for 3 different values of the Demand Factor.

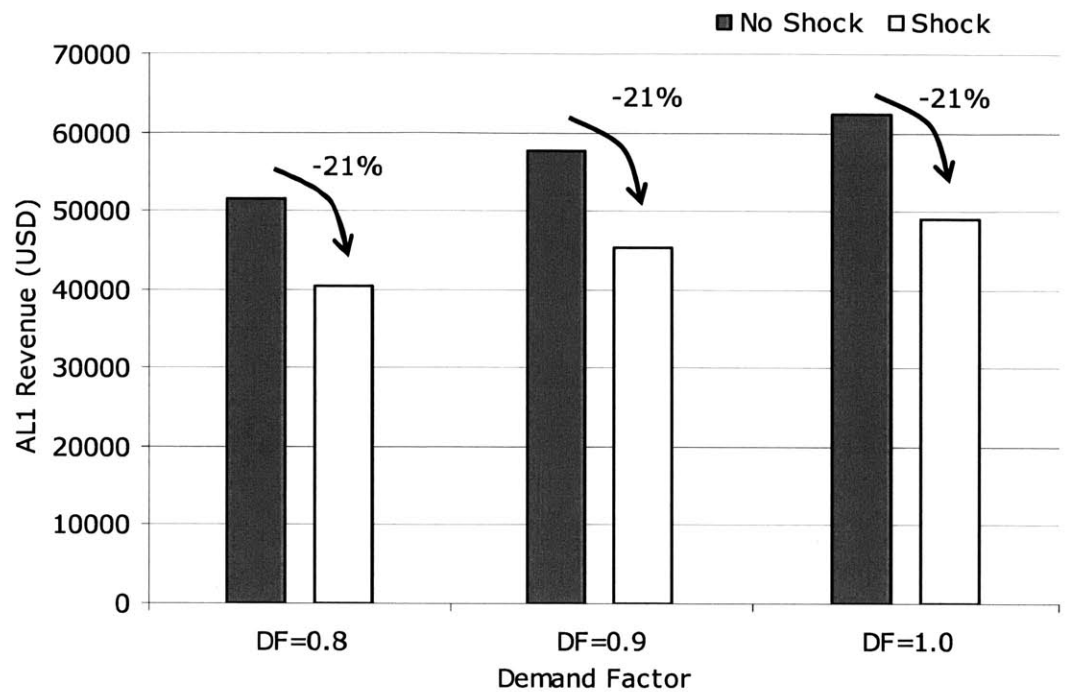

Figure 37: Revenue of AL1 as a Function of Demand Factor, Without and With Shock

Irrespective of the Demand Factor value, the drop in revenue caused by the shock in the restriction represents a decrease by 21%.

Finally, on *Figure 38*, we present a comparison of various cases: EMSRb vs. EMSRb without shock, EMSRb vs. EMSRb with shock, FCFS vs. FCFS with shock, and EMSRb vs EMSRb with shock and the Belobaba-Weatherford Sell-up algorithm ("SU" on the figure) we have mentioned in Chapter 3.

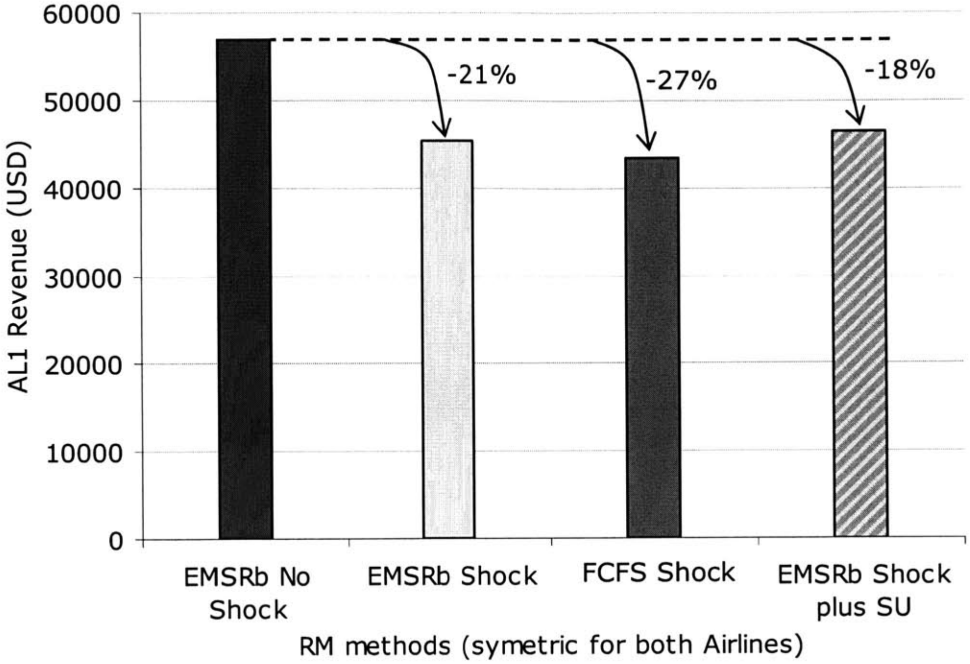

Figure 38: Revenue of AL1 With Various RM Methods

Again, we notice the decrease in revenue in the case of EMSRb when the shock occurs (first two bars), but it is interesting to consider the effect of the Belobaba-Weatherford sell-up algorithm, which allows recovering a small amount of the drop in revenue. The decrease is reduced from -21% to -18% when the sell-up algorithm is used. Nevertheless, this is not enough to recover the entire losses.

4.1.2. Performance in Network D

In this section we will briefly mention the results obtained by **Gorin**[44] regarding the removal of restrictions and advance purchase requirements in the environment of the network D when a Low-cost airline enters 10 non-stop markets and competes with the Airline 1 in these local markets.

The environment is that of Network D as described in Chapter 3, with and without new entrant, with removal of restrictions and advance purchase requirements, as well as compression of fare ratios for the 3 airlines in the 10 markets.

Consider the case when both network carriers are using EMSRb. When the LCC enters the 10 markets using EMSRb and the incumbents match restrictions and advance purchase requirements, the revenues of Airline 1 decrease by 8% on average, from $1,235,000 down to $1,135,000. This illustrates and somewhat

[44] Gorin, *Assessing Low-Fare Entry in Airline Markets: Impact of Revenue Management and Network Flows,* PhD Thesis, MIT, Cambridge, MA, 2004

quantifies the spiral-down effect in this particular case. Looking at the local loads in the 10 markets, we notice that bookings actually dropped in Y, B and M classes, whereas Q bookings increased by 140% (buy-down phenomenon due to the removal of the fences between fare products - restrictions and advance purchase requirements).

If we compare with the exact same case except that Airline 1 uses DAVN, we notice that the advantage brought by DAVN over EMSRb is still the same in terms of percent as it was without LCC: revenues increase by 1.54%. In the next sections, DAVN will generally be the reference for the RM methods.

4.2. Performance of Q-Forecasting (QF)

This section is divided in two sub-sections. In the first one, we will analyze the performance of the Q-Forecasting (QF) method based on Frat5 values that we set as input before running PODS. In the second sub-section, we will show the performance of the sell-up estimator, which dynamically estimates the Frat5 values during the run, and compare it with the results obtained in the first sub-section. The use of QF with input Frat5 by an airline will be generally reflected in the graphs with diagonal stripes, while the use of QF with the sell-up estimator will be emphasized by horizontal stripes.

4.2.1. Input

In this entire sub-section, we will only consider the Network D (cf. section 3.1.2) in the "new entrant" configuration introduced above. Airline 1 will be the only airline which will use the Q-Forecasting method (apart from the LCC when it uses EMSRQ2, but we will not comment the effects of the forecasting method on this airline). The Frat5 set used in this subsection is the "B" set (cf. 3.4.3), which is a medium-aggressive set.

Consider the case when AL1 and AL2 are using DAVN, AL3 (LCC) is using EMSRb (very basic RM especially in unrestricted environments), and the Demand Factor is set to 1. On *Figure 39* is shown the impact of the entrance of the LCC in the network on AL1 and AL2's revenues, first when AL1 uses traditional PU forecasting, and then when it uses Q-Forecasting.

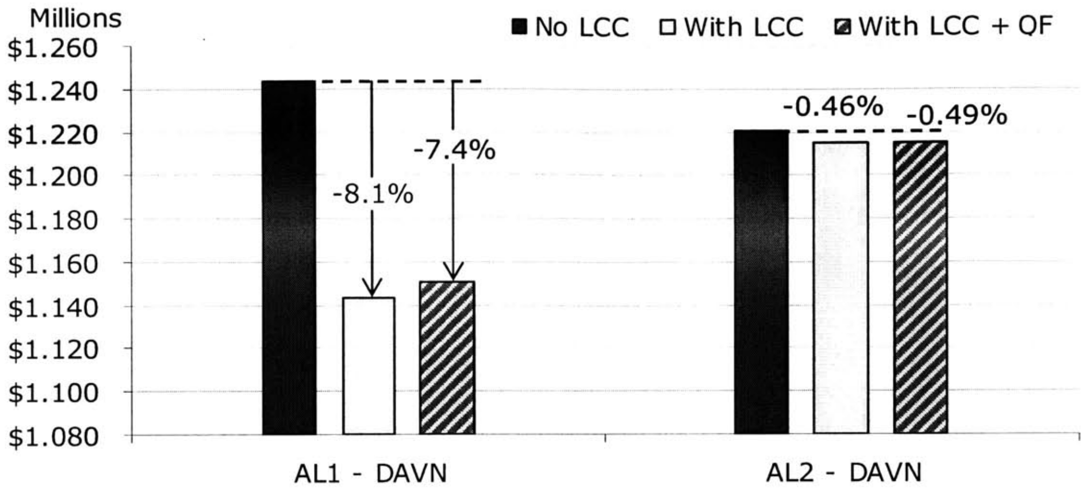

Figure 39: Impact of Q-Forecasting on AL1 and AL2 Revenues

As previously shown, the entrance of the LCC in the 10 markets causes a decrease in revenue for AL1 by 8.1% when it matches its competitor (but still uses Pick-up – PU – forecasting). When AL1 switches to Q-Forecasting, the drop in revenue is reduced to 7.4% with respect to the case when there is no LCC. In other terms, considering only the case of Network D with 3 airlines, when AL1 changes from PU to QF, its revenues increase by 0.62% from $1,143,400 to $1,150,470. If AL2 suffers from the entrance of the LCC (-0.46% in revenues), we see that the impact of AL1 using QF is not really significant (-0.49% in revenues when AL1 switches to QF).

Figure 40 highlights the effect of QF on the average load factors (across their networks) of the 3 airlines, under the same conditions as in the previous graph.

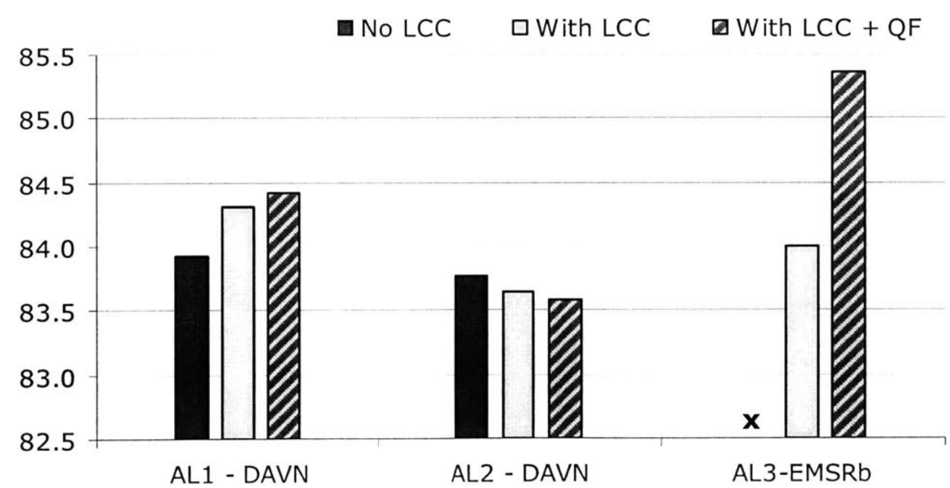

Figure 40: Impact of Q-Forecasting on AL1, AL2 and AL3 Network Loads

The loads of AL1 increase very slightly in the cases when it competes with the LCC. On contrary, AL2's loads decrease slightly. Like AL1, AL3's loads increase, but very significantly this time, when AL1 changes from PU to QF (AL3 captures many Q passengers denied by AL1 due to QF).

In order to understand the changes in AL1's loads, we focus now on the passengers the airline carries on the 10 legs where it competes with the LCC. We can see on *Figure 41* the local passengers of AL1 (opposed to connecting passengers).

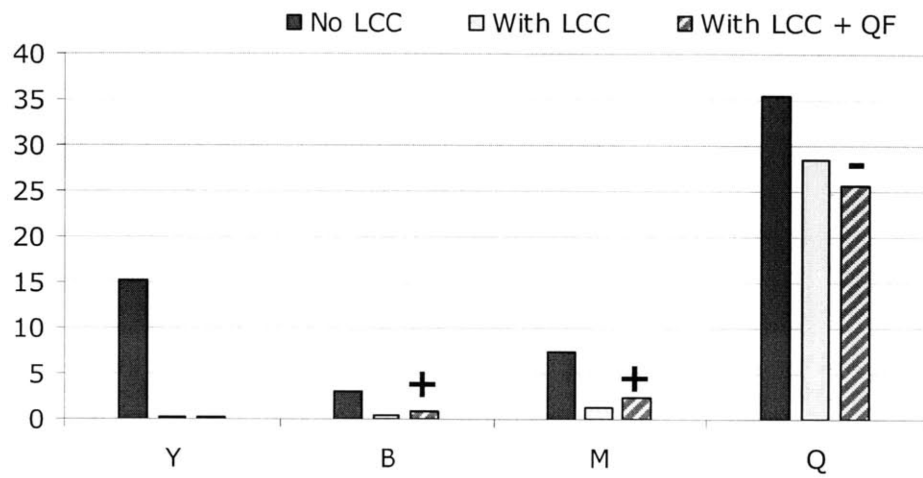

Figure 41: Impact of Q-Forecasting on AL1 Local Loads in the NEM

The QF method worked as predicted. After the drop in all classes caused by the entrance of the LCC, QF helps protecting more higher-class seats (increase in Y, B and M), while reducing the number of Q bookings. On the contrary, the LCC's Q loads increase as we noticed on *Figure 40*. Yet this did not explain the increase in AL1's network loads; so let's consider the breakdown of the loads on these 10 legs with respect to local and connecting passengers (*Figure 42*).

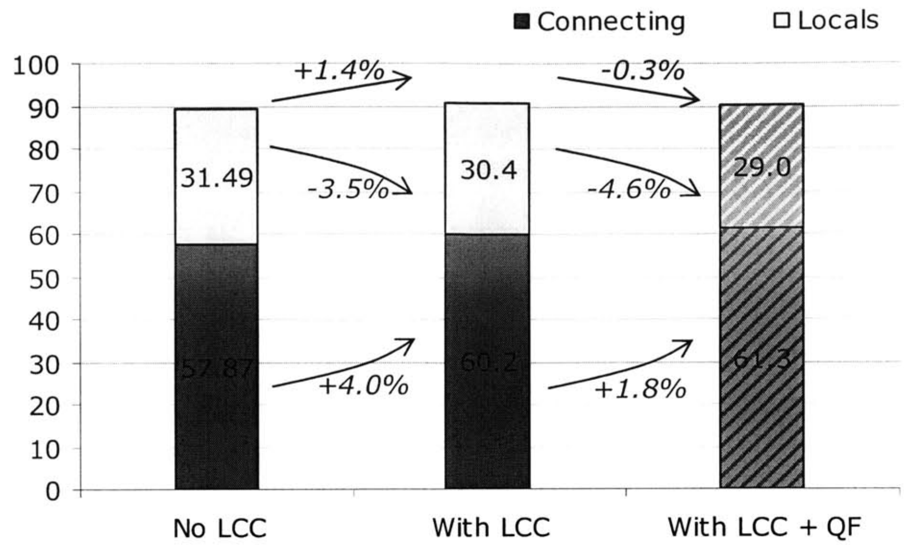

Figure 42: Impact of Q-Forecasting on AL1 Passenger Types in the NEM

If the number of local passengers has decreased (by 4.6%) as we saw previously when AL1 switches from PU to QF, the number of connecting passengers increases on these legs by 1.8%, especially when AL1 uses QF. Therefore, on total, the loads increase with respect to the base case. The breakdown of the connecting passengers would also show that AL1 carries more Y, B and M connecting passengers with respect to Q connecting passengers. So AL1 also benefits from QF at the network level. Other network effects would probably explain why the network loads increase slightly between PU and QF while total loads in the NEM decrease slightly (-0.3%) at the same time.

Figure 43 and *Figure 44* show the Q and Y booking curves (as a function of Time-Frames) when AL1 is using PU and when AL1 is using QF, differentiating local and connecting passengers.

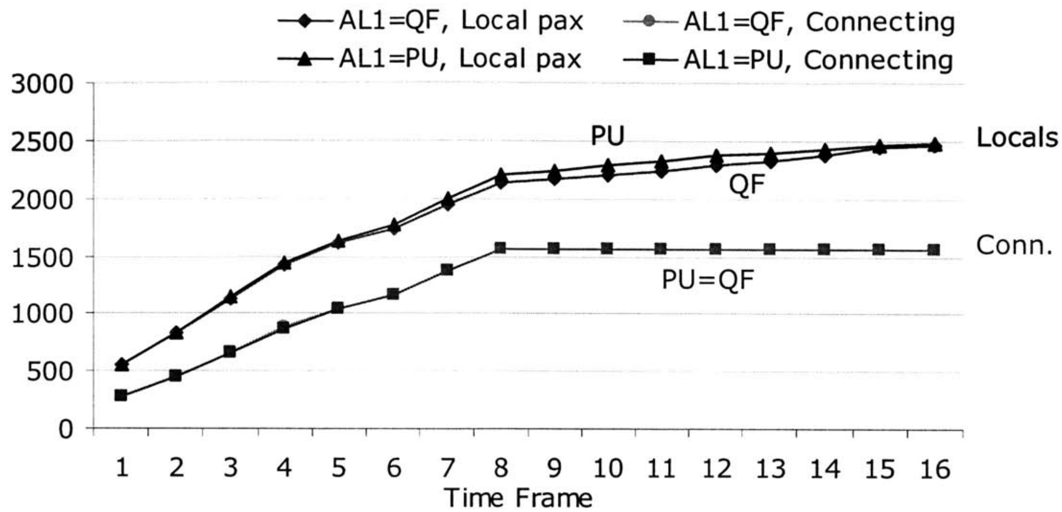

Figure 43: Daily Number of Bookings for AL1 in Q Class

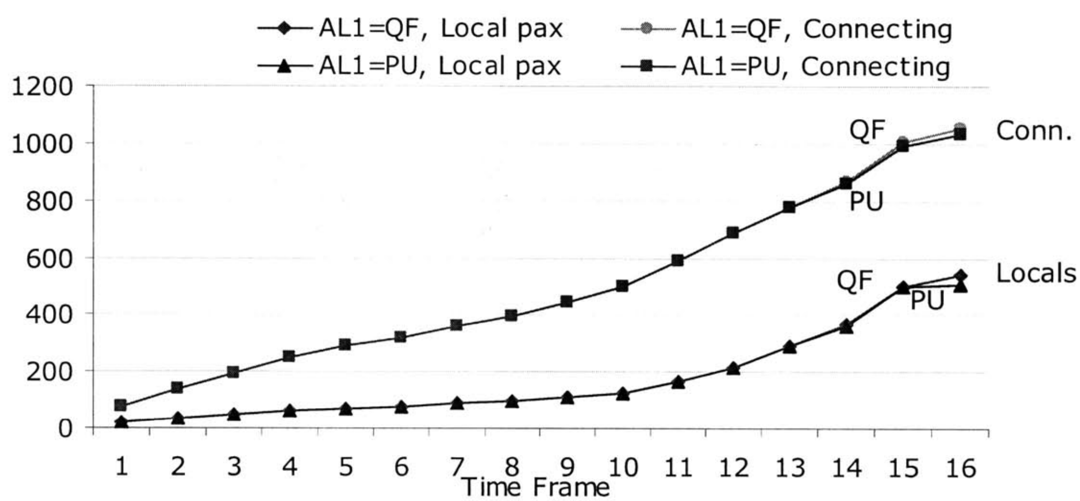

Figure 44: Daily Number of Bookings for AL1 in Y Class

In Q class, we see that QF has prevented some local bookings after TF_7, very early in the process. Regarding the Y class, we notice that QF allowed some extra bookings in the very last time-frames for both local and connecting passengers.

So Q-Forecasting increased AL1's revenue by 0.62%, and worked by protecting more seats for high-classes (Y, B and M), denying more Q-booking requests. It was applied to local passengers, for whom AL1 competes with the LCC, but its impact also affected connecting passengers, given that more seats are available for them in the aircraft due to the presence of the new entrant, which captures more local passengers.

We consider now the Revenue Management method used by the LCC. In the previous graphs, we saw the case when it was using EMSRb. As said in the previous sections, EMSRb in unrestricted environments leads inevitably to a spiral-down in revenues. So we want here to check if Q-Forecasting also improves AL1 revenues when the LCC is using for instance the Fixed Threshold RM method (with thresholds set to .75/.60/.45). Again in this case, AL1 and AL2 are using DAVN, the demand factor is set to 1.0, and AL1's input Frat5 is "B". We present on *Figure 45* the Revenue change when AL1 switches from PU to QF, for AL1 and AL3, given AL3's RM method.

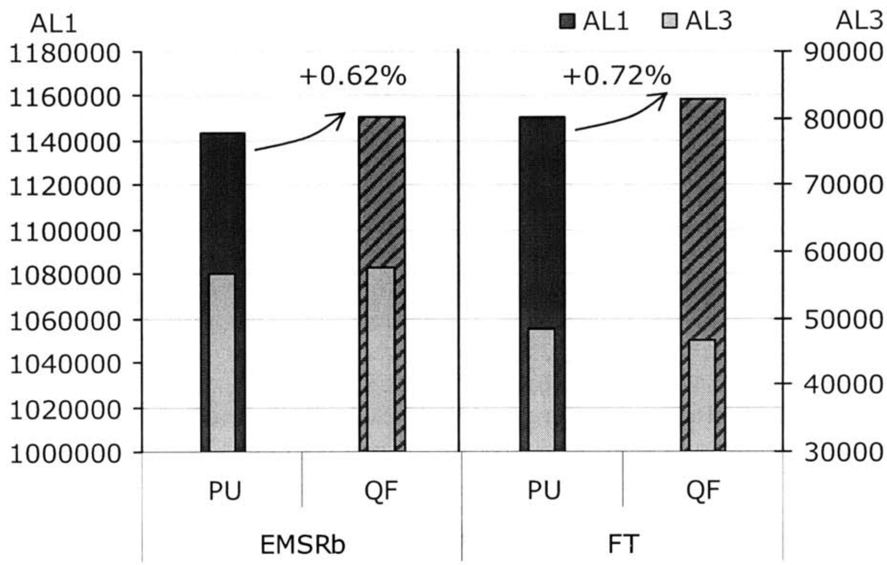

Figure 45: Impact of the LCC's RM on AL1's Revenue Gains Due to QF
The first line of the X axis refers to the forecasting method used by AL1, the second line refers to AL3's RM

The first 2 sets of columns act as a reminder of the results we have already shown. The other 2 sets represent the equivalent results when the LCC is using FT instead of EMSRb. As shown, the benefits of QF for AL1 are slightly higher in this second case (+0.72% increase in AL1's revenues), while AL3 is not taking advantage from this change. The corresponding ALF are not very different from the values shown in the case when AL3 was using EMSRb, except that AL1's loads are slightly higher when AL3 is using FT.

Looking at AL1 and AL3's fare class mixes for the local passengers in the 10 NEM (*Figure 46*), we notice the same mechanisms as for the EMSRb case. AL1 is protecting more high-class seats and denies more Q booking requests, while AL3's loads are declining very slightly, especially in the highest classes.

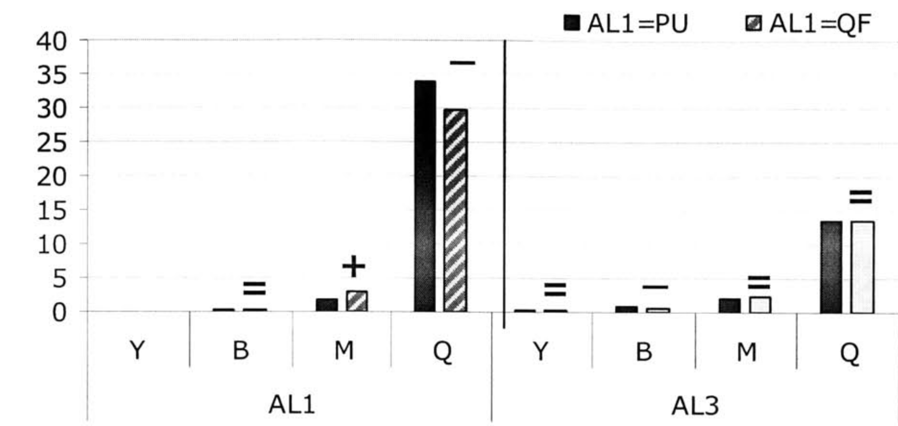

Figure 46: Impact of QF on AL1 and AL3 Local Loads when AL3 uses FT

In fact, similar impacts are observed when AL1 uses QF, whether AL3 uses FT, AT, EMSRb or EMSRQ2 (as introduced in sub-section 3.1.5). Only the amplitude of the improvement varies slightly from one case to the other. When AL3 uses EMSRQ2 for instance, the increase in AL1 revenue is reduced to 0.54%. EMSRQ2 is indeed more adapted to unrestricted environments than EMSRb, thus the LCC is more effectively competing against the incumbent.

Another case we wanted to consider is when AL1 switches to other RM for which QF is applicable, for instance ProBP and HBP. *Figure 47* shows the benefits of QF for AL1 when the airline applies QF to DAVN, HBP and ProBP. In this case, AL2 uses EMSRb and AL3 uses EMSRQ2. The Demand Factor and the Frat5 sets remain the same.

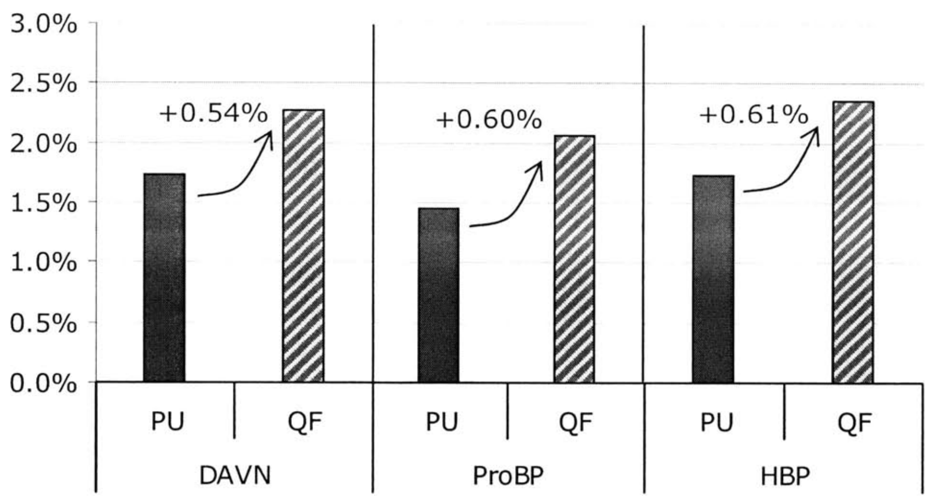

Figure 47: Impact of AL1's RM on the Revenue Gains Due to QF With Respect to EMSRb/EMSRb/EMSRQ2
The first line of the X axis refers to the forecasting method used by AL1, the second line refers to its RM method.

The benefits of QF are not very sensitive to AL1's RM method, even if they are a little higher when AL1 uses HBP and ProBP than they were with DAVN. Overall, the benefits remain in the range of an increase by 0.55 to 0.60% in revenue.

The last relevant sensitivity to analyze regarding the QF-input method is the impact of the Frat5 values we are using as inputs. So we are now considering a less aggressive set of Frat5 - "D" - which should protect fewer seats for the highest classes. The most aggressive set of Frat5 is not necessarily the set which leads to the highest revenue. Remember the goal is to protect enough seats for high-fare, late-booking passengers, but also to avoid flying empty seats, which would in the end lead to a decrease in revenue. So we do not want necessarily to be too aggressive regarding the Frat5 values.

In this case, we consider AL1 using DAVN, AL2 using EMSRb and AL3 using EMSRQ2, for a Demand Factor of 1.0. *Figure 48* is just a reminder of the Frat5 sets shapes as a function of time-frame. The less aggressive set is indeed below the other.

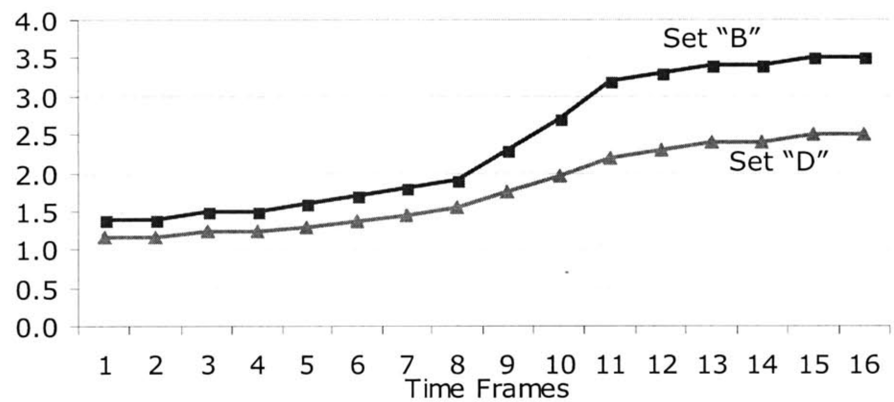

Figure 48: Frat5 sets "B" and "D" as a function of Time-Frame

Given these sets, we consider *Figure 49*, on which are presented the change in AL1's revenue when the input Frat5 is modified.

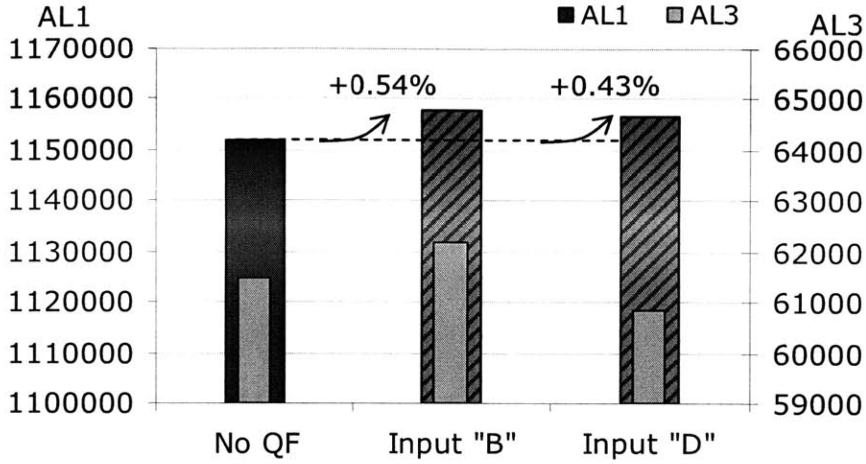

Figure 49: Impact of a change in Frat5 on AL1 and AL3 Revenues

The graph shows that the improvement provided by QF is lower for the Frat5 set "D", which happens to be too flat and therefore protects fewer seats for the higher classes. Nevertheless, it still represents an increase in revenue by 0.43% with respect to PU forecasting.

4.2.2. Estimator

Following the previous sub-section, our main interest is in the impact of QF in Network D. Nevertheless, some of the simulations were done in the single market case, especially the sensitivity analyses which require a lot of runs. So we will present some results in both the single Market case and in the case of Network D.

We consider first the Single Market case, for which *Table 12* summarizes the comparison in terms of performance between several RM methods. We deal again with an unrestricted environment (no advance purchase requirement either), and 8 fare classes as described in sub-section 3.2.2. In a given row of the table, which corresponds to a given RM method (with asterisk), we compare, in terms of revenue gains with respect to FCFS, which RM used by the competitor (columns) performs best. In this case, the sell-up estimator is used by EMSRQ2 as specified in the table.

	FT	AT	EMSRb	EMSRQ2-Est
FT*	+35%	+35%	+3%	**+43%**
AT*	+36%	+37%	+3%	**+44%**
EMSRb*	+48%	**+52%**	+0%	+10%
EMSRQ2-Input*	**+32%**	+28%	+1%	+28%

Table 12: Revenue Gains of Each RM Method with Respect to FCFS

We notice that depending on the RM*, there is no unique RM which performs best against all methods. Nevertheless, QF with sell-up estimator performs best against FT and AT, the most common RM methods used by LCC today, increasing revenues up to 44% with respect to FCFS, and it does well against EMSRQ2. The method is therefore competitive with respect to other ones presented here. It only does not perform well against EMSRb because in this environment, its competitor's revenues spiral-down and the estimator is unable to estimate sell-up rate.

We now return to Network D to analyze in more detail the performance of the sell-up estimator with respect to PU forecasting and with respect to QF with input Frat5. AL1 and AL2 are now both using DAVN, and AL3 is using FT (Thresholds at .75/.60/.45). The demand Factor is set to 1.0. *Figure 50* presents the revenue gains with respect to PU forecasting provided by QF-Input (QF with Input Frat5 "B") and QF-Est (QF with Sell-up Estimator) when the LCC uses FT or EMSRQ2.

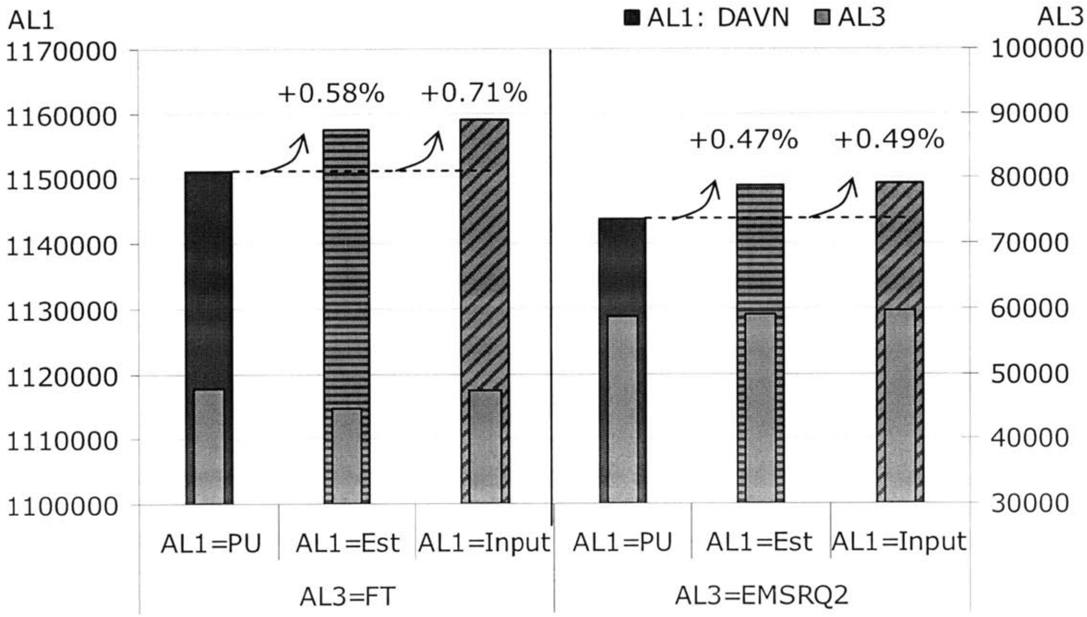

Figure 50: Performance of the Sell-up Estimator in Terms of Revenues

In the case when AL3 uses FT, the sell-up estimator leads to an increase in revenue by 0.58%, which is slightly below the gains provided by QF-Input, but nevertheless significant. The benefits amplitudes are lower in the case when the LCC uses EMSRQ2, but the same ranking is observed between QF-Est and QF-Input. AL3 indeed increases significantly its revenues when it uses EMSRQ2 relative to FT, so the improvements in AL1 revenues are lower.

In order to highlight the involved mechanisms and compare with those of QF-Input, we observe the corresponding loads on the 10 legs (case AL3=FT) and especially the breakdown between local and connecting passengers, as presented on *Figure 51*.

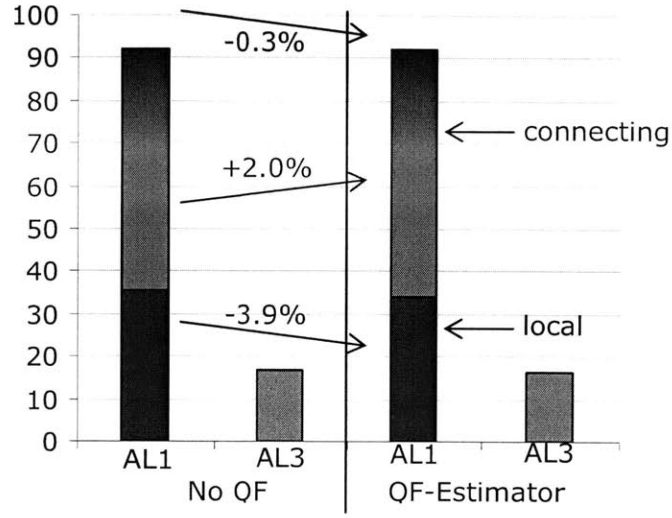

Figure 51: Impact of the Estimator on AL1 Local and Connecting Pax in the NEM

We notice that the exact same things happen, which are the decrease in local passengers and the increase in connecting passengers. The decrease in local passengers corresponds to a significant decrease in local Q passengers which outweighs the increase in local high-class passengers. On the other hand, the increase in connecting passengers is due to an increase in connecting high-fare-class passengers on these legs. These phenomena are shown on the following *Figure 52* and *Figure 53*.

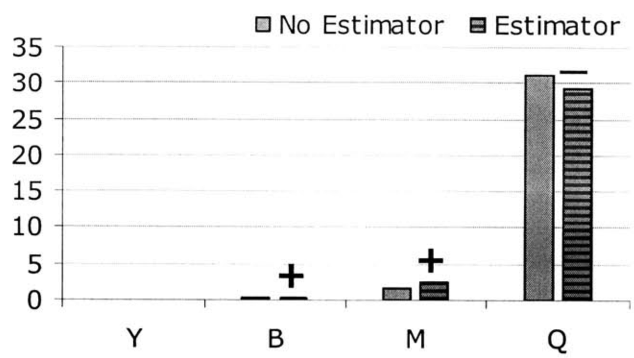

Figure 52: AL1 Local Loads in the NEM

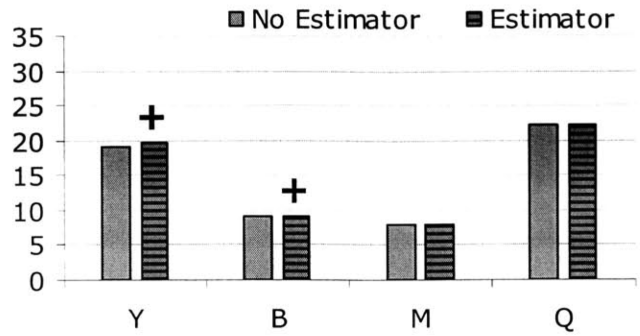

Figure 53: AL1 Connecting Loads in the NEM

In the next graph - *Figure 54* - we will compare the Estimated Frat5 with the typical Frat5 set that we were using in the QF-Input cases ("B").

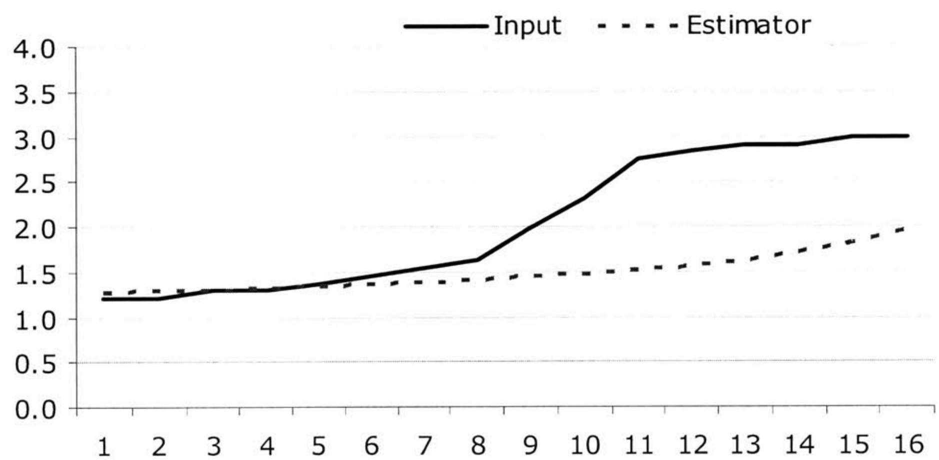

Figure 54: Estimated Frat5 Compared with Medium Input

As we can observe, the estimated Frat5 values are lower than the input values. Notice that the reported estimated values are averaged over time-frames since that they are computed at each of them. We point out also that the shape is not the same. We do not recognize the "S" shape which characterizes the input Frat5, but we observe a very flat curve until TF_{13} when it starts increasing slightly more significantly. This corresponds to the booking curve shape, which shows some extra bookings in Y, B and M classes in the last time frames.

We now come back to the Single Market case in order to analyze the Demand Factor sensitivity. To do this, we consider AL1 using EMSRQ2 with QF and the sell-up estimator, and AL2 using various RM methods. The fare products are unrestricted for both airlines (8 fare classes), and there is no advance purchase requirement. *Figure 55* presents the curves of AL1's revenue as a function of DF, against the different RM methods its competitor uses (FT, AT, EMSRQ2 - with input Frat5 "B" - and EMSRb).

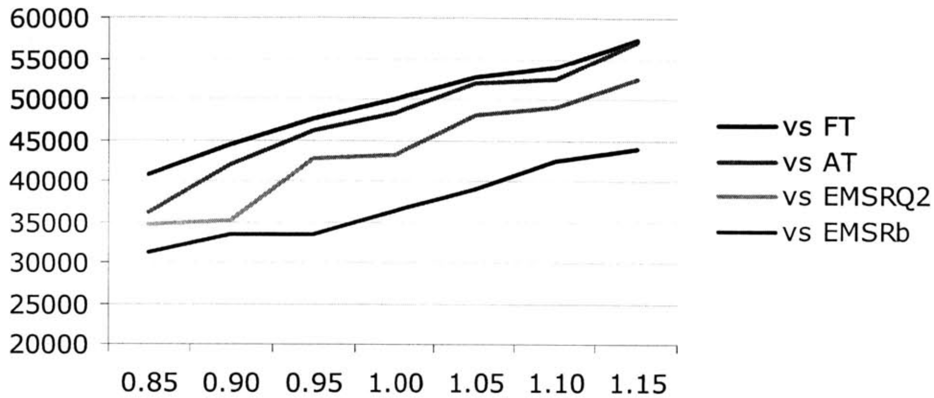

Figure 55: Sensitivity to Demand Factor in terms of Revenues

As expected, the revenues increase when demand increases. We see that the estimator performs best against FT and AT, whatever the Demand Factor is. Against EMSRQ2, the performance is below the others, given that this RM is more sophisticated than the Threshold methods, so both methods compete for the high fare passengers with equivalent tools, and in the end, they have to share. Against EMSRb, the problem is different. In unrestricted environments, EMSRb behaves almost like FCFS. Therefore, all passengers book the lowest fares which are available for a long period of time in AL1's inventory, and the estimator does not observe enough occurrences of sell-up. Thus, given the mechanisms of the estimator, it cannot predict high sell-up rates and in the end protects few seats for the high-fare classes. We also notice that the ranking between methods is not changed when the demand increases.

Figure 56 is the corresponding graph in terms of AL1's load factor. Again, as predicted, the ALF increases when the demand increases. At lower demand factor (0.85 to 0.95), EMSRQ2 with sell-up estimator captures more passengers against FT than against other RM methods, but starting from DF=0.95, EMSRQ2's loads are almost equal, whether competing against FT, AT or EMSRb. Against EMSRQ2 with input Frat5, again, both RM methods behave the same way, they focus on higher fare passengers, and compete for them in priority, denying more Q booking requests since sell-up occurs often.

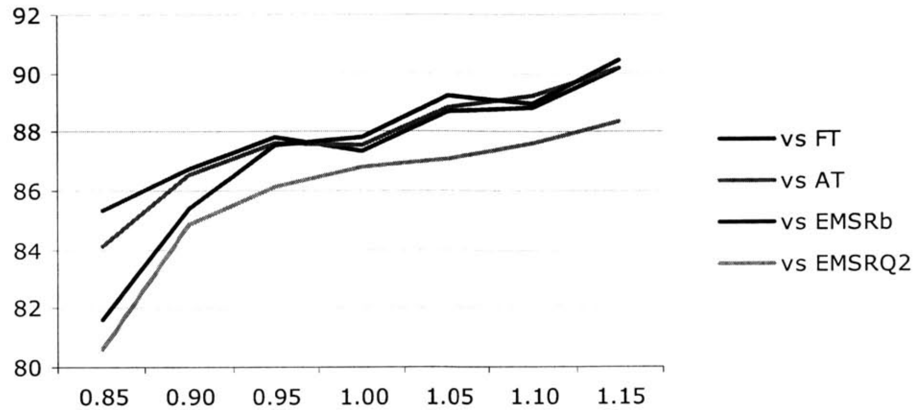

Figure 56: Sensitivity of Demand Factor in terms of Loads

In order to understand the mechanisms behind the increase in revenues, we observe in more detail AL1's fare class mix evolution when the demand factor increases, as presented on *Figure 57*.

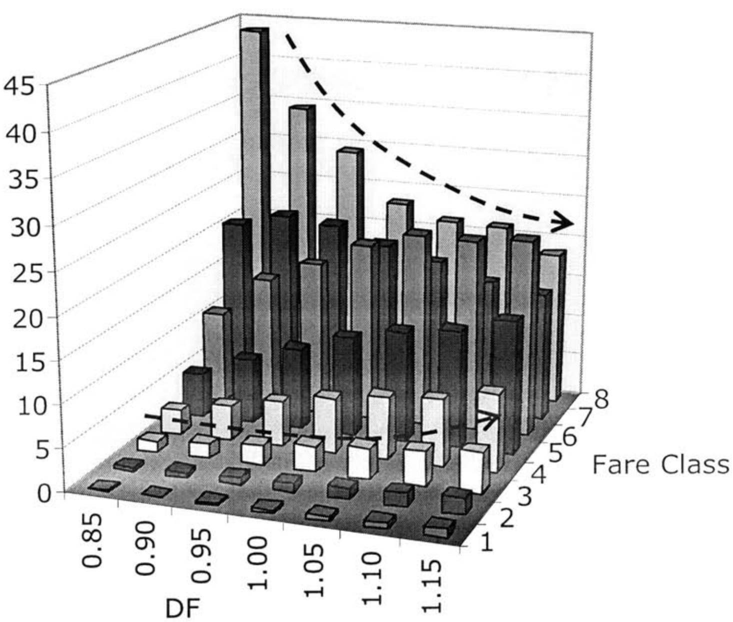

Figure 57: AL1 Fare Mix Trend as a Function of DF

This graph presents a good illustration of how the Q-Forecasting method works: as demand increases, the number of high-fare passengers also increases, so sell-up occurrences are more frequent. The consequence is that estimated Frat5 values rise too, and AL1 is protecting more seats for high fare-class passengers. This translates on the graph into a significant decrease in loads of classes 7 and 8, while classes 1 to 6 see their loads increase quite substantially, as indicated by the arrows. The increase is even significant in classes 1 and 2.

The mechanisms can be seen as well if we look at the change in the estimated Frat5 when the demand increases, as presented on *Figure 58*. On this graph, the higher the demand is, the darker the line.

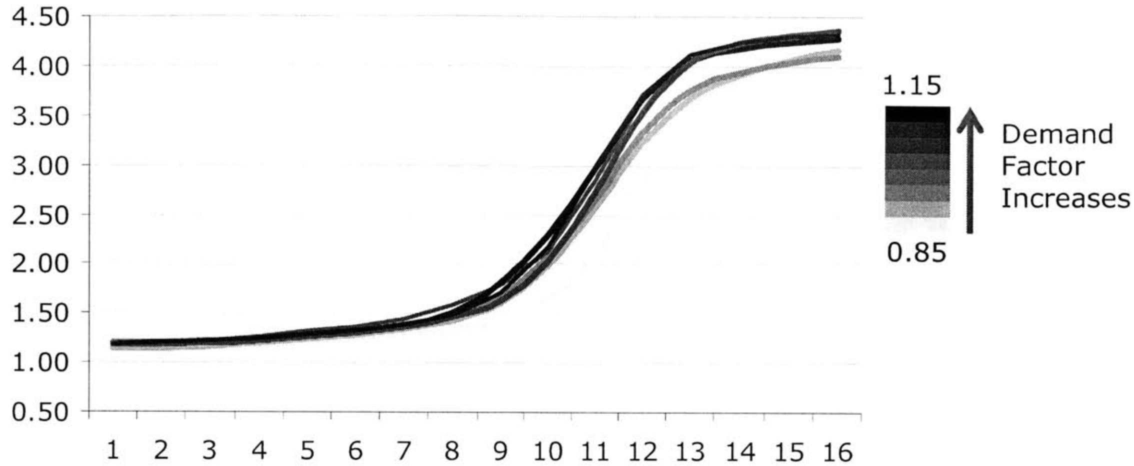

Figure 58: Changes in Frat5 (as a Function of TF) When DF Increases

In concordance with the previous observations, when the demand increases, the number of high fare class passengers rises, the frequency of observed sell-up rises, and the estimator predicts higher Frat5 values to protect more seats for late-booking, high-fare passengers.

4.2.3. Conclusion

The Q-Forecasting method improves the performance of traditional revenue management methods (DAVN, ProBP, HBP) in unrestricted environments by increasing their total network revenues on average by 0.60% in the cases considered in this section. It corresponds to an increase in revenue in the 10 NEM by 8.8%. The magnitude of the improvement depends on the demand level and on the LCC's Revenue Management method as well. The Q-Forecasting method actually protects more seats for high-fare, late-booking passengers in the 10 legs where the LCC entered, as well as for the connecting passengers flying on these legs, implying a non-negligible network effect. The sell-up estimator, even if it cannot exactly reach the performance of specific input Frat5 sets, performs relatively well in comparison, which is crucial given that only an estimator can work in the real world (where one does not want to have to enter arbitrary input Frat5 sets which should be different for each market and departure).

4.3. Performance of Fare-Adjustment Methods (FA)

In this section, we will analyze the performance and mechanisms of the Fare Adjustment methods, without using the Q-Forecasting method (see section 4.4 for combinations of both methods). Given the complexity in the number of adjustable parameters, and the similarities between methods, we will consider in this thesis only the KI Fare Adjustment method using the time-frame sell-up constant. The other methods (KI with BTC sell-up constant, MR with TF sell-up constant and MR with BTC sell-up constant) give approximately the same results and the mechanisms they involve are almost the same. There is no sell-up estimation properly implemented at this point in PODS in order to dynamically predict the Frat5 values for the Fare Adjustment methods, so we will focus on the use of input Frat5. The use of the FA method will be generally emphasized in the next figures by vertical stripes.

4.3.1. Input Frat5 values

As shown in section 3.5.4, the sell-up considered in the case of Fare Adjustment is "unconditional" whereas it is conditioned on the closure of the Q-class in the case of Q-Forecasting. Therefore, the Frat5 values used for Fare Adjustment methods have to be much lower than those used for Q-Forecasting. That is the reason why we define here some very simple (linear) Frat5 sets, named after low-case letters in order to distinguish them from the QF Frat5 sets (from h to z). *Figure 59* presents the shapes of these Frat5, and *Table 13* the value of the each Frat5 set at TF_1 and TF_{16}.

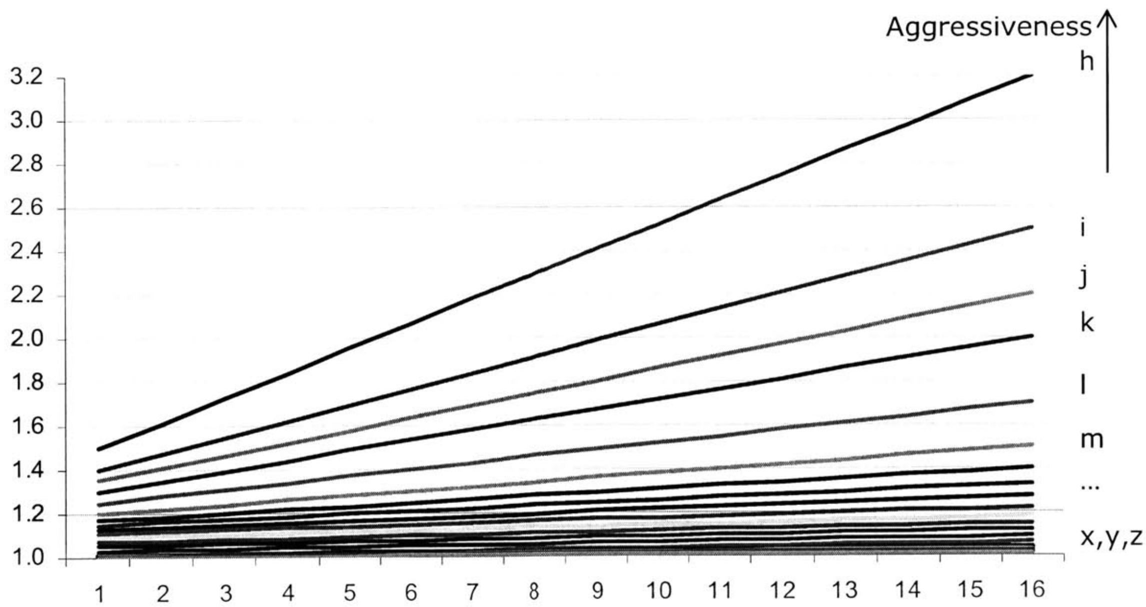

Figure 59: Fare Adjustment Frat5 Curves as a Function of Time-Frame

Set	h	i	j	k	l	m	n	o	p	q
TF1	1.500	1.400	1.350	1.300	1.250	1.200	1.170	1.150	1.130	1.110
TF16	3.200	2.500	2.200	2.000	1.700	1.500	1.400	1.330	1.270	1.220

Set	r	s	t	u	v	w	x	y	z
TF1	1.090	1.070	1.050	1.030	1.010	1.007	1.005	1.003	1.002
TF16	1.180	1.150	1.120	1.090	1.060	1.040	1.030	1.020	1.010

Table 13: Values of the FA-Frat5 Sets at TF_1 and TF_{16}

We will first analyze in detail the results obtained with one given set of Frat5 ("k"), which is the Frat5 sets which provided the best results, before looking at the sensitivity to a change in Frat5.

4.3.2. Impact of Fare Adjustment Methods on DAVN

We consider the case of Network D with the LCC entering the 10 non-stop markets, complying with the description given in sub-section 3.2 in terms of fare products (no restriction, no advance purchase requirements, and compressed fare ratios). AL1 uses DAVN without QF, AL2 uses DAVN, and AL3 uses AT with a target load factor of 80%. The FA method is applied to AL1, in the 10 NEM only, and uses the Input Frat5 "k" (which varies linearly from 1.3 to 2.0) described in the previous sub-section. The Demand Factor is set to 1. On *Figure 60* is presented the impact of the FA method on AL1 and AL3's revenues relative to the case when AL1 does not use FA ("base case").

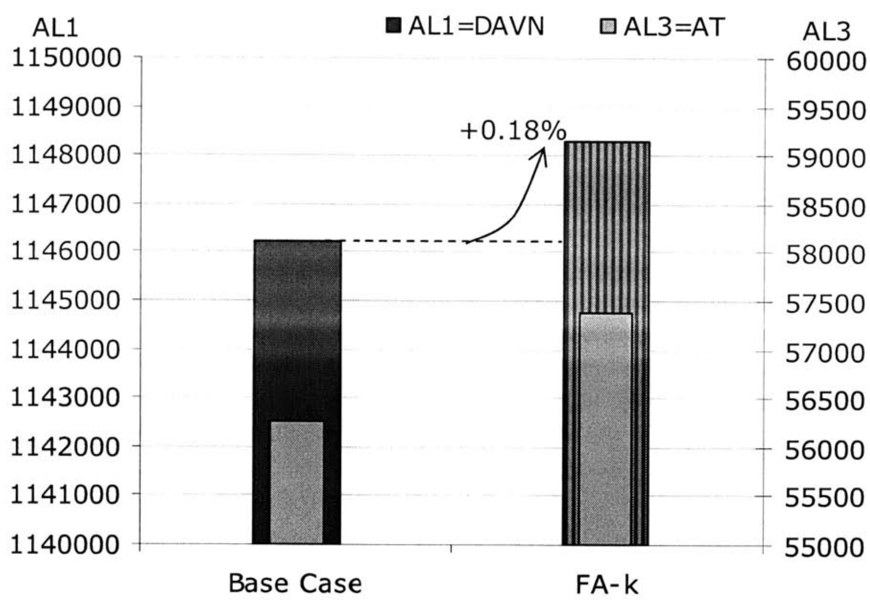

Figure 60: Impact of the FA Method on AL1 and AL3 Revenues

The utilization of the Fare Adjustment method by AL1 with DAVN allows the airline to increase its total network revenue by 0.18%. We will see later that the range of the revenue improvement, like in the case of Q-Forecasting, varies depending on the Frat5 set used as input. As one can notice, the new entrant benefits from the FA (its revenue increases by 2.01 %). The corresponding network average load factors for both airlines increase as well, by 0.20% for AL1 and by 1.57% for AL3. We will see in more detail how the method worked in the following graphs.

Like we did in the case of QF, we observe the effect of the FA method on the loads on the 10 affected legs. On *Figure 61* is shown the impact on AL1's connecting and local passengers, as well as AL3's local passengers.

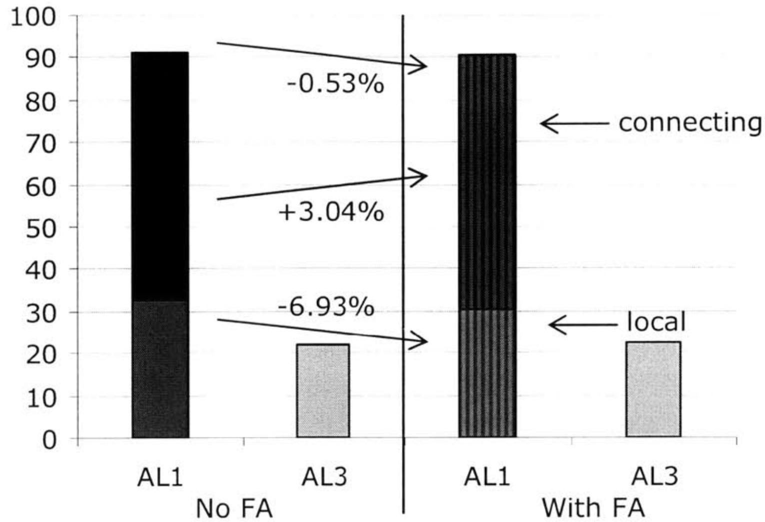

Figure 61: Impact of the FA Method on the Local/Connecting Passengers

Looking at AL1 total loads on these legs, we notice a slight decrease in passengers by 0.53%. But if we consider the breakdown relative to connecting and local passengers, we realize again that the method did not act the same way on both categories. The local loads decreased by 6.93% whereas the number of connecting passengers increased by 3.04%. Thus, we observe, like in the case of Q-Forecasting, some non-negligible network effects. The loads of AL3 increased at the same time by 1.55%. *Figure 62* and *Figure 63* present another breakdown of AL1's passengers traveling on these legs: the change in (respectively) local and connecting passengers fare mixes.

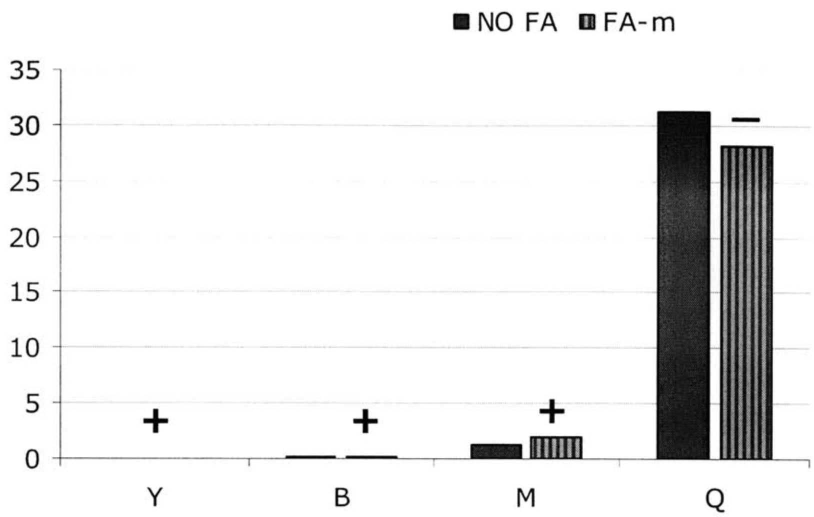

Figure 62: Impact of the FA Method on AL1 Local Loads in the NEM

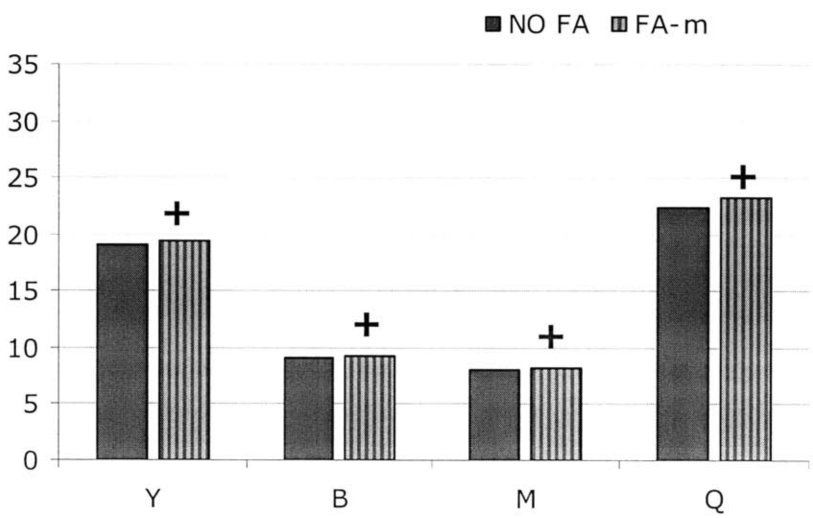

Figure 63: Impact of the FA Method on AL1 Connecting Loads in the NEM

The changes in local loads are similar to the ones observed for QF. The FA method leads to a decrease in Q bookings, whereas it increases the bookings in Y, B and M classes. On contrary, regarding connecting passengers, the FA method leads to an increase in all fare classes loads, even in Q.

To understand the mechanisms of the method, we observe the graph of the adjusted fares (= ODF-PEcost) for the 4 classes, as a function of Time-Frame, shown on *Figure 64*. The values reported here are averaged adjusted fares of the actual adjusted fares on the 10 legs. The Y fare is not adjusted, so it does not vary.

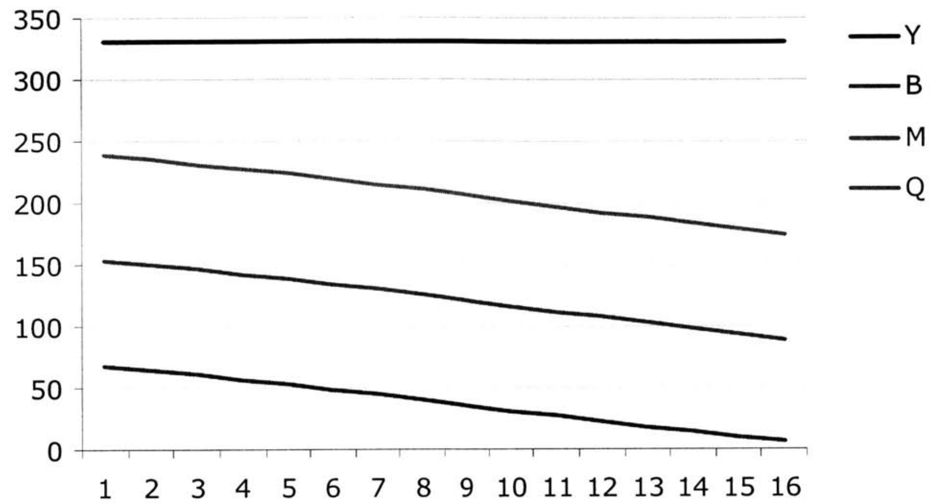

Figure 64: AL1 Adjusted Fares in the NEM

The graph shows that the adjustment changes during the booking process. This observation complies with the fact that the Frat5 values vary as well. The adjusted fares decrease as we approach the departure date, meaning that we send them into lower buckets, which are closed earlier. In other terms, the lower the adjusted cost is, the earlier the corresponding booking request will be denied. The adjustments applied here are very similar, in terms of amount of money, from one class to the other. The Y class, to which the adjustment is not applied, will benefit from the earlier closure of the lower classes. The Q-class is almost closed down by the method at TF_{16}.

4.3.3. Sensitivity to Input Frat5

Like in the case of Q-Forecasting, it is interesting to observe the impact of a change in the Frat5 values. So we run several times the same network with the different input Frat5 presented in section 4.3.1, from h (the most aggressive set) to z. On *Figure 65* is shown the evolution of AL1's revenue when the aggressiveness of the Frat5 decreases.

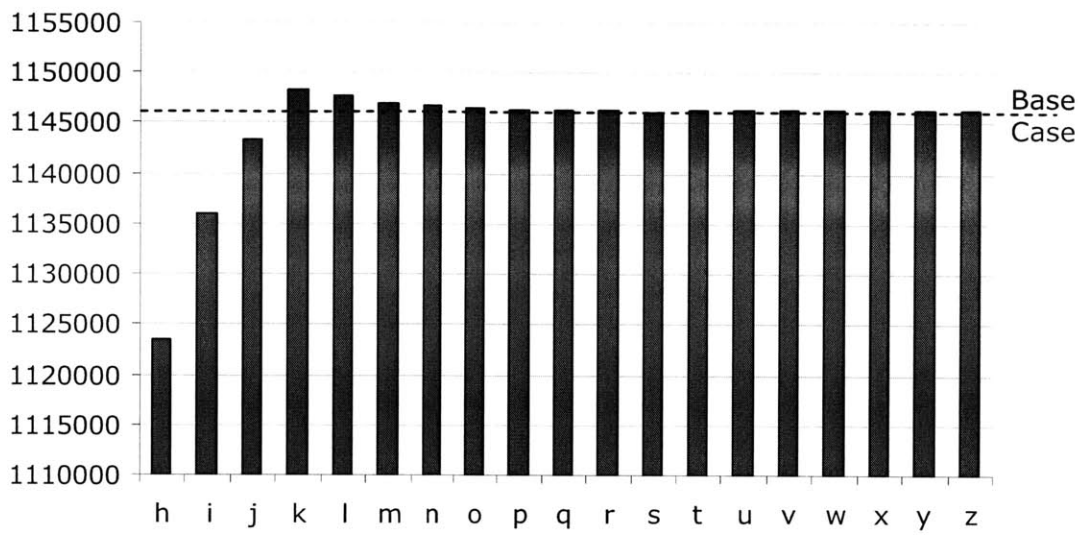

Figure 65: Sensitivity to FA Frat5

The sensitivity analysis shows the presence of a maximum in AL1's revenue - $1,148,295 for the Frat5 set "k" (whose values go from 1.3 to 2.0). In the base case (No FA), the revenue of AL1 was $1,146,207, which is approximately the level reached here for the Frat5 r to z. The maximum represents an increase by 0.18% with respect to the base case. We notice that the revenue drops very fast when we increase the aggressiveness from set k (left hand side of the graph). On contrary, for less aggressive Frat5, we reach a plateau at the level of the base case (No Fare Adjustment, no Q-Forecasting).

In order to understand the previous graph, we observe how the adjusted fares in the 10 NEM change when the Frat5 become more or less aggressive. On *Figure 66*, *Figure 67* and *Figure 68*, we see the averaged adjusted fares respectively for the B, M and Q classes. The darker the line, the *less* aggressive the Frat5. The curves corresponding to the Frat5 "k" (optimal found in this analysis) are thicker.

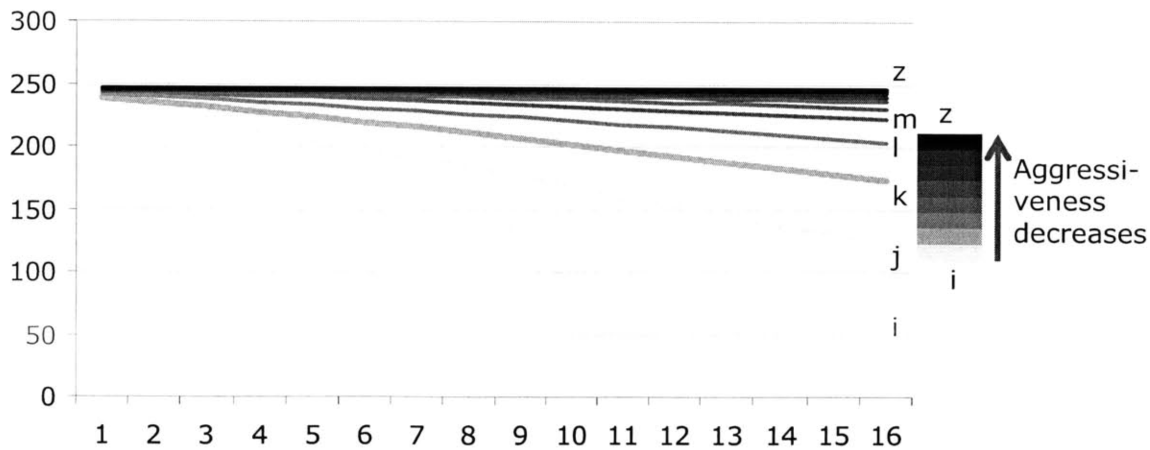

Figure 66: Adjusted B Fares in the 10 NEM

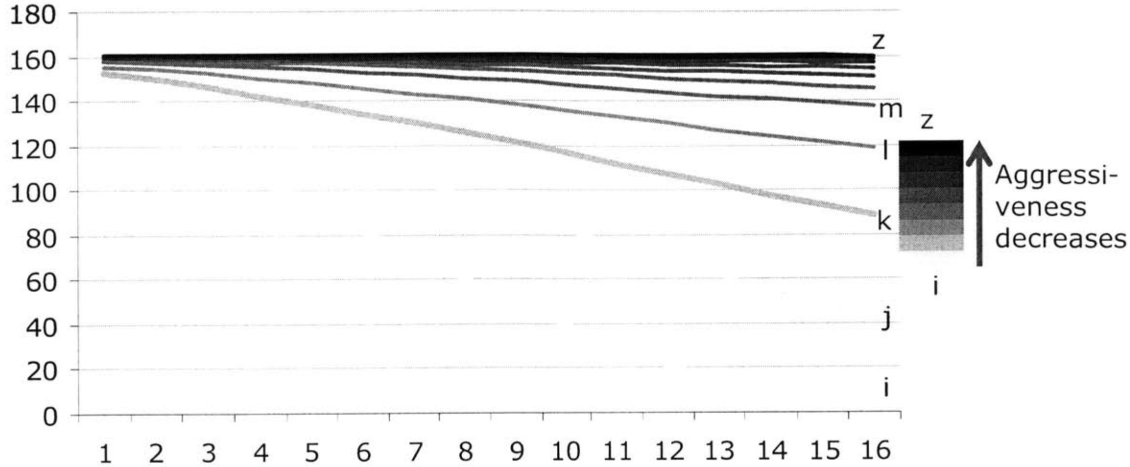

Figure 67: Adjusted M Fares in the 10 NEM

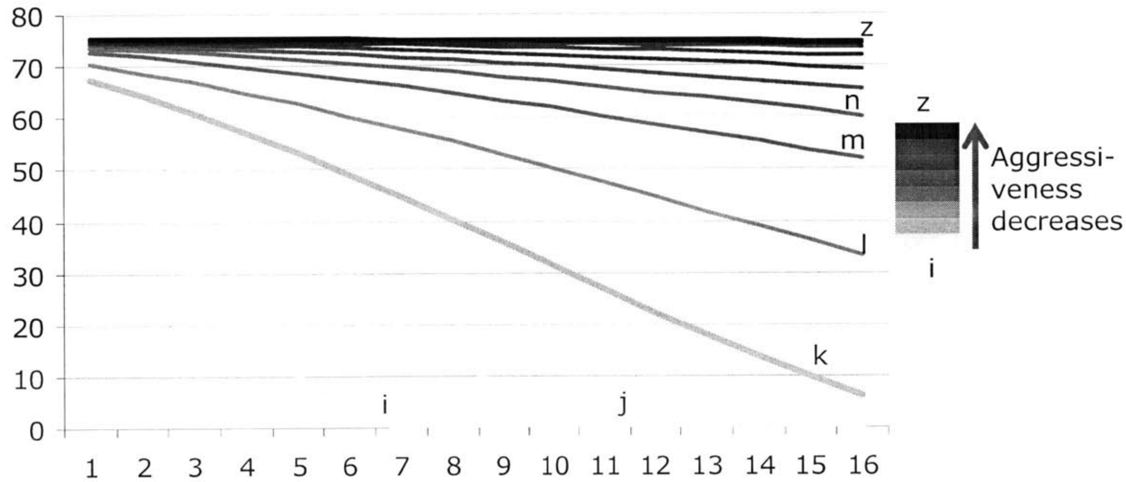

Figure 68: Adjusted Q Fares in the 10 NEM

For a given fare class, the adjusted fares decrease when the Frat5 become more aggressive (what we observe when we slice one of these graphs for a given Time-Frame). The aim being to protect more seats for the higher classes, lower adjusted fares make the class unavailable. For the less aggressive sets (r to z), the adjustments are almost not visible, and the effects are not significant with respect to the base case. Consequently, the corresponding revenues are at the same level as the base case. On contrary, the most aggressive Frat5 are very aggressive, even reducing the M and Q fares to $0. As a consequence, the lower classes are closed down early, probably too early, and too many seats are protected for the highest classes. Thus, revenues drop fast. The optimum Frat5 set "k" is a tradeoff between both extreme cases.

4.3.4. Conclusion

The Fare Adjustment methods help improving DAVN in unrestricted environments, by increasing its total network revenue by about 0.20% with respect to the case "No QF / No FA" under the conditions considered in this section. Again, it corresponds to an increase by 3% in the 10 NEM. The revenue is quite sensitive to the input Frat5 around the optimal set. The method actually decreases the local fares in the unrestricted environment before running the optimizer and sends them in lower buckets which are closed earlier to protect more seats for the high-fare customers. The magnitude of the decrease in adjusted fares depends on the input Frat5.

4.4. Combination of Q-Forecasting and Fare Adjustment Methods

4.4.1. Using the same Input Frat5

Before examining the differences between the Frat5 for Q-Forecasting and the Frat5 for Fare Adjustment methods, we used the same set for both methods. On *Figure 69* are shown the Frat5 sets which will be used in this section.

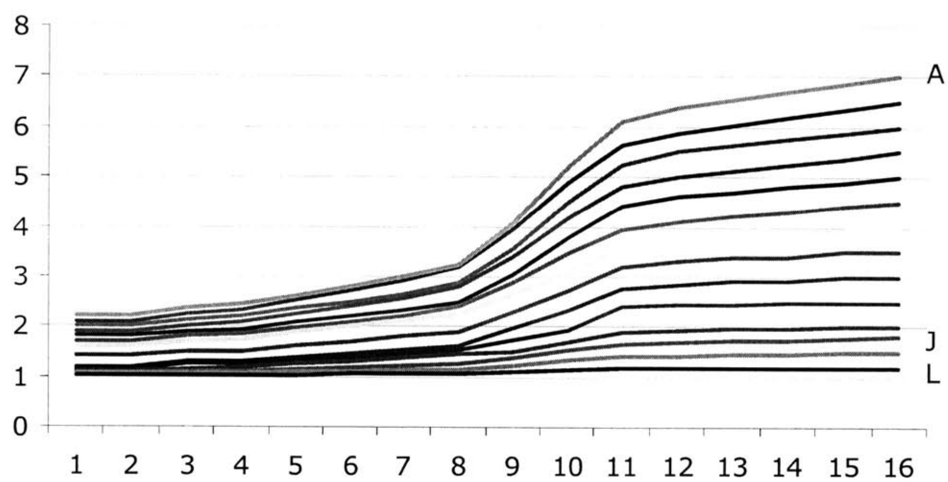

Figure 69: Frat5 Sets Used For Both QF and FA Methods

Figure 70 presents the results obtained in terms of AL1 and AL3 revenues for the case DAVN/DAVN/AT (with target load factor at 80%) at demand level 1.0 when the same Frat5 set – J – is used by both methods.

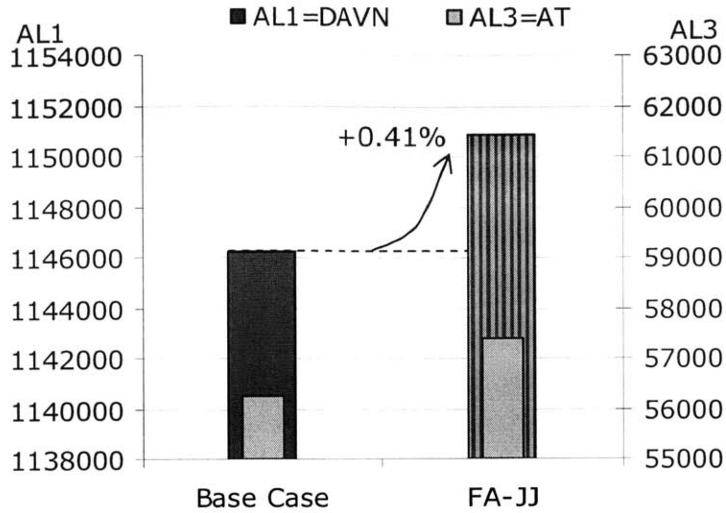

Figure 70: Revenue of AL1 and AL2 When AL1 Uses Frat5 Set "J" For QF and FA

The use of both Q-Forecasting and Fare Adjustment methods with the same input Frat5 leads to an increase in AL1's revenue by 0.41%, which is a good improvement. At the same time, AL3 improves its revenue by 2.68%. The next figure - *Figure 71* - presents the breakdown of AL1 and AL3 loads on the 10 legs in terms of local and connecting passengers.

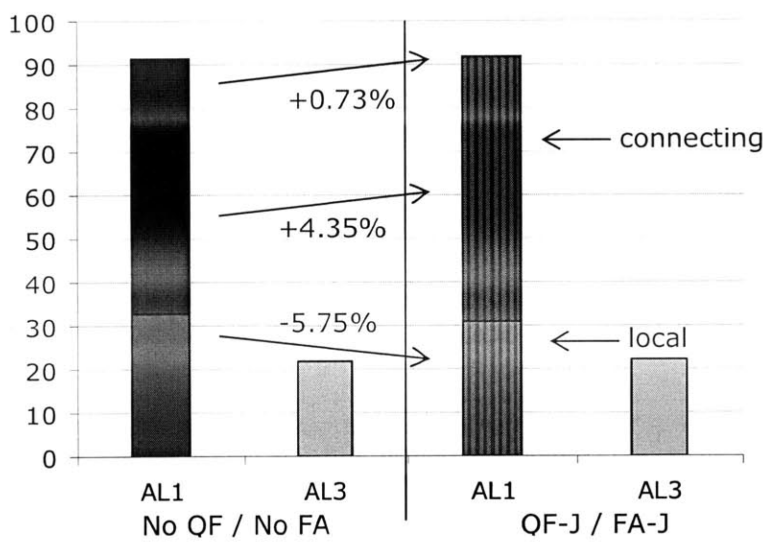

Figure 71: AL1 and AL3 Local and Connecting Passengers

AL1 total loads on the 10 legs increased by 0.73%. Nevertheless, the effects of QF+FA were not the same on both local and connecting passengers. The local loads decreased by 5.75% whereas the connecting increased by 4.35%. Given that the latter are more numerous, the total results into an increase in loads. Simultaneously, AL3 increased its loads by 1.55%.

Figure 72 and *Figure 73* present AL1's loads on the 10 legs in terms of fare class mixes, respectively for the local and for the connecting passengers.

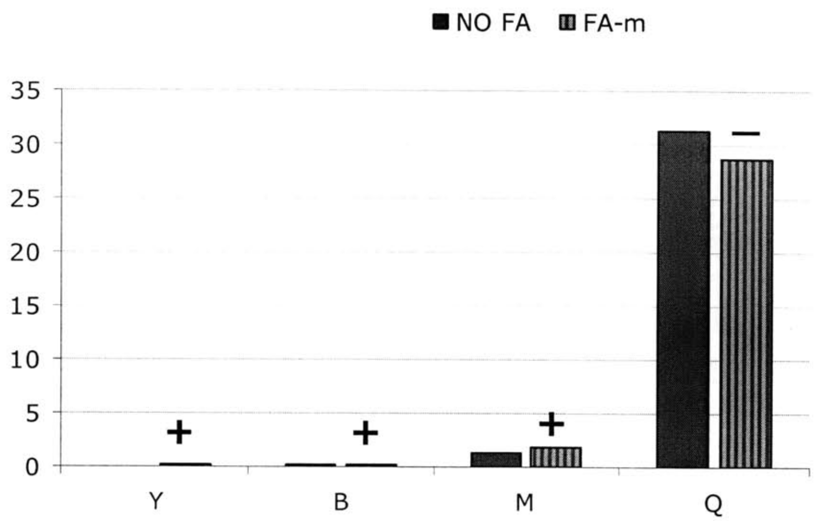

Figure 72: AL1 Local Passengers Fare Class Mix on the 10 Legs

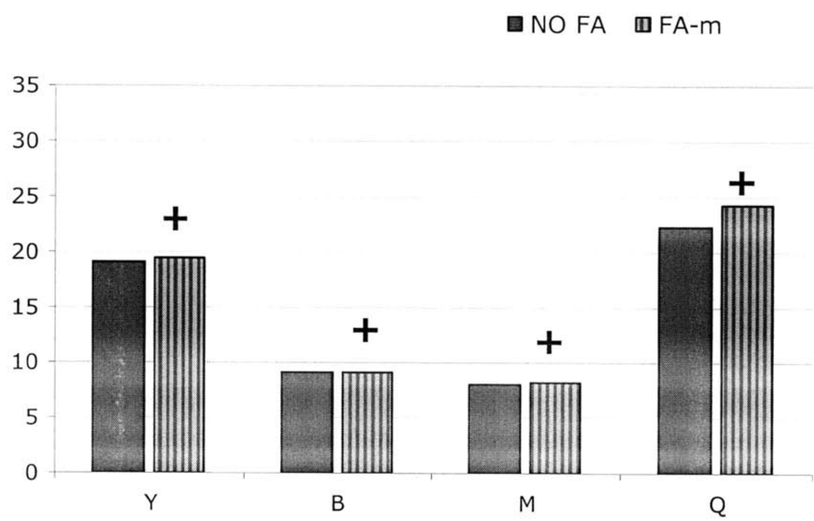

Figure 73: AL1 Connecting Passengers Fare Class Mix on the 10 Legs

Like in the case when only the Fare Adjustment method was applied, we observe a significant decrease in Q local passengers, whereas the Y, B and M loads increase. Regarding connecting passengers, the loads increase in all fare classes, leading to the total increase observed on *Figure 71*.

We show on *Figure 74* the behavior of Adjusted Fares across time-frames, for the 4 fare classes.

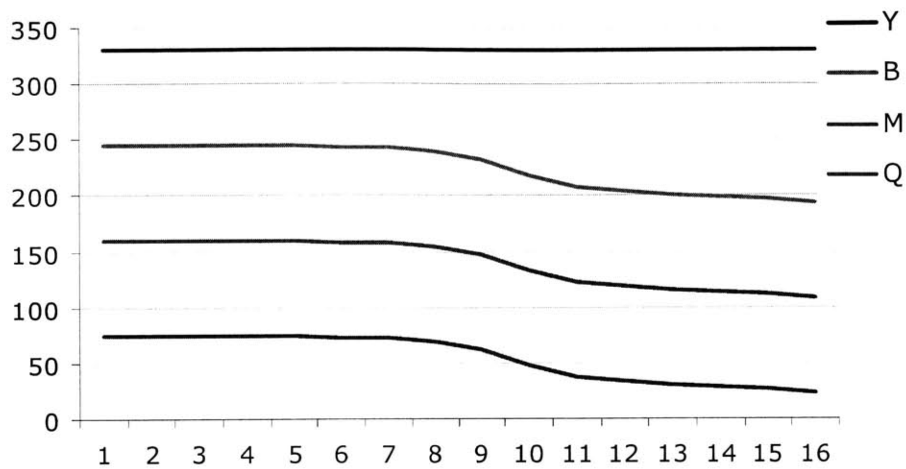

Figure 74: Adjusted Fares of AL1 in the 10 NEM

The Adjusted Fares decrease as time goes, and we recognize in the slope of the curve a relation with the opposite of the Frat5 slope. None fare class was closed down due directly to Fare Adjustment method in this case, even if the adjustments are more significant than the ones we observed in the previous section.

We analyze now the sensitivity to the Input Frat5, observing on AL1's revenue change relative to the Frat5 presented on *Figure 69*.

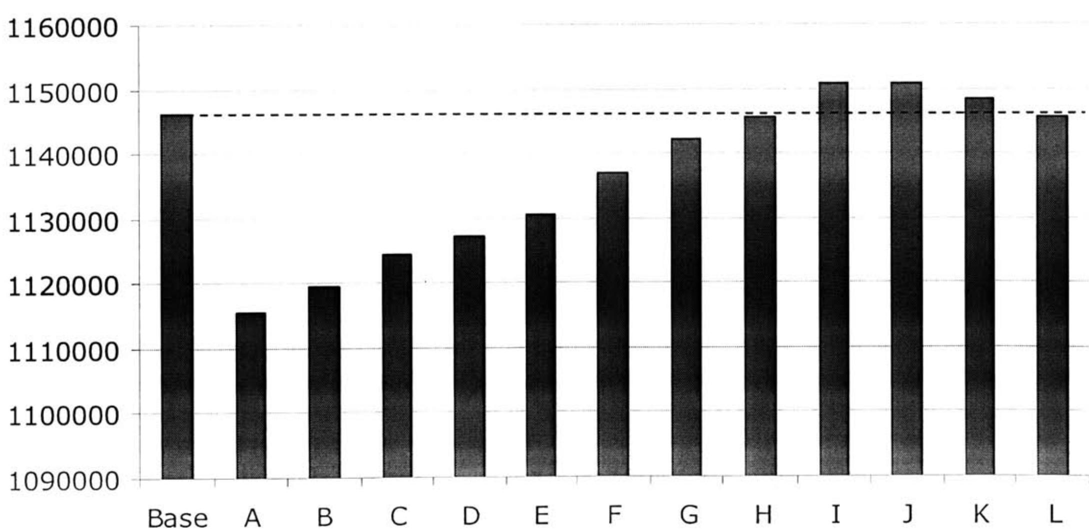

Figure 75: Evolution of AL1 Revenue When Frat5 Sets Vary From A to L

The first bar on the graph represents AL1's revenue in the base case - No QF /No FA. In the other cases, the letter corresponds to the Frat5 set used by both Q-Forecasting and Fare Adjustment methods. We observe a bell-shape with a maximum in revenue obtained for the Frat5 set "J". This case, as presented above, represents an increase by 0.41% in revenue with respect to the base case. Around the optimal set, the revenues quickly drop below the base case level. On *Figure 76*, *Figure 77* and *Figure 78* are shown the corresponding adjusted fares, respectively for the B, M and Q fares.

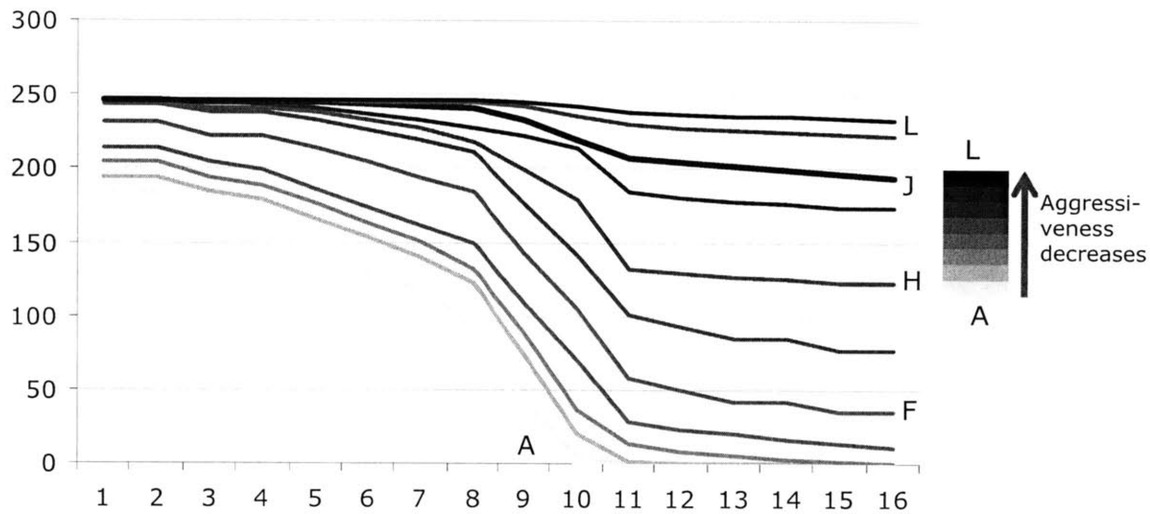

Figure 76: AL1 Adjusted B Fares in the 10 NEM

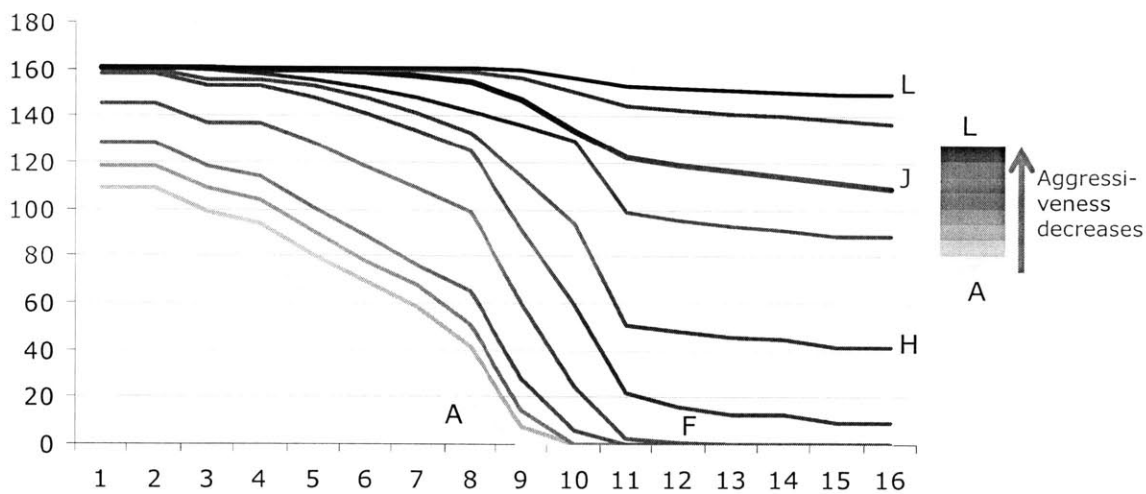

Figure 77: AL1 Adjusted M Fares in the 10 NEM

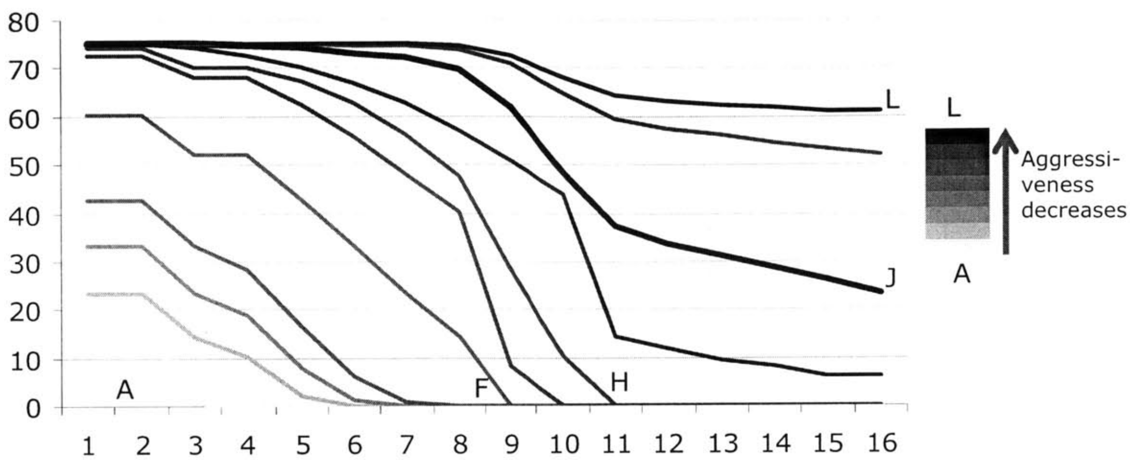

Figure 78: AL1 Adjusted Fares in the 10 NEM

The adjusted fares are globally lower than the graphs we observed in the case of the Fare Adjustment methods used alone, given that the Frat5 sets are more aggressive in this case. Consequently, in many cases, the Q, M and even B classes are closed down directly by the Fare Adjustment method, very early in the booking process. This leads to the pretty low revenues we observed on *Figure 75*. The Frat5 set "J" is emphasized on each graph with a thicker line.

On *Figure 79*, we consider the case when AL1 uses the FA method with different sets of Input Frat5 (the sets are shown on *Figure 80*). We observe then the benefits provided by the additional use of Q-Forecasting (again with *the same* input Frat5 as the FA method). We notice that in all cases, QF allows an improvement in AL1's revenue, even if it does not reach the revenue obtained by QF with Input "A" alone. The benefits of QF are more significant when the Frat5 set is aggressive (sets "A" and "C") than when the Frat5 are more adapted to FA alone (sets "E", "F" and "G").

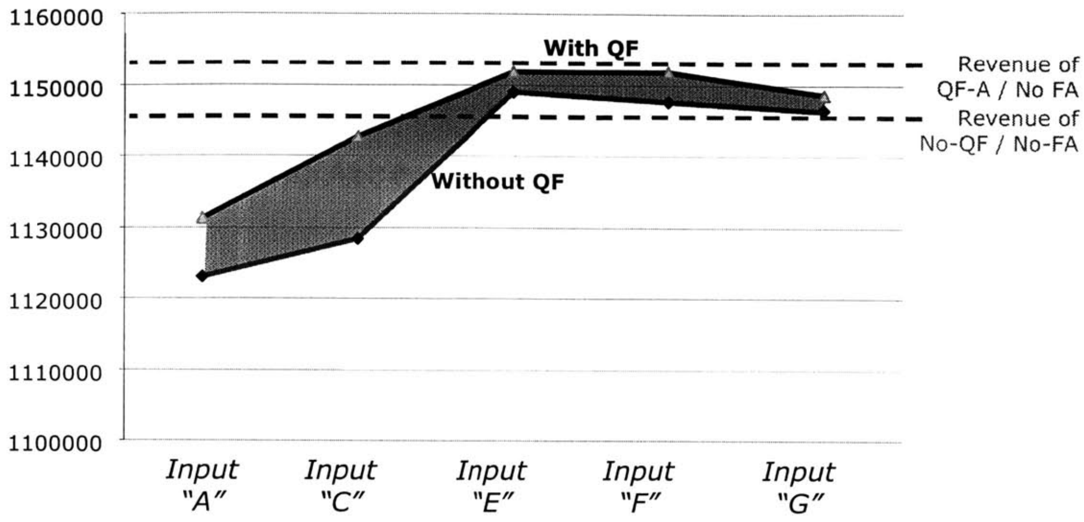

Figure 79: AL1's Revenue Sensitivity to Frat5 When Using FA and Additional Benefits Due to QF

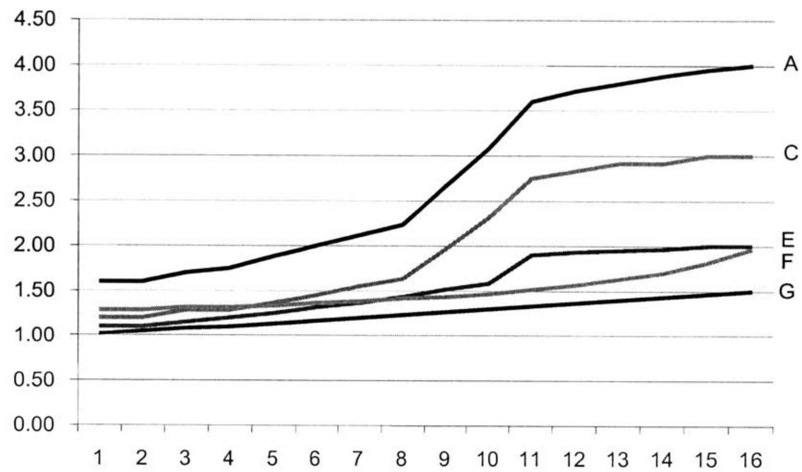

Figure 80: Frat5 sets used as input in the Figure 79

The problem pointed out by these graphs is that even if obviously the use of both Q-Forecasting and Fare Adjustment methods benefits to DAVN in unrestricted environments, we clearly do not reach the optimal case when using the same input Frat5 for both methods. The optimal Frat5 for Q-Forecasting is quite aggressive, and when we add the Fare Adjustment method using the same Frat5, revenues drop under the base case level. Symmetrically, the optimal Frat5 for is the Fare Adjustment method quite flat, and when we add the Q-Forecasting using the same Frat5, revenues drop under the base case level too. The highest revenue obtained when both method use the same Frat5 set is still under the highest revenue obtained with Q-forecasting alone. The corresponding Frat5 set is at an intermediate level between the optimal - but aggressive - Frat5 set used for Q-Forecasting alone and the optimal - but very flat - Frat5 set used for Fare Adjustment methods alone.

4.4.2. Matrices of inputs

In order to look for a combined benefit from the Q-Forecasting method and from the Fare Adjustment method, we have added in PODS the possibility to use different Frat5 sets for each method. Therefore, we are able to apply relatively high Frat5 to QF, while FA uses lower values because, as explained in Chapter 3, the FA-Sell-up is not conditioned on the closure of the Q-Class (therefore sell-up rates are lower). We are considering here the Frat5 sets presented on *Figure 81* and on *Figure 82*.

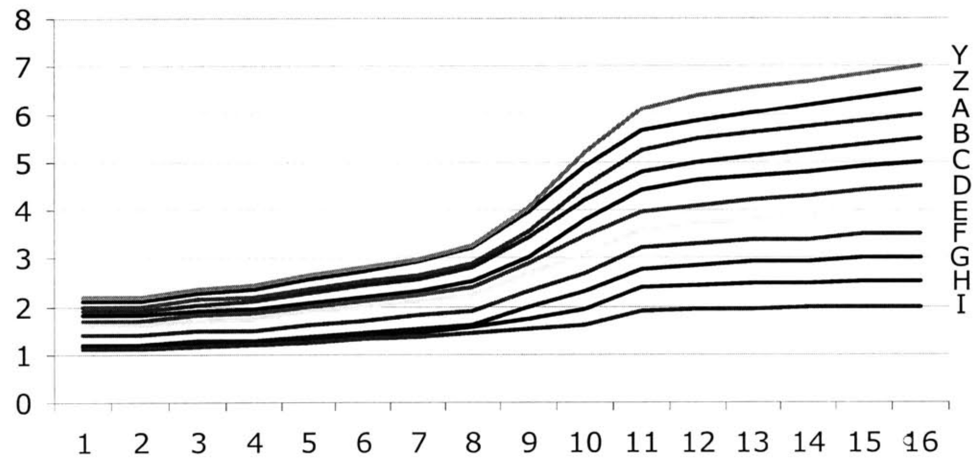

Figure 81: QF Time-Frame Frat5 Sets Used in the Matrices

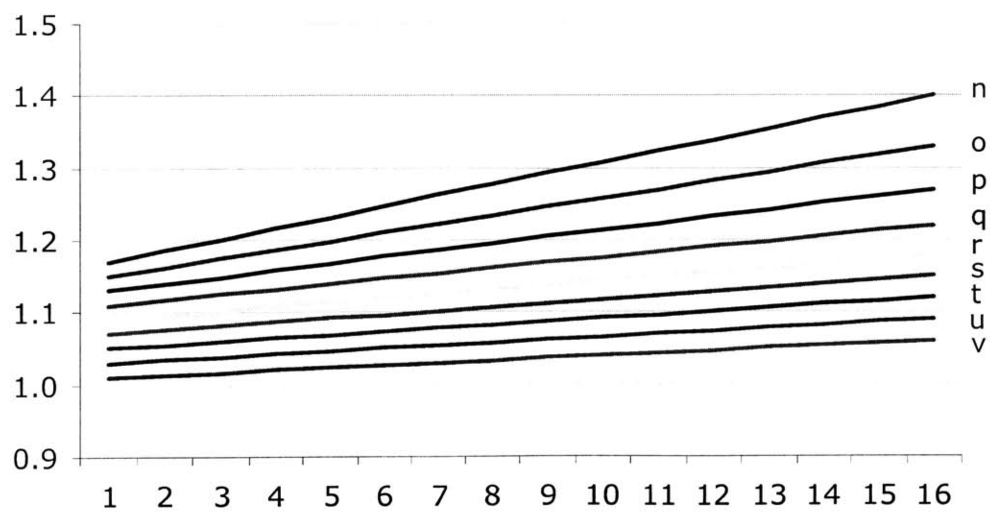

Figure 82: FA Time-Frame Frat5 Sets Used in the Matrices

Figure 83 presents the change in AL1 revenue when the QF Frat5 sets vary from Y to I and the FA Frat5 sets vary from n to v.

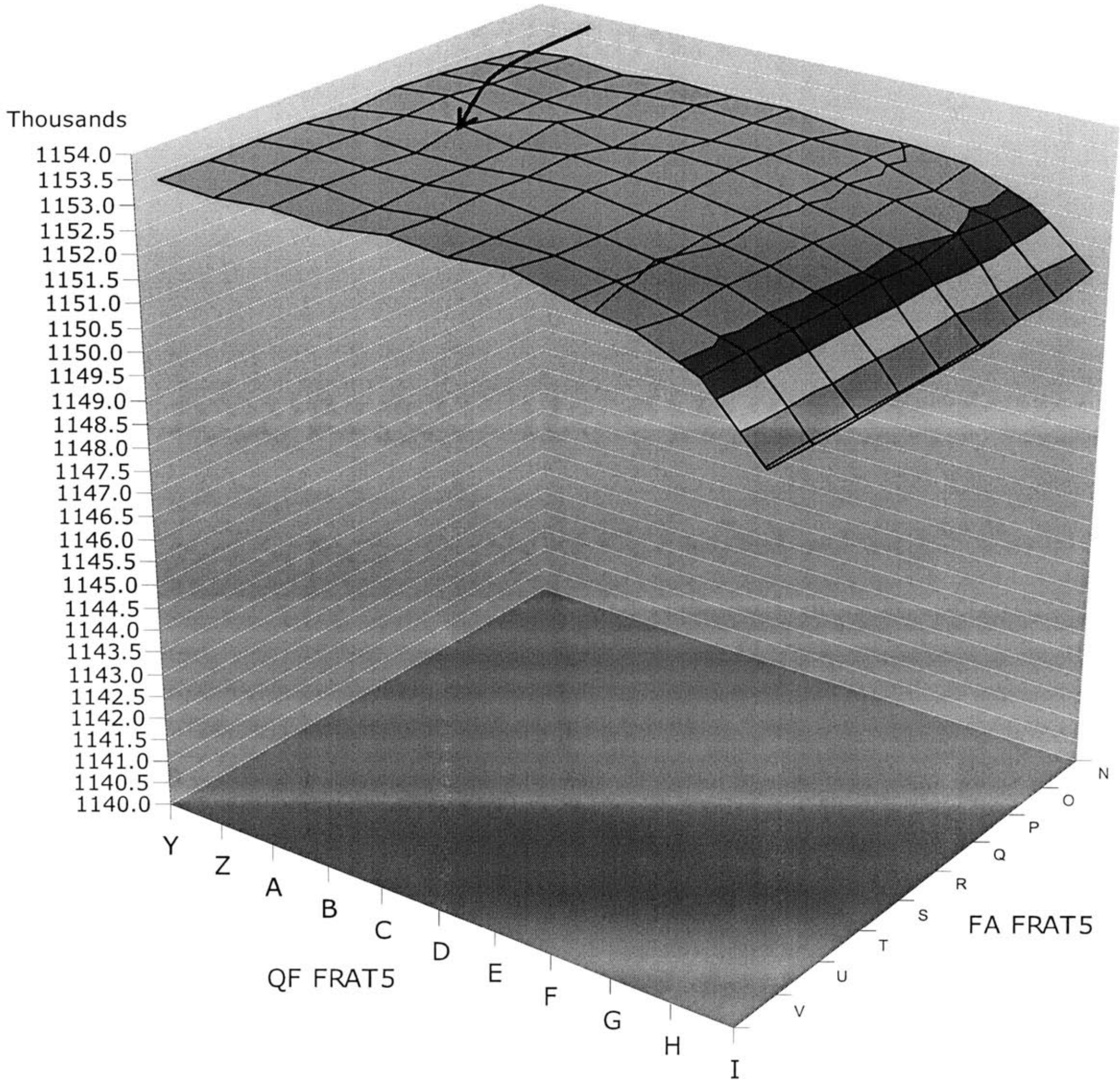

Figure 83: AL1 Revenue Depending on QF and FA Frat5

The obtained "surface" presents a large "plateau" for all QF Frat5 inputs more aggressive than sets G and F. There is little sensitivity on this surface, and generally little sensitivity with respect to the FA Frat5. The revenue obtained in the base case (No QF / No FA) for the same conditions is $1,146,207. The best total network revenue obtained in the matrix above reaches $1,153,461 (sets A,r), which is 0.63% (equals 9.2% revenue increase in the 10 NEM) above the base case.

In order to analyze the impacts of a change in the demand, we change the demand Factor from 1.0 down to 0.9 on *Figure 84*, and from 1.0 up to 1.1 on *Figure 85*, everything else remaining equal.

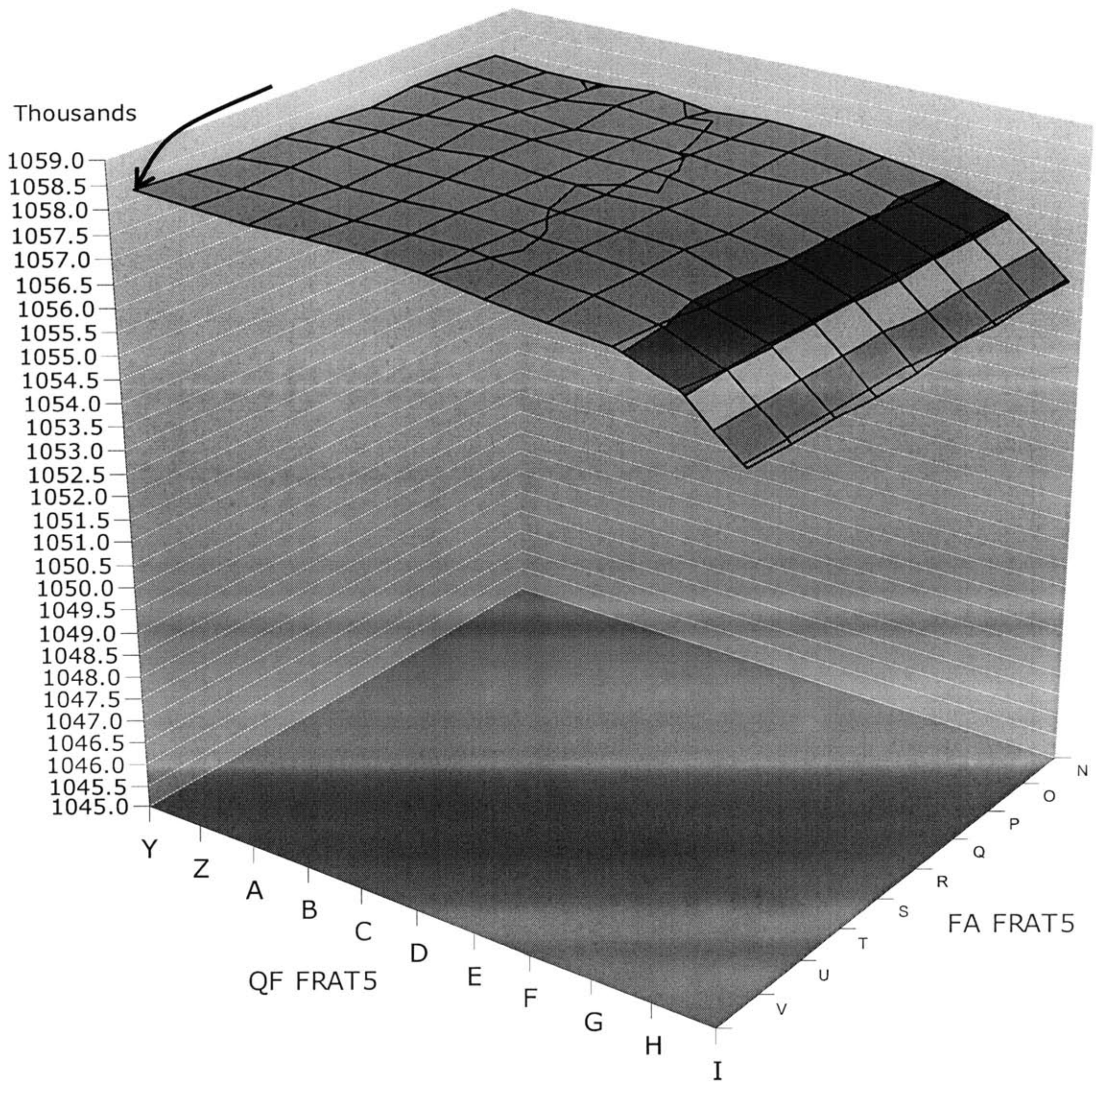

Figure 84: Impact of a Decrease in Demand (DF=0.9)

Again, the surface as a plateau shape, with the same zone of low-sensitivity to both input Frat5. AL1 total network revenues in the base case (No-QF / No-FA, and DF=0.9) were $1,052,650 whereas the optimal of this matrix reaches $1,058,207 (sets Y,v) which is 0.53% higher (equals 7.8% revenue increase in the 10 NEM).

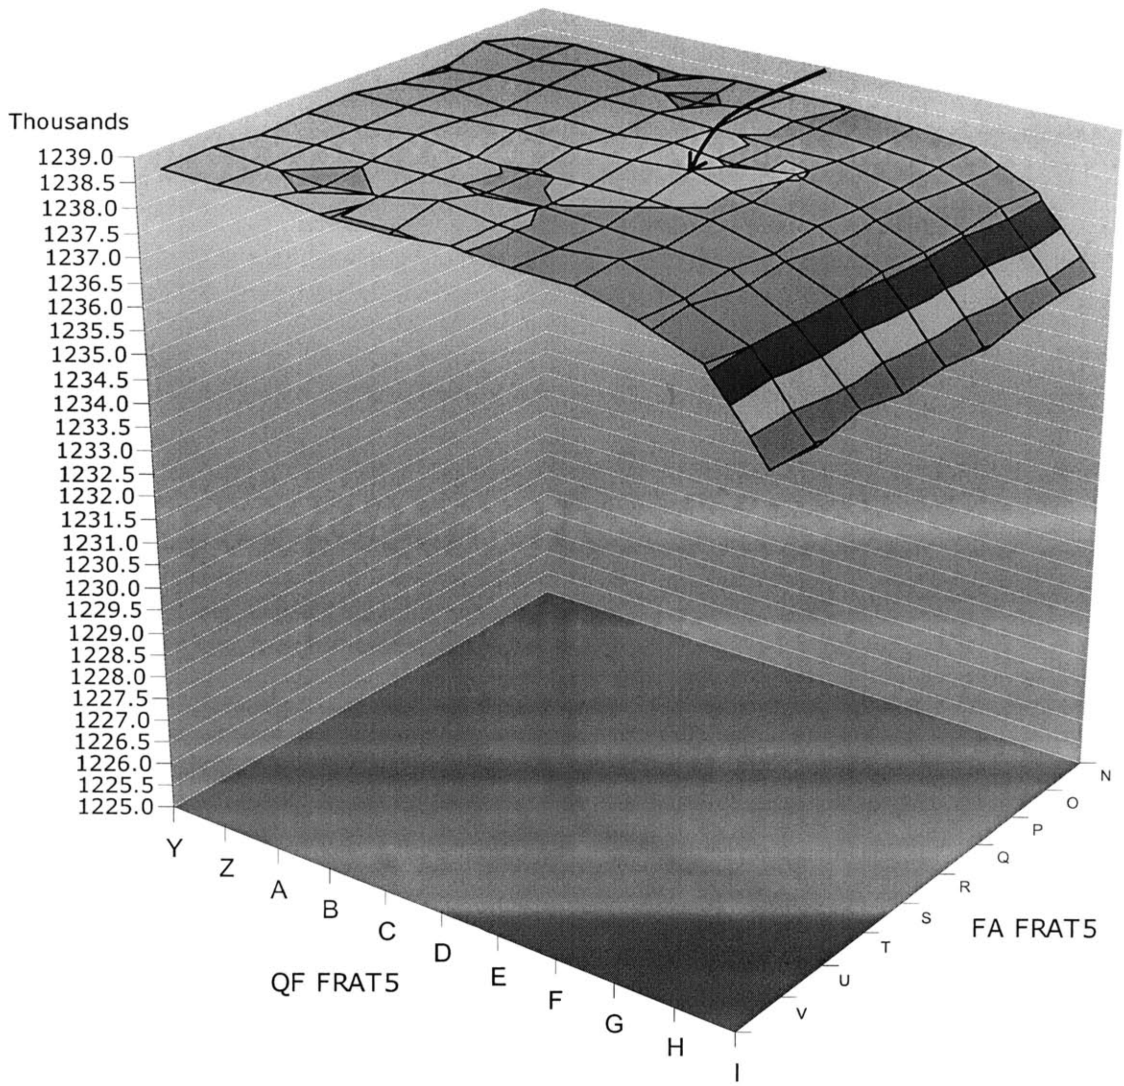

Figure 85: Impact of an Increase in Demand (DF=1.1)

Again, the surface as a plateau shape, with the same zone of low-sensitivity to both input Frat5. AL1 revenues in the base case (No-QF / No-FA, and DF=1.1) were \$1,229,680 whereas the optimal of this matrix reaches \$1,238,714 (sets E,r) which is 0.73% higher (equals 10.7% revenue increase in the 10 NEM). We observe that the optimal Frat5 couple is less aggressive at higher demand, given that it is easier to capture more high-fare class passengers at high demand.

This time we consider another parameter, the capacity of the New Entrant that we decrease from 30 seats down to 20 seats. On *Figure 86* we see the impact on AL1's revenue.

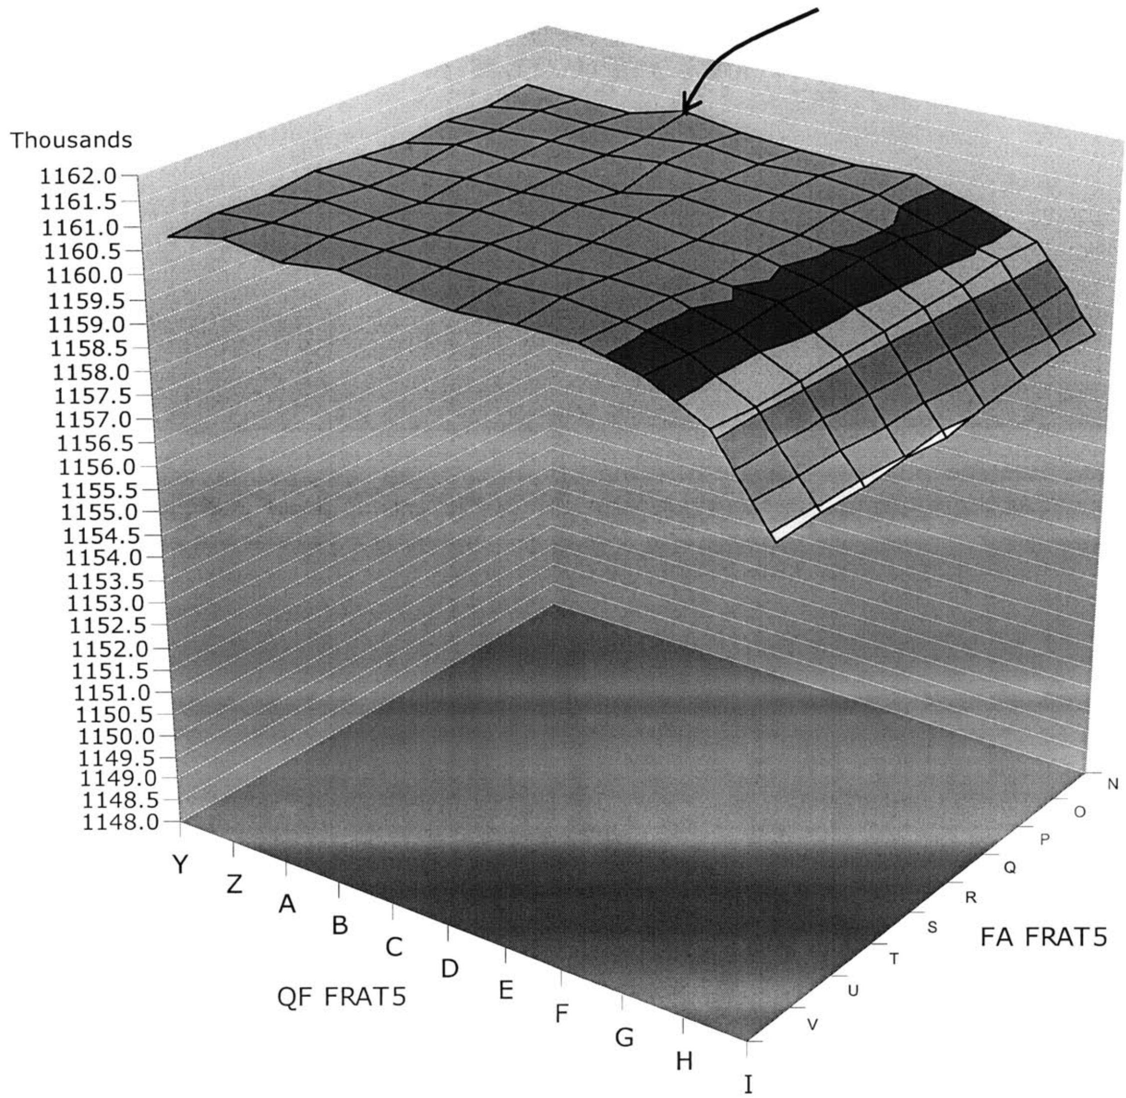

Figure 86: Impact of a Decrease in the LCC's Capacity (from 30 to 20)

Once again, the surface as a plateau shape, with the same zone of low-sensitivity to both input Frat5. AL1 revenues in the base case (No-QF / No-FA, LCC=20 seats) were $1,151,422 whereas the optimal of this matrix reaches $1,161,008 (sets B,n) which is 0.83% higher (equals 12.2% revenue increase in the 10 NEM).

4.4.3. Conclusion

On *Table 14* is presented a summary of the results obtained in this section. It regroups, in each case, the corresponding base case and best case total network revenues, the improvement in % and the Frat5 couples to which the best case refers.

	DF=0.9	DF=1.0	DF=1.1	Cap=20	No LCC
Base Case	$ 1,052,650	$ 1,146,207	$ 1,229,680	$ 1,151,422	$ 1,170,426
Best Case	$ 1,058,207	$ 1,153,461	$ 1,238,714	$ 1,161,008	$ 1,181,383
Improvement	0.53%	0.63%	0.73%	0.83%	0.94%
Frat5 Sets	Y,v	A,r	E,r	B,n	Y,t

Table 14: Matrices Summary

The improvement in total network revenue in the DF=1.0 case is by 0.63%, which is higher than the benefits provided by the fare adjustment method alone and slightly higher and the improvement provided by Q-Forecasting alone. It corresponds to an increase by 9.3% in the 10 NEM. On *Figure 87* is shown the matrix of QF-Frat5 x FA-Frat5 on which are located the optimums of the different cases.

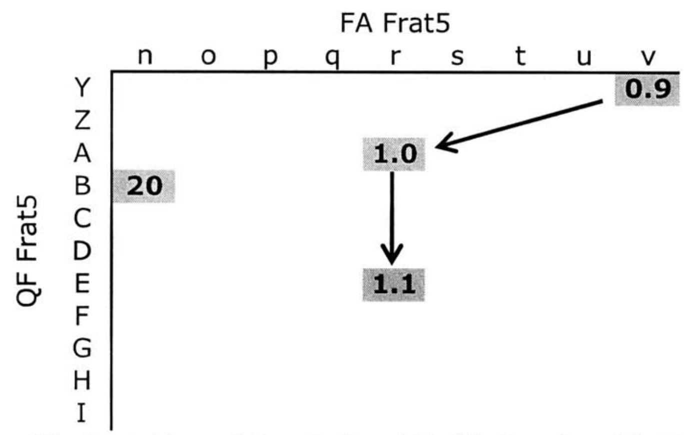

Figure 87: Evolution of the Optimal Frat5 Couple with the DF

When demand increases, more and more high-fare passengers are showing-up. Given that AL1's capacity remains the same, the methods try to be more aggressive and protect more seats for these passengers relative to the lower classes. Consequently, we see that from DF=1.1 to DF=1.0 and then to DF=0.9, the optimum moves toward more aggressive QF-Frat5. The case when the capacity of the LCC is reduced to 20 seats is, to a certain extent, equivalent from AL1's point of view, to an increase in demand, and therefore the optimum tends towards less aggressive Frat5 sets as well (with respect to the DF=1.0 case).

Like in the case of Q-Forecasting, we need a sell-up estimator to be able to use the Fare Adjustment methods in the real world. The use of both FA and QF Frat5 estimators should bring additional benefits, even if their magnitude may not be as high as the improvements we obtained with the input Frat5.

5. Conclusion and Future Research Directions

5.1. Summary of Findings

The objective of this thesis was to explain and measure the performance of two new alternatives aimed to help airline Revenue Management methods to deal with unrestricted fare structures and recover from the spiral-down of revenue that traditional RM methods lead to. We explained the mechanisms of each method, the way they are linked to each other and why they are complementary and not redundant. We also described how they are implemented in the Passenger origin Destination Simulator (PODS) with which the performance measures and analyses would be done. It was also the goal of the thesis to present diverse analyses obtained from the simulations results to highlight the mechanisms presented before.

The results of the simulations show that it is possible to partially recover from the drop caused by the removal of the fare fences in a subset of markets, using either the Q-Forecasting, the Fare Adjustment method, or better, both simultaneously. The two methods act at different steps in the process of DAVN; a network Revenue Management method. The first acts in the forecasting process, addressing the problem raised by the violation of the fundamental "independence of demand by fare class" assumption of traditional Revenue Management systems. The second acts within the booking limit optimizer, introducing the Marginal Revenue as input to the optimizer instead of the OD fares, accounting for the sell-up probability. Both methods indeed account for sell-up potential, which is the base for the concepts presented in this thesis. The thesis shows as well that it is possible to build a sell-up estimator – in the case of Q-Forecasting – which, based on historical records, can predict the sell-up rates and can provide satisfying results compared to the version in which we input the rates.

On *Figure 88* is presented a summary of the contributions brought by each method separately, and then together, assuming in each case the use of the Frat5 set which led to the highest revenue in its own category. The Frat5 sets are a series of Fare ratios, relative to the Q fare, at which 50% of passengers are willing to buy a ticket at a given day of the booking period. They represent the sell-up potential in PODS. Theses cases correspond to results shown and discussed in this thesis, but it is interesting to compare them from a broader point of view.

In all cases, we consider AL1 and AL2 using DAVN (Displacement Adjusted Virtual Nesting), AL3 using AT (Adaptive Threshold, a dynamic version of the Threshold Revenue Management method) with an 80% load factor target and a demand factor equal to 1.0. The methods – Q-Forecasting and Fare Adjustment – are used only by AL1. The following figure reports AL1's revenues, using the methods as indicated on the X-axis.

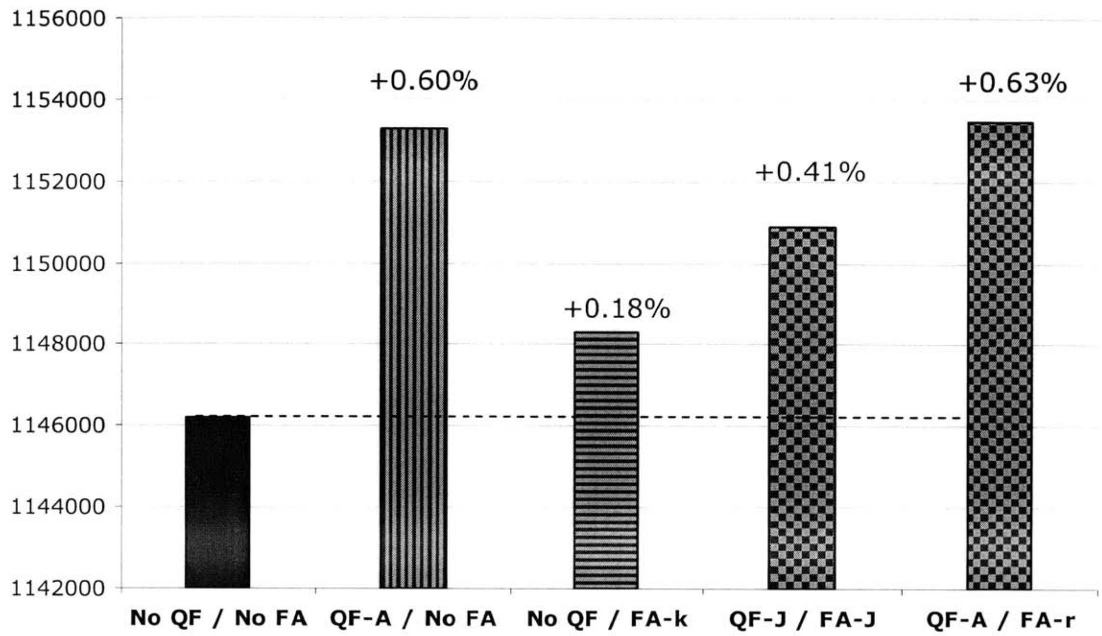

Figure 88: AL1 (DAVN) Revenue in Each Case Using QF and/or FA

The first bar represents the base case, without Q-Forecasting or Fare Adjustment methods (resulting from the drop in total network revenues by 8% due to the entrance of a low-cost carrier in the 10 non-stop markets).

The second bar - vertical stripes - corresponds to the case when the airline uses only Q-Forecasting, with the input Frat5 set "A". In this case, revenues are higher by 0.60% (equals 8.8% revenue increase in the 10 NEM). The improvements obtained with the sell-up estimator for QF-Frat5 (without FA) were just slightly below this one, around 0.50%.

The third bar - horizontal stripes - represents the case when the airline uses the Fare Adjustment method alone, with input Frat5 "k". The benefits in terms of total network revenue with respect to the base case are by 0.18% (equals 2.6 % revenue increase in the 10 NEM).

When both methods are using the same input Frat5 "J" simultaneously - 4th bar -, the improvement reaches 0.41% (equals 6% revenue increase in the 10 NEM).

Finally - last bar - when the airline uses DAVN with both methods and different Frat5 sets, we obtain the highest improvement in revenue, by 0.63% (for the QF-Frat5 "A" and the FA-Frat5 "r" - equals 9.2% revenue increase in the 10 NEM).

Even if the improvement difference is not very large between Q-Forecasting alone and the use of both methods simultaneously (from 0.60% to 0.63%), the second case actually leads to the highest increase in revenue. Moreover, this configuration appears to be very stable in terms of sensitivity to Input Frat5, whereas each method used independently was more sensitive.

Both methods manage – to some extent – to capture more higher-class local passengers on the 10 legs with unrestricted fare products by closing down earlier the lower virtual buckets in the DAVN process. Moreover, they lead to network effects by increasing the number of connecting passengers (given that the Low-Cost Carrier still captures a significant number of local passengers). Q-Forecasting especially increased the number of high-class connecting passengers, whereas Fare Adjustment method increased the loads in all connecting classes.

We have shown in the thesis in detail the mechanisms involved in the different methods, based on representations of estimated Frat5, adjusted fares, sensitivity analyses with respect to demand or to input Frat5 sets.

5.2. Future Research Directions

The first suggestion for future research directions is a direct follow-up to this thesis. The estimation of the various forms of sell-up based on historical data is crucial. It is indeed the only way for implementing these two alternatives in the actual Revenue Management Systems used by airline companies in the "real world". We have seen that the Q-Forecasting Frat5 estimator performs relatively well compared with the Q-Forecasting input Frat5. An equivalent (in terms of performance as well) is needed for the estimation of the Fare Adjustment Frat5. This estimator is currently in development in PODS but it is not yet dramatically competitive with the input Frat5. The best results are obtained when we combine the 2 methods, so the real world needs both estimators.

The environments in which we tested the 2 methods bring the second research direction. It would be indeed interesting to measure the performance of these methods in fare product structures which are not entirely unrestricted. This corresponds to cases when a same airline offers some differentiated and some undifferentiated fare products in the same structure, such as in the US where "SimpliFares" are now common. In this situation, Q-Forecasting and Fare Adjustment should be applied to the undifferentiated classes whereas the restricted classes would be managed by traditional Revenue Management methods.

A parallel direction lies into the more and more obvious segmentation of demand between product oriented and price-oriented demand. Ideally, one would like to forecast each type of demand independently, in order to capture the maximum revenue from the product-oriented demand (by managing them with traditional revenue Management methods) and then apply the alternative methods presented in this thesis to the price-oriented method for which the traditional method do not perform very well. Using either only traditional Revenue Management methods or only Q-Forecasting and Fare Adjustment methods for both types of demand would clearly not be optimal. The problem is obviously to properly estimate the demand by type, and then to be able to offer the proper product to each demand.