Revenue Management Based on Dynamic Programming in Unrestricted and Simplified Fare Structures

By

Thierry Vanhaverbeke

Ingénieur des Arts et Manufactures (ECP 2005)

Submitted to the Department of Civil and Environmental Engineering
in partial Fulfillment of the Requirements for the Degree of Master of Science
in Transportation

At the

Massachusetts Institute of Technology

June 2006

1st Copy

Signature of Author ....
Department of Civil and Environmental Engineering
January 27, 2006

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

Accepted by ....
Nigel Wilson
Chairman, Transportation Education Committee

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

Revenue Management Based on Dynamic Programming in Unrestricted and Simplified Fare Structures

By

Thierry Vanhaverbeke

Submitted to the Department of Civil and Environmental Engineering on January 27, 2006, in partial Fulfillment of the Requirements for the Degree of Master of Science in Transportation

ABSTRACT

In the past few years, many low-cost airlines have been created around the world. Legacy carriers used to apply revenue management (RM) to fully-restricted fare structures with segmented demand. These low-fare competitors with simplified or unrestricted fare structures often manage to capture an important part of the demand in the markets they enter, forcing legacy carriers to change their fare structures. However, because of this change in fare structures, most RM forecasters and optimizers that were previously used by the legacy carriers are no longer effective.

The goal of this thesis is to describe and test two RM methods based on dynamic programming in fare structures with few or no restrictions. The first one, the "Lautenbacher" approach (DP-LB), was developed for fully-restricted environments but, when associated with an appropriate forecaster, may be used in less-restricted fare structures. The second one, the "Gallego-Van Ryzin" approach (DP-GVR), was developed for unrestricted fare structures. In this thesis we describe these two methods, and analyze the performance of both in comparison with the performance of traditional RM methods used by legacy carriers, based on results obtained with the Passenger Origin-Destination Simulation (PODS).

In our simulation DP-LB leads to results very close to those obtained with traditional RM methods. In unrestricted fare structures DP-GVR appeared to be theoretically appealing. The simulation results obtained with DP-GVR against a competitor using no RM or against a competitor using the same RM method were better than those obtained with other RM methods. But against advanced competitors, DP-GVR got worse results than other RM methods. We showed that this variability in the performance of DP-GVR was related to its sensitivity to forecasts of probabilities of sell-up, forecasts that are difficult to estimate because they depend on the RM methods and actions of the competition. Further study with improved estimation of probabilities of sell-up will be required to determine whether this method can lead to stable improvement over traditional RM methods in unrestricted fare structures.

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

Acknowledgments

First of all, I would like to thank my thesis and research advisor, Dr Peter Belobaba, for his valuable help during the time I spent at MIT. He has helped me understand the basics of revenue management and introduced me to the small group of initiated people who know how to use the Passenger Origin-Destination Simulation that gave me the experimental basis for this thesis. With a lot of patience he gave me precious advice to improve the presentations for the PODS meetings and to achieve this thesis.

I would also like to thank Craig Hopperstad for his support, for his patience when answering by email with specific explanations all the questions I asked him and for his helpfulness when there was any need for "hard-wiring" concerning the simulator. His role is crucial in the development of the Passenger Origin-Destination Simulation.

This research is funded by the PODS Consortium, composed by several airlines and I would like to thank the revenue management people at these airlines for their helpful comments during the PODS meetings.

Working with a team of students in the International Center for Air Transportation lab was a very nice experience and I would to thank all those I met there and especially the other students who were part of the PODS research. Thanks to Richard Cléaz-Savoyen who taught me what I needed to know to continue his research and to Maital Dar, Gregory Zerbib, Valenrina Soo, Mike Reyes and Emmanuel Carrier. Special thanks to Jon, Jeff, Maital and Mike for accepting to proof-read portions of my thesis.

Finally, I would like to thank my family and friends, for supporting me throughout the time I was away from France, and Valérie, for her understanding, love and support.

TABLE OF CONTENTS

ACKNOWLEDGMENTS 5

TABLE OF CONTENTS 6

LIST OF FIGURES 8

LIST OF TABLES 10

ABBREVIATIONS AND NOTATIONS 11

DEFINITIONS 12

1. INTRODUCTION 13

1.1. Revenue Management in the Airline Industry 13

1.2. Fare Structures and Low-cost Competitors 15

1.3. The Concept of Sell-up 16

1.4. Algorithms for Revenue Management 17
1.4.1. Static Algorithms 18
1.4.2. Dynamic Algorithms 19

1.5. Goal and Structure of the Thesis 20

2. LITERATURE REVIEW 21

2.1. Traditional Revenue Management Methods 21

2.2. RM Methods Based on Dynamic Programming 24

2.3. Standard DP (Lautenbacher) 25

2.4. The Gallego-Van Ryzin (GVR) Model 26

3. THE PASSENGER ORIGIN DESTINATION SIMULATION 29

3.1. General Presentation 29
3.1.1. The Networks 29
3.1.2. The Demand 31
3.1.3. The Runs 33

3.2. Detruncation and Forecasting **34**
3.2.1. Traditional Methods 34
3.2.2. Q-Forecasting 35
3.2.3. Hybrid Forecasting 38

3.3. Traditional RM Methods in PODS **40**
3.3.1. Threshold RM Methods 40
3.3.2. EMSRb 41
3.3.3. DAVN 42

3.4. DP RM Methods in PODS **42**
3.4.1. DP-Lautenbacher 43
3.4.2. DP-GVR 43
3.4.3. Example based on DP-GVR 44

4. RESULTS **46**

4.1. Unrestricted Fare Structures **46**
4.1.1. Base Case in Network D 46
4.1.2. Against a Traditional RM Method (Network D) 47
4.1.3. Against the Same Method (Network D) 49
4.1.4. Against EMSRb with Q-Forecasting (Network D) 51
4.1.5. Q-Forecasting and Traditional RM Methods 55
4.1.6. Against AT80 (Network D) 60
4.1.7. Sensitivity of DP-GVR to Input FRAT5s 62
4.1.8. Summary of Results Obtained in Unrestricted Network D 63

4.2. Results in Unrestricted Network R **64**

4.3. Simplified Fare Structures **71**
4.3.1. Base Case in Network D 71
4.3.2. Benefit of Hybrid Forecasting for EMSRb and DAVN 72
4.3.3. Use of Hybrid Forecasting with DP 73
4.3.4. Summary of Results in Semi-restricted Network D 75

5. CONCLUSION **77**

5.1. Summary of Findings **77**

5.2. Further Studies **79**

LIST OF FIGURES

Figure 1: Nested protection of seats 15
Figure 2: Traditional Revenue Management Process 18
Figure 3: States and transitions 19
Figure 4: The Lautenbacher DP 26
Figure 5: The Gallego-Van Ryzin DP 28
Figure 6: Single market 29
Figure 7: Network of Airline 1 30
Figure 8: Network of Airline 2 30
Figure 9: Network R 31
Figure 10: Arrival curves 32
Figure 11: PODS booking process for traditional RM methods 33
Figure 12: Booking Curve Detruncation 34
Figure 13: Q-Forecasting Process 36
Figure 14: Shape of the probability of sell-up 37
Figure 15: Input FRAT5s 37
Figure 16: Scope of Hybrid-Forecasting rules 39
Figure 17: Classification of AP bookings with the path rule of Hybrid Forecasting 39
Figure 18: Determination of Booking Limits with EMSRb 41
Figure 19: Fare class mix of Airline 1 – Base Case in unrestricted Network D 47
Figure 20: Results obtained against EMSRb with regular forecasting 47
Figure 21: Comparison of the local fare class closures of the 2 airlines 48
Figure 22: Fare class mixes of Airline 1 using EMSRb with QF or DP-GVR 49
Figure 23: Results obtained against the same method 50
Figure 24: Local fare class closure of both airlines using DP-GVR 50
Figure 25: Fare class mixes of both airlines using EMSRb with QF or DP-GVR 51
Figure 26: Results obtained against EMSRb with QF 52
Figure 27: Comparison of the local fare class closures of the 2 airlines 52
Figure 28: Fare Class Mixes obtained when Airline 2 uses EMSRb with QF 53
Figure 29: Forecasts of equivalent Q-demand to come for DP-GVR with QF 54
Figure 30: Various fare class closures and their corresponding fare class mixes 55
Figure 31: Obtaining equivalent Q-demand for EMSRb in one time-frame 56
Figure 32: Obtaining all Potential demand from estimated Q-demand 57
Figure 33: Observed demand-to-come for a given departure 58
Figure 34: Currently used forecasted demand-to-come with EMSRb 58
Figure 35: DP-GVR determination of lowest open fare class 59
Figure 36: Getting the right bookings-to-come 60
Figure 37: Results obtained against AT80 61
Figure 38: Comparison of the local fare class closures of the 2 airlines 61
Figure 39: Fare Class Mixes obtained when the 2nd airline uses AT80 62
Figure 40: Fare Class Closure and Cumulative Bookings by class 63
Figure 41: Fare Class Mix of Airline 1 64
Figure 42: Revenues and load factors of all 4 airlines 64
Figure 43: Revenues of all 4 airlines when Airline 1 uses various RM methods 65
Figure 44: Revenues of all 4 airlines when using various RM methods (FRAT5s A) .. 66
Figure 45: Results of Airline 1 - EMSRb with QF 66
Figure 46: Results of Airline 1 - DP-LB with QF 66
Figure 47: Results of Airline 1 - DP-GVR 67
Figure 48: Fare Class Mix of Airline 1 67
Figure 49: Revenues and load factors of all 4 airlines 67

Figure 50: Revenues of all 4 airlines when using various RM methods 68
Figure 51: Revenues of all 4 airlines when using various RM methods (FRAT5s A) .. 69
Figure 52: EMSRb with QF against Eb-QF/AT80/Eb-QF 69
Figure 53: DP-LB with QF against Eb-QF/AT80/Eb-QF 69
Figure 54: DP-GVR against Eb-QF/AT80/Eb-QF 70
Figure 55: Fare class Mix of Airline 1 – Base Case in semi-restricted Network D 72
Figure 56: Revenues of Airline 1 using EMSRb with HF against FCFS 72
Figure 57: Revenues of Airline 1 using DAVN with HF against FCFS 73
Figure 58: Fare class mixes of Airline 1 using HF (FRAT5 A4) 73
*Figure 59: Fare class Mix of Airline 1 – DP Base Case in semi-restricted Network D*74
Figure 60: Revenues of Airline 1 using DP with HF against Eb with std. Fcst. 74
Figure 61: DP/HF1 (IAP0) vs DP/HF1 (IAP0) 75
Figure 62: Historical data 81

LIST OF TABLES

Table 1: Fare Structures 32
Table 2: Average fares in Network D 33
Table 3: Threshold percentages 41
Table 4: Probabilities of sell-up in each time-frame 44
Table 5: Computing the expected revenue for DP-GVR 44
Table 6: Lowest fare class to be open in each decision period 45
Table 7: Summary of all best revenue increases over the base case in Network D .. 63
Table 8: Simplified Fare Structure 71
Table 9: Results obtained in semi-restricted network D 75

ABBREVIATIONS AND NOTATIONS

AT	Adaptive Threshold
AP	Advance Purchase
BKG	Bookings
BTC	Bookings-To-Come
DAVN	Displacement Adjusted Virtual Nesting
DP	Dynamic Programming Method
DP-GVR	Dynamic Programming Approach introduced by Gallego-Van Ryzin
DP-LB	Dynamic Programming Approach introduced by Lautenbacher
Eb	EMSRb
EMSRb	Expected Marginal Seat Revenue Method
F_i	Fare of fare class i, with $F_i > F_{I+1}$
FA	Fare Adjustment
FC	Fare Class
FCFS	First Come First Served
FCST	Forecast
FT	Fixed Threshold
Frat5	Fare Ratio to Q at which 50% of the passengers accept to sell-up
HF	Hybrid Forecasting
KI	Karl Isler
LCC	Low-Cost Carrier
LP	Linear Program
NEM	New Entrant Market(s)
O-D	Origin-Destination
PAX	Passengers
PODS	Passenger Origin-Destination Simulator
QF	Q-Forecasting
R_i	Restrictions i
RM	Revenue Management
TR	Total Revenue
TF	Time-Frame
WTP	Willingness-To-Pay

DEFINITIONS

Bookings-to-come

Forecast of bookings occurring in the current time-frame and in all the following time-frames until departure

Detruncation

Transforming actual bookings to potential demand by estimating what the demand would have been if there were no capacity constraints

Fare class

Product proposed by an airline and associated with a fare and restrictions. Booking in a given fare class may be possible only at given moments of the reservation period

Load factor

Number of passengers divided by the number of seats

O-D market

Demand corresponding to all the passengers wanting to travel from a given origin to a given destination during a given period of time.

Reservation period

The period preceding the departure of a flight, during which passengers can book tickets for seats made available by the RM systems of competing airlines

Revenue management

Control of seat inventory that includes forecasting demand and optimizing the pricing of seats in order to maximize revenues

Sell-up

Decision to buy a highest fare than the desired one

Yield

Average fare paid by passengers, by mile flown

1. INTRODUCTION

The goal of this thesis is to test the efficiency of revenue management methods based on dynamic programming in unrestricted and simplified fare structures. We will use a software package developed by Hopperstad, the Passenger Origin Destination Software, which we will present in Chapter 3 to simulate both airline competition and passenger demand in defined markets.

1.1. Revenue Management in the Airline Industry

Revenue management is part of the airline planning process. We first define the role of revenue management in this process, which precedes the booking of a seat by a passenger. This short presentation is inspired by **Barnhart, Belobaba and Odoni**[1].

Fleet Planning is a long-term planning that corresponds to the management of the aircraft fleet of an airline. Fleet planning evaluates the required number of aircraft and their types according to the fit between the operating characteristics of the available aircraft and the market requirements. Specifying the flight legs that will be served and the departure time of each flight corresponds to *Network Design* and *Schedule Design*. The airlines choose to serve various routes and they decide which frequencies and schedules they want on those routes.

Fleet Assignment assigns aircraft types to all the flights the airline has decided to schedule. This may be done in a way that minimizes the assignment costs, which include both operating costs, which are costs of flying a given flight leg with a given aircraft type, and spill costs, which are revenue loss when there is more demand than capacity on a given leg. *Aircraft Maintenance Routing* makes sure to route all the individual aircraft while satisfying all maintenance requirements. Then *Crew Scheduling* assigns crews, which include pilots and sometimes flight attendants, to the aircraft while trying to minimize crew costs. The crew allocation makes sure a crew is available for each flight at a minimum cost while satisfying many work rules.

Pricing corresponds to selecting a set of prices that will be charged for seats. Indeed not all seats are available at the same price. The airline will offer different fares on each O-D market. Many passengers think that the airlines are constantly changing their prices and that they may choose any fare for a given O-D market while in fact the airlines only have a set of available fares and they are closing fare classes at given points in time.

Revenue Management corresponds to the control of the inventory of seats. After having chosen a set of fare classes the airline decides how many seats it wants to make available in a fare class at a given time in the reservation process. If all fare classes were available under the same conditions people would of course buy the cheapest tickets. In order to differentiate between the fare classes the airlines use restrictions and advance purchase requirements. This practice is called *Differential pricing*. Each airline has a central reservation system which records all bookings and

---

[1] BARNHART, BELOBABA, ODONI, *Applications of Operation Research in the Air Transport Industry*, Transportation Science, Vol. 37, 2003

provides the number of seats that are still available at a given moment. The duty of revenue management is to provide the central reservation system with parameters that are used to control inventories. This task is called *Seat Inventory Control*.

Through the reservation process some fare classes can be closed according to the number of seats that have already been booked in order to save seats for higher-fare-paying passengers. Revenue Management determines at any time during the reservation process whether a request for a seat in a given fare class should be accepted or not in order to maximize revenues. The difficulty is to get a good load factor on a flight without losing passengers with high willingness-to-pay. Indeed in comparison with other industries revenue management for airlines considers perishable goods: if a seat is not sold prior departure the corresponding good is lost and it will bring no revenue to the airline. Thus if too many seats are saved for high-yield passengers then planes may depart with empty seats which represent a loss of revenues. Since business passengers tend to book much closer to departure than leisure passengers the airlines can gradually close lower fare classes to capture this demand with higher willingness-to-pay. By using such fare structures with different restrictions and fares the airlines manage to partially segment the demand. Yet if they let too many people book in the lower fare classes then they may lack seats to capture all the demand with high willingness-to-pay that comes at the end of the booking process.

At a given time in the reservation process concerning a specific flight, there is a given number of seats that are still available. A certain amount of this remaining capacity will be made available for each fare class. Partitioned inventory would correspond to allocating a certain number of seats to each fare class. It is easy to see the drawbacks of such a method. With partitioned inventories, higher fare classes may be closed while seats are still available for lower fare classes. That is, someone may not be able to book a seat in a high fare class while there are still seats available! This would result in potential revenue loss.

Instead most airlines are using nested booking inventory control (*Figure 1*). For a given fare class, seats are protected from bookings in lower fare classes. For example if 20 seats are saved for fare class 3 then only passengers from higher fare classes, 1, 2 and 3, will be able to book those seats. A request for fare class 1 will always be accepted as long as there are seats remaining.

Recently most US legacy airline companies have had to adapt to new market conditions, which include new distribution channels, continuing growth of low cost carriers, decreasing willingness-to-pay of business passengers, the impact of September 11, and increasing cost of fuel, so any increase in revenues can be very beneficial. Among the most important changes that legacy carriers have to face are new fare structures with fewer restrictions and compressed fares. In markets where legacy carriers have to face the entrance of a low-cost carrier, the legacy carriers have to choose either to match the fare structure of their low-cost competitor or to keep their old fare structure. Most legacy aircraft deciding to match have to adapt their revenue management system to those new fare structures and that is the reason why we will focus our study in markets with unrestricted and simplified fare structures.

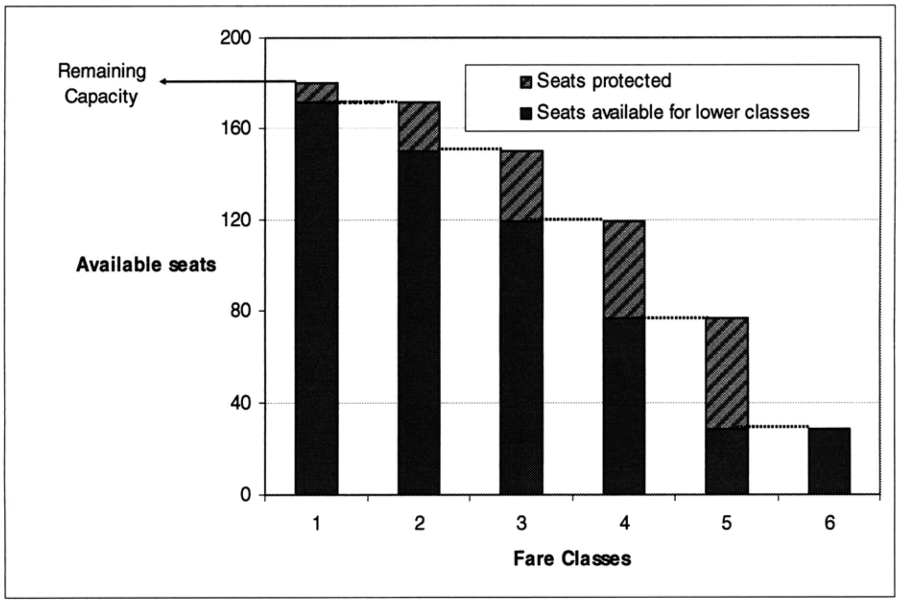

Figure 1: Nested protection of seats

1.2. Fare Structures and Low-cost Competitors

Legacy carriers use restrictions and advance-purchase requirements to segment demand. These restrictions can, for example, be Saturday night requirements, change fees, or non-refundability of cancelled tickets. Advance purchase corresponds to the automatic closure of a fare class at a given point of the reservation period. Business passengers are more sensitive to those restrictions. It can motivate their decision to pay a higher price.

Before the deregulation of the US airline industry in 1978, the Civil Aeronautics Board controlled all fares. The fares were based on distance and on the average industry costs. Consequently airlines were mainly competing on service, trying to capture passengers trough better quality of service and higher frequency. Their fare structures were stable and simple as the airlines had not much flexibility in establishing those fare structures. After deregulation fare structures became much more complex as airlines were then free to decide of their fare structure and fares. In each market the airlines established their own fare structure. They used different restrictions for each fare class in order to segment demand and the use of differential pricing became crucial. For many years most legacy carriers then used fully-restricted fare structures. For details about restrictions, business class and first class please refer to the work of **Bohuntisky**[2].

In recent years low-cost companies have developed in the US and elsewhere around the world with lower fares than legacy carriers and with less restricted or even unrestricted fare structures. In many cases legacy carriers have to match at least partially their new low-cost competitors on fare products if they do not want to lose markets. Not matching may lead to spiral-down, meaning that all passengers do not

[2] BOHUTINSKY, *The Sell-up Potential of Airline Demand*, MS thesis, MIT, 1990

want to sell-up to higher fare classes anymore and prefer to buy-down in the lowest open fare classes which leads airlines using traditional optimizers to accept more and more people in the lowest fare classes.

Many low-cost carriers do not segment demand by using those restrictions. They only close fare classes at different times of the reservation period. As low-cost carriers have no restriction such as Saturday night requirements more people are able to buy low fares. These unrestricted fare structures are attracting demand but this demand is less willing to sell-up to pay a high fare. So when a low-cost carrier is entering a market where there were only legacy carriers, the low-cost carrier may manage to capture an important part of the existing demand by allowing people to book lower fares at any point of the reservation period. On the other side the legacy carriers see the willingness-to-pay of their passengers decrease.

One solution for legacy carriers to partially match the low-cost competitors is to use simplified fare structures. Those fare structures have fewer restrictions and compressed fares, which means that the ratio of the higher fare to the base fare is lower than for regular fares in fully-restricted environments. In this thesis we will study what is happening both in unrestricted fare structures and simplified fare structures.

1.3. The Concept of Sell-up

People wishing to go from a given origin to a given destination at a given point in time are often offered several ways to travel. There may be other ways to travel than air transportation: highways, trains, buses. But even if we only focus on the air transportation in most cases passengers are facing various possibilities and have to take the decision of the best one to pick.

We first present the approach based on a traditional view of revenue management in fully-restricted fare structures as introduced by **Charania**[3]. In this approach the passenger may have an idea of the fare he would like to pay, the airline he would like to travel with and whether or not he is ready to take a ticket that requires a connection and we can consider that he will check on the phone or in a travel agency if his desired fare is available. If his desired fare is available then the passenger may decide to directly book a ticket or to first have a look at other existing fares in order to maybe find a cheapest one. If the desired fare is not available in most cases the passenger will look at other possibilities:

- He may look for a seat on another flight in the same fare class and with the same airline.
- He may look for a seat on another flight with or without connections in the same fare class but with a different airline.
- He may choose not to travel.
- He may decide to look at the next fare class for the same flight in the same airline. If this fare is still not available then the passenger may choose any of the 3 previous possibilities.

[3] CHARANIA, *Incorporating Sell-up in Airline Revenue Managements*, MS thesis, MIT, 1998

We see that if his desired fare is not available then the passenger may decide to sell-up in a higher fare class to the same airline or to a different airline. These possibilities of sell-up are ignored by most traditional RM methods.

In simplified and unrestricted fare structures the consumer decision process has changed. In such fare structures, most passengers would like to travel at the lowest available price and have no precise idea of the fare product they are looking for. One passenger may be ready to pay a higher price in order to avoid connections, to travel with his preferred airline or to get some advantages (refundability for example) but the differences between the various fares he is offered are less important than in fully-restricted fare structures and consequently they are not driving his choice. There is a given demand for each O-D market that would be ready to travel at the lowest fare offered by the airlines on this O-D market. We can then consider that there are particular probabilities that a given person who would travel in this O-D market at the lowest available fare will sell-up to book a ticket in a given fare class and that these probabilities of sell-up depend on several parameters describing the market situation.

Bohuntisky[4] studied the impact of sell-up on overall revenues. According to her study the probabilities of sell-up are higher when demand is higher for specific flights or periods of the day (peak periods). This is consistent with the fact that when demand is higher there is a larger amount of passengers with high willingness-to-pay and then for the same capacities airlines may be able to capture more demand with high willingness-to-pay while trying to spill demand with low willingness-to-pay in other airlines. Sell-up is also more important in higher fare classes. Finally passengers are more willing to sell-up within an airline when this airline dominates markets while in competitive markets they can decide to fly with other airlines.

1.4. Algorithms for Revenue Management

There are two main parts in managing the inventories on an airline network: forecasting demand, either on legs or on origin-destination paths, and closing fare classes at the right time, according to those forecasts. Traditionally these two parts are done separately and forecasts are used as input for the optimization part (*Figure 2*).

In RM systems, forecasts are revised throughout the reservation process, based on the number of current bookings that have been made at a given time. So the period before departure is divided into time frames between which new values for remaining capacities in each fare class are computed. These forecasts are then unconstrained to take into account all the demand that was unable to book because of capacity constraints. If a fare class is closed nobody will be able to book in the class and the recorded demand will be less than the actual demand. Consequently if this data was used by the forecaster the demand for the class would be underestimated.

The unconstrained forecasts are then used by the optimizer to determine which policy should be followed during the reservation period with the aim to reach maximum revenues. The forecasts can be forecasts of bookings occurring within a time-frame or forecasts of bookings-to-come, which means forecasts of all bookings

[4] BOHUTINSKY, *The Sell-up Potential of Airline Demand*, MS thesis, MIT, 1990

occurring between now and departure. Most traditional RM methods use forecasts of bookings-to-come. The optimizer may use leg-based forecasts or path-based forecasts. Leg-based forecasts are forecasts by flight legs. But people traveling on the same flight may have different origin and destination. Path-based forecasts are forecasts for each possible path. They could be rolled up in leg-based forecasts by adding for each leg the forecasts of all paths that include this flight leg. In some cases path-based forecasts may need to be converted to leg-based forecasts if the optimizer is only using leg-based forecasts.

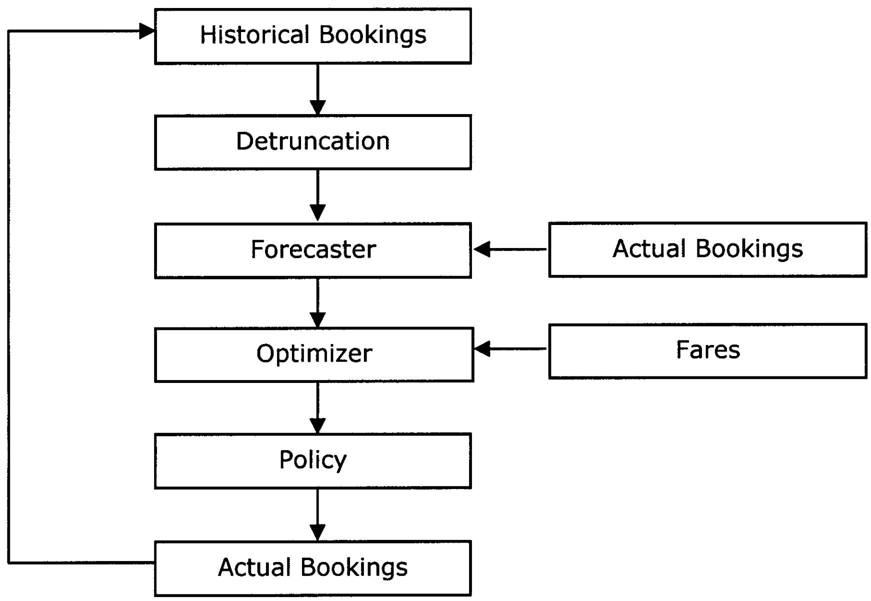

Figure 2: Traditional Revenue Management Process

The main current distinction between optimizer methods is based on whether or not they enable Origin-Destination control (for references on O-D control please refer to the Literature Review and especially to the work of **Williamson**[5,6]). But there exist two broader categories that differ according to the demand model that is used. Some just consider forecasts of bookings-to-come without taking into account the arrival pattern of passengers. Others consider at several points in time what the probability that a passenger will book a ticket in a given fare class is. As introduced by Lautenbacher and Stidham we will use the terms *static* and *dynamic* to differentiate between these two methods, respectively.

1.4.1. Static Algorithms

Using no revenue management corresponds to First Come First Served. In this case no fare class is closed by revenue management and people book seats according to restrictions and advance-purchase requirements. Revenue management adds control

[5] WILLIAMSON, *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*, PhD thesis, MIT, 1992
[6] WILLIAMSON, *Comparison of Optimization techniques for Origin-Destination Seat Inventory Control*, Master's Thesis, MIT, 1988

over the inventories by changing the availability of seats in the various fare classes according to forecasts based on historical bookings and on actual bookings occurring during the reservation process.

The static approach is the one that has been traditionally used by RM systems. It makes no assumption on the arrival pattern of passengers. Static methods use the forecasted bookings-to-come during the remaining period before departure as variables but they do not take into consideration the arrival patterns of those bookings.

At a given point in time the expected marginal value of a seat is the theoretical minimum fare that should be accepted to lead to optimal revenues. So it is the fare the airline should be ready to accept to let someone book a seat if it wants to maximize its revenues. The assumption that is used by most static approaches, to compute marginal revenues for example, is that lowest classes book first, that the passengers arrive by order of increasing willingness-to-pay.

1.4.2. Dynamic Algorithms

Dynamic algorithms for revenue management take into consideration the arrival order of passengers. In dynamic revenue management models demand for each fare class is modeled as a stochastic process. The reservation period is divided into small decision periods. During each decision period the probability to see more than one booking request is negligible, so the decision periods should become shorter when demand is supposed to be higher. Each time a request for booking arrives, the decision to either accept or refuse the booking must be made (*Figure 3*).

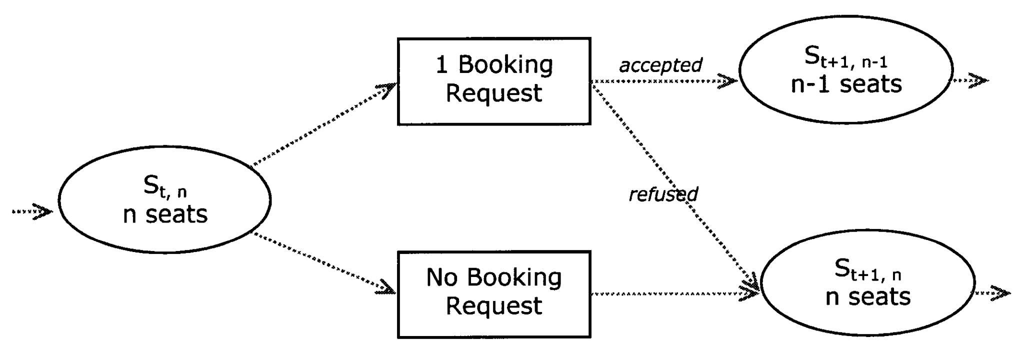

Figure 3: States and transitions

If the request is accepted then there is one less seat available and the airline earns the fare corresponding to the class in which the booking occurred. Determining the optimal policy requires using a dynamic programming algorithm. The decisions correspond to changes between states $S_{t, n}$ that are indexed by the remaining time until departure, t, and on the remaining capacity, n. All decisions are made according to forecasted probabilities to have a booking request at given points in time. The decision to close a fare class is highly dependent on the comparison of the sell-up rate between 2 fare classes and the fare ratio between those 2 fare classes. The output of a dynamic programming algorithm to be used for inventory control can include a set of classes that should be opened at a given points in time on a daily basis, the lowest class that should be opened, a bid price or booking limits. A bid price is the minimum amount an airline is ready to accept to let someone book a ticket at a given point in time. Dynamic programming algorithms demand a high level of computation as the effect of individual decisions has to be evaluated.

The dynamic programming algorithms should not be mistaken with dynamic pricing algorithms. Those algorithms consider that the decision variable should be the fare, not the set of open classes. Some of them for example consider that the expected marginal revenue used by some revenue management methods should be the direct fare requested to let someone book a seat.

1.5. Goal and Structure of the Thesis

Traditionally, the airlines have been using static revenue management methods and most work done in the PODS consortium has been focused on these methods. In this thesis we will study the performance of two methods based on dynamic programming. We will see if they can lead to greater revenues than traditional methods, what the limits of those new methods are and what can be done to improve their performance. The remaining chapters of this thesis will be a Literature Review, a Presentation of PODS (Passenger Origin-Destination Simulation), the Results, and a Conclusion.

In the literature review of Chapter 2, we present some of the most important papers that are related to traditional airline revenue management and some recent work done in the field. Then, we introduce the papers related to the dynamic approach to revenue management and describe the two methods that have been implemented in the Passenger Origin Destination Simulation. In Chapter 3 we present the PODS simulator and explain its conception as well as define the various specific settings that are important in understanding the results of our research. We first briefly describe the general structure of the simulator before focusing on specific details relevant to our research. In Chapter 4, we analyze the performance of the two dynamic programming methods that are studied in this thesis by looking at the outputs obtained by running various tests with the simulator. We will look at revenues, load factors, fare class mixes and closure of fare classes. In Chapter 5, we will present our conclusions and propose avenues for further research.

2. LITERATURE REVIEW

This chapter describes the various papers that have been written about traditional Revenue Management and the application of Dynamic Programming to Revenue Management.

2.1. Traditional Revenue Management Methods

As we previously stated, the goal of revenue management is to maximize revenues with an efficient management of inventories through advanced purchase requirements, restrictions and closure of fare classes. Revenue management really started to develop after airline industry deregulation in the U.S. in 1978.

At first some airlines started to use overbooking models to decrease the effect of no-shows as revenue loss. Overbooking means selling more seats than are actually available on given legs. Overbooking can prevent having too many empty seats on a flight when people are not showing up, but overbooking more seats than necessary can lead to having to deny boardings. For further reference, refer to the work of **Rothstein**[7,8]. In 1972 **Littlewood**[9] published a solution to the problem of a single leg with two fares. In 1987 and 1989 **Belobaba**[10,11] described a basic algorithm for multiple fare classes with nested control and named it Expected Marginal Seat Revenue heuristic. This algorithm is used by many airlines around the world.

However, these methods were developed for a single leg and consequently they do not take into consideration network effects. The first revenue management (RM) methods focused on local requests by forecasting demand and managing inventories on single legs. These RM methods were used on single leg but they could still be used with forecasts of bookings by path. We would just need to sum on each leg the forecasts for all paths that include this leg. Yet the RM methods could be improved as well.

Revenue management methods based on Origin-Destination (O-D) forecasts and/or Origin-Destination inventories were developed later. O-D forecasts correspond to forecasts of demand for a given O-D itinerary which means the number of passengers wanting to travel from a given origin to a given destination on a particular routing. RM methods that are O-D based can make efficient use of O-D forecasts. The first idea was to nest all the local and connecting fares on a given leg in various buckets. On each leg all local and connecting fare classes are nested in a given number of buckets that group fares that are close in value and the EMSRb algorithm is used to determine when a bucket should be closed. When a bucket is closed on a leg then all the fare classes it contains are closed at the same time. For

---

[7] ROTHSTEIN, *An Airline Overbooking Model*, Transportation Science, Vol. 5, 1971
[8] ROTHSTEIN, *Airline Overbooking: Fresh Approaches Needed*, Transportation Science, Vol. 9, 1975
[9] LITTLEWOOD, *Forecasting and Control of Passenger Bookings*, AGIFORS, 1972
[10] BELOBABA, *Air Travel Demand and Airline Seat Inventory Management*, PhD Thesis, MIT, 1987
[11] BELOBABA, *Application of a Probabilistic Decision Model To Airline Seat Inventory Control*, Operations Research, Vol. 37, 1989

example, **Williamson**[12,13] and **Vinod**[14] developed such algorithms. Taking into account network effects is especially important in hub-and-spoke patterns. Before the deregulation of the airline industry most airlines used point-to-point nonstop service. But after deregulation many airlines shifted to a hub-and-spoke network that enabled them to serve more city pairs and that presented some advantages such as the possibility to base most facilities (for maintenance, crews...) at the hub. In those hub-and-spoke networks people traveling on a same flight may in fact have many different origins or destinations.

The first optimization algorithms were greedy approaches that always gave preference to higher paying passengers, meaning that a connecting passenger may displace two locals while bringing less overall revenue to the airline. In order to take into account the network value of an O-D fare rather than just the price of the O-D ticket, the Displacement-Adjusted-Virtual-Nesting (DAVN) algorithm addressed this problem by removing from the fares the expected displacement cost that is induced on other legs when a request for a connecting passenger is accepted on a given leg. It obtains these expected displacement costs by solving a linear programming model. Then leg by leg the fares are put in virtual buckets.

In 1982 **Glover et al**[15] gave a mathematical network formulation as a minimum network cost flow problem. The aim of their model is to maximize revenues with flows on paths that are not violating the capacity constraints and that are less than forecasted demands. One limitation of the model is its assumption of deterministic demand. In 1986 **Wollmer**[16] extended this model by using probabilistic demand. In 1990 **Curry**[17] proposed a mathematical programming formulation for multiple flight legs with continuous demand distribution extending the work by Glover et al. However, the capacities could not be shared among different O-D paths in this algorithm. In 1992 **Wollmer**[18] introduced an optimal policy for more than 2 classes on a single leg when the lowest class books first with discrete demand distribution. **Brumelle and Mc Gill**[19] provided booking policies for a network with multiple nested fare classes. In fact, all three papers required that the classes book in increasing order of fare. In 1995 **Robinson**[20] presented a model that made no assumption on the order of arrival of the fare classes.

Other revenue management methods have also been developed using, for example, bid-prices. As we introduced earlier a bid price is the minimum amount an airline accepts to let someone book a seat at a given point in time. They are set for each leg and a booking request is accepted only if the fare is greater than the sum of the bid

---

[12] WILLIAMSON, *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*, PhD thesis, MIT, 1992

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

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

[15] GLOVER ET AL, The passenger mix problem in the scheduled airlines, Interfaces, Vol. 12, 73-79, 1982

[16] WOLLMER, *A Hub-Spoke Seat Management Model*, unpublished company report, Douglas Aircraft Company, McDonnell Douglas Corporation, Long Beach, CA, 1986

[17] CURRY, *Optimal Airline Seat Allocation with Fare Classes Nested by Origins and Destinations*, Transportation Science, Vol. 24, 193-204, 1990

[18] WOLLMER, *An airline seat management model for a single leg route when lower fare classes book first*, Operations Research, v.40 , 26-37, 1992

[19] BRUMELLE and McGILL, *Airline Seat Allocation with Multiple Nested Fare Classes, Operations Research*, Vol. 41, 127-137, 1993

[20] ROBINSON, *Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes*, Operations Research, Vol. 43, 252-263, 1995

prices along the path. Most traditional methods of revenue management use bookings limits by fare class in order to save seats for higher fare classes that request one booking limit for each fare class. These booking limits need to be computed by fare class and they may be set on legs or on paths. Bid price approaches appear to provide methods that are more intuitive and simpler to use.

Unfortunately, an effective control relying on bid-prices requires frequent updates. There is no booking limit preventing planes of filling-up with passengers who maybe do not bring the highest revenue contribution to the network. If a low fare is the current bid-price the plane may fill-up with low-fare-paying people and on the next update many seats will be booked that could have been saved for higher-yield passengers. Some work on bid prices can be found in **Talluri and Van Ryzin**[21]. In 2000 **Van Ryzin and Mc Gill**[22] presented a simple adaptive method that does not require independent forecast nor optimization. It uses the historical bookings to directly adjust booking limits.

Widely used assumptions among traditional revenue management method are the independence of demand among the various fare classes and segmentation of demand between leisure and business passengers. These assumptions are not true anymore due to the market changes imposed by the appearance of less restricted fare structures used by low-cost airlines. The use of traditional revenue management methods leads to spiral-down.

In simplified fare structures more passengers are likely to buy-down. If two available fares have similar restrictions there is indeed no reason for people to pay a higher fare. Then forecasts of higher fare classes are decreased as people buy down to lower fare classes if those classes are not closed by the revenue management. This means there are fewer seats saved in higher fare classes, which enables even more people to buy-down. This phenomenon is known as spiral-down.

Another assumption is that pricing and revenue management should be done sequentially. Pricing decisions are used as input for revenue management while the pricing decisions will have an effect on the demand that will be forecasted by revenue management. So it may be interesting to incorporate all decisions in the same process. Dynamic pricing is in part exploring this possibility. However, we will not study dynamic pricing in this thesis.

The effect of possible substitution, meaning other airlines with their own network, capacities and revenue management policies, is often ignored in revenue management studies. Some work on the effect of available substitution can be found in **Mahajan and Van Ryzin**[23] and in **Smith and Agrawal**[24]. **Netessine and Shumsky**[25] and **Zhang and Cooper**[26] more specifically study algorithms for parallel flights when demand depends on the booking policy of another airline.

---

[21] TALLURI and VAN RYZIN, *An Analysis of Bid-Price Controls for Network Revenue Management*, Management Science, Vol. 44, 1577-1593, 1998

[22] VAN RYZIN and McGILL, *Revenue Management Without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels,* Management Science, Vol. 46, 760-775, 2000

[23] MAHAJAN and VAN RYZIN, *Stocking Retail Assortments under dynamic consumer choice*, Operations Research, Vol. 49, 334-351, 2001

[24] SMITH and AGRAWAL, *Management of Multi-item Retail Inventory Systems with Demand Substitution*, Operations Research, Vol. 48, 50-64, 2000

[25] NETESSINE and SHUMSKY, *Revenue Management Games: Horizontal and Vertical Competition*, Management Science, Vol. 51, 2005

Some overviews of all the revenue management history have been done by **Barnhart, Belobaba and Odoni**[1], as well as **McGill and Van Ryzin**[27].

2.2. RM Methods Based on Dynamic Programming

Most traditional revenue management methods assume a sequential arrival of the various fare classes. Yet if this is not true those methods may lead to non-optimal solutions. Methods based on dynamic programming consider the actual demand arrival pattern. Revenue management models based on dynamic programming consider whether or not to accept a request that arrives at a particular time in the reservation process which can translate into whether or not a fare class should be opened at a given time.

Lee and Hersch[28] described a discrete-time dynamic programming model which does not require any assumption on the arrival patterns for passengers of the different fare classes and which accepts multiple seat bookings. Demand for each fare class is modeled through a Poisson process. The whole booking process is modeled as a Markov Decision Process which means that at a given point in time the state of the system is only dependent on the remaining time until departure and the remaining capacity. **Subramanian, Stidham and Lautenbacher**[29] extended the model developed by **Lee and Hersch** to incorporate overbooking, cancellation and no-shows. **Liang**[30] and **Van Slyek and Young**[31] provided a solution to the model of **Lee and Hersch** in continuous time.

One negative aspect of dynamic programming algorithms is the high computation time required to solve them. Consequently some research has been done to approximate those algorithms. In 1998 **Chen et al**[32] have shown how existing linear programming models and regression splines could be combined to estimate the value function of a Markov Decision Process formulation of the revenue management problem. **De Boer et al**[33] proposed a model based on stochastic programming models extending models developed by Williamson. In 2003, **Bertsimas and Popescu**[34] proposed an algorithm based on approximate dynamic programming that uses bid-prices obtained from a linear programming relaxation. More recent work

---

[26] ZHANG and COOPER, *Revenue Management for Parallel Flights with Customer-Choice Behavior*, Operations Research, Vol. 53, 415-431, 2005

[27] McGILL and VAN RYZIN, *Revenue Management: Research Overview and Prospects*, Transportation Science, Vol. 33, 1999

[28] LEE and HERSH, *A Model for Dynamic Airline 1212Seat Inventory Control with Multiple Seat Bookings*, Transportation Science, Vol. 27, 252-265, 1993

[29] SUBRAMANIAN, STIDHAM and LAUTENBACHER, *Airline Yield Management with Overbooking, Cancellations, and No-Shows*, Transportation Science, Vol. 33, 147-167, 1999

[30] LIANG, Solution to the Continuous Time Dynamic Yield Management Model, Transportation Science, Vol. 33, 117-123, 1999

[31] VAN SLYKE and YOUNG, Finite Horizon Stochastic Knapsacks with Applications to Yield Management, Operations Research, Vol. 48, 155-172, 2000

[32] CHEN et al, *A Markov Decision Process Based Approach to the Airline YM Problem*, working paper, http://citeseer.ist.psu.edu/chen98markov.html, 1998

[33] DE BOER et al, *Mathematical Programming for Network revenue Management Revisited*, European Journal of Operational Research, Vol. 137, 2002

[34] BERTSIMAS and POPESCU, *Revenue Management in a Dynamic Network Environment*, Transportation Science, Vol. 37, 257-277, 2003

includes research by **Bertsimas and de Boer**[35] who combined a stochastic gradient algorithm and approximate dynamic programming to improve existing booking limits obtained with a nested booking-limit policy.

In this thesis we will focus on two main RM methods that utilize dynamic programming. The first model has been developed by **Lautenbacher** and the second approach will be the one introduced by **Gallego/Van Ryzin**.

2.3. Standard DP (Lautenbacher)

Lautenbacher/Stidham[36] published a discrete-time finite horizon Markov decision process (MDP) formulation to solve the problem of a single-leg inventory. To solve their formulation they use dynamic programming with backward computation. The reservation period is divided in N decision periods such that no more than one request may arrive during one period. This means that the probability to see more than one request arrive during the same decision period is negligible. The capacity on the leg is C.

At a given point in time a request may arrive for any of the K fare classes (*Figure 4*). Demands for each fare class are supposed to be generated through independent Markov processes. The fares are as follow: $p_1>p_2>...>p_K$. The remaining capacity, the remaining time in the reservation period and the requested fare class are all taken into consideration when deciding whether to accept or reject a request. The maximum expected revenue is noted $R_n(b)$ with n, n-1, ..., 2, 1, 0, the remaining decision periods before departure, 0 being the departure time, and b the number of reservations that have been accepted, the current realized bookings.

$P_{f,n}$ = Probability to have one booking request for class f during decision period n
$P_{0,n}$ = Probability to have no booking request for any class during decision period n

From this we can find the following equation for the maximum revenue to obtain in time-frame n when b bookings already occurred:

$$R_n(b) = \sum_{f=1}^{K} P_{f,n} \, . \, \max\{ R_{n-1}(b+1) + p_f , R_{n-1}(b)\} + P_{0,n} \, R_{n-1}(b).$$

This equation summarizes the fact that for each fare class f there is a probability P_{fn} that a request may occur in this fare class and we may then accept or reject this request. If the request is accepted then we gain p_f and the maximum expected revenue for the next time frame will be $R_{n-1}(b+1)$. If the request is rejected then the maximum expected revenue for the next time frame will be $R_{n-1}(b)$.

[35] BERTSIMAS and DE BOER, *Simulation-Based Booking Limits for Airline Revenue Management*, Operations Research, Vol. 53, 90-106, 2005
[36] LAUTENBACHER/STIDHAM, *The Underlying Markov Decision Process in the Single-Leg Airline Yield-Management Problem*, Transportation Science, Vol. 33, 136- 146, 1999

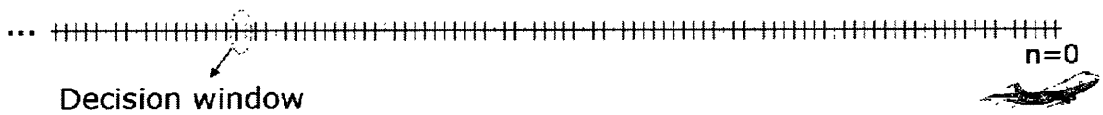

Fare Class		Probability	Expected revenue if accepted	Expected revenue if refused
1		$P_{1,n}$	$R_{n-1}(b+1) + p_1$	$R_{n-1}(b)$
2		$P_{2,n}$	$R_{n-1}(b+1) + p_2$	$R_{n-1}(b)$
	⋮	$P_{f,n}$	$R_{n-1}(b+1) + p_f$	$R_{n-1}(b)$
K		$P_{K,n}$	$R_{n-1}(b+1) + p_K$	$R_{n-1}(b)$
Nobody shows up		$P_{0,n}$	$R_{n-1}(b)$	

Figure 4: The Lautenbacher DP

The boundary conditions are $R_0(b)= 0$ if $x\leq C$ and $R_0(b)= -R(x-C)$ if $x>C$ (with R bigger than any p_f). A request for class f coming in decision period n will be is to be accepted if $R_{n-1}(b+1) + p_f > R_{n-1}(b)$. This can be rewritten as $p_f > R_{n-1}(b) - R_{n-1}(b+1)$. This quantity, which is a bid-price, is the minimum amount to pay to get on the plane. This means that a request in a given fare class will be accepted if the fare is greater than the difference in expected revenue due to one less seat available in the next time-frame. $\Delta_{n-1}(b) = R_{n-1}(b) - R_{n-1}(b+1)$ is the expected marginal revenue corresponding to a seat in decision period n-1 when there is b seats that have been booked already. We can obtain the optimal policy by backward induction. Note that the formulation is equivalent to defining the booking limits by fare class such that $B_{fn} = \min\{b\geq 0 \,/\, \Delta_{n-1}(b) > p_f\}$. A request is to be accepted only if $0 \leq b \leq B_{fn}$.

The Lautenbacher/Stidham paper also shows how static models and dynamic models can be put in relation.

2.4. The Gallego-Van Ryzin (GVR) Model

Gallego/Van Ryzin[37] modeled the problem of airline inventories with the objective to determine how to dynamically price inventories for a stochastic and price-sensitive demand.

The stochastic number of accepted requests is modeled as a Poisson process. It is expressed as a rate, $\lambda(p,s)$, the number of requests by unit of time, which is function of the current price p and the elapsed time s from the first day of the reservation

[37] GALLEGO and VAN RYZIN, *Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons*, Management Science, Vol. 40, 999-1020, 1994

period. The remaining capacity of the airline on a given leg is n. The total time of the reservation period is T. This time is divided in intervals T_i such that the probability $o(T_i)$ that more than one request arrived during T_i is negligible. The probability to sell one seat during T_i is $\lambda(p,s) * T_i$.

In a first paper Gallego and Van Ryzin explain how to find an optimal pricing policy for this problem. N_t is the total number of seats sold up to time t. If $dN_t=1$ then the airline has sold a seat and it receives p_t. The revenue rate is $p_t * \lambda(p_t,s)$.

The expected revenue is is $J_u(n,T)=E_u[\int_0^T p_t dN_t]$ with u a given pricing policy.

We define $J^*(n,t)=\sup_u(J_u(n,t))$. We can show that $\partial J^*(n,t)/\partial t = \sup_\lambda [r(\lambda) - \lambda*(J^*(n,t) - J^*(n-1,t))]$ and that this equation has a unique solution when the demand function satisfies given requirements. It may be difficult to find an optimal solution to this equation. Gallego and Van Ryzin adapted this model to a stochastic demand with a given set of prices.

In a second paper **Gallego/Van Ryzin**[38] explain how to specifically apply this model to revenue management in an airline and test it.

Gallego/Talluri[39] provide the following Bellman equation for $J_t(x)$:

$$J_t(x) = \max_S\{ \sum_{f\in S} [\lambda.P_f(S).(p_f + J_{t-1}(x-1))] + (\lambda.P_0(S) + 1 - \lambda).J_{t-1}(x) \}$$

Using the same notations as for the Lautenbacher formulation we introduce:
- λ: probability of a request arrival
- S: a subset of fares chosen among {1,2,..,K}
- $P_f(S)$: probability that a customer choose product f when S is the subset of available fares
- $P_0(S)$: probability that a customer that makes a request do not choose any of the available fares

Note that for all S

$$- \sum_{f\in S} P_f(S) + P_0(S) = 1$$

This formulation takes into consideration the possibility of buy-down or sell-up while the Lautenbacher formulation assumes independent fare classes. The Gallego-Van Ryzin formulation first considers the arrival of random passengers through λ. These passengers are supposed to possibly sell-up to any open class or to spill out. The probabilities of sell-up depend on the set of open fare classes. With the Lautenbacher model we only consider whether or not to accept a request occurring in a given fare class. These request arrivals are supposed to be independent of the set of open fare classes so this model assumes independent demand for the different fare classes.

Isler[40] adapted this formulation to the case that considers as a decision variable the lowest class to be open (*Figure 5*). When there is no restriction between classes,

[38] GALLEGO and VAN RYZIN, *A Multiproduct Dynamic Pricing Problem and its Application to Network Yield Management*, Operations Research, Vol. 45, 24-41, 1997
[39] GALLEGO and TALLURI, *Revenue Management Under a General Discrete Choice Model of Consumer Behavior*, Management Science, Vol. 50, 15-33, 2004
[40] ISLER, *Swiss International Air Lines Ltd. Note*, 2004

among the various fares that are available, bookings will only happen in the lowest open class.

... n=0

Decision window ⟷ Pr( )=λ, Pr(nobody)=1- λ

Fare Class		Probability of sell-up	Expected revenue if accepts to sell-up	Expected revenue if refuses to sell-up
1		$Pr_{1,n}$	$J_{n-1}(x-1) + p_1$	$J_{n-1}(x)$
2		$Pr_{2,n}$	$J_{n-1}(x-1) + p_2$	$J_{n-1}(x)$
	⋮	$Pr_{f,n}$	$J_{n-1}(x-1) + p_f$	$J_{n-1}(x)$
K		$Pr_{K,n}$	$J_{n-1}(x-1) + p_K$	$J_{n-1}(x)$
Wants nothing		$Pr_{0,n}$	$J_{n-1}(x)$	

Remaining seats

Figure 5: The Gallego-Van Ryzin DP

One can see that this approach is close to the approach of Lautenbacher. However, the Gallego-Van Ryzin approach takes into account the possibility of sell-up.

$J_n(x) = \max_f\{ \lambda.(Pr_{f,n}.(J_{n-1}(x-1) + p_f) + (1- Pr_{f,n}).(J_{n-1}(x))) + (1 - \lambda).J_{n-1}(x) \}$
In this formulation $Pr_{f,n}$ is the probability in decision period n that a customer will buy the fare f given that he would buy the lowest fare. We introduce $\Delta J_n(x) = J_n(x) - J_n(x-1)$ which is the bid price. Then $J_n(x) = J_{n-1}(x) + \lambda. \max_f\{ Pr_{f,n}.(p_f - \Delta J_{n-1}(x)) \}$.

If we assume that the sell-up probabilities are of the form $\exp(-b(F-p_K))$ where F is the fare then we can compute the optimal price to charge:
$F_n^*(x) = \Delta J_n^*(x) + 1/b_n$

We should change the lowest open fare class from f to f+1 when:
$Pr_{f,n} * (p_f - \Delta J_n(x)) = Pr_{f+1,n} * (p_{f+1} - \Delta J_n(x))$

So class f will be open only if:
$(Pr_{f+1,n} * p_{f+1} - Pr_{f,n} * p_f) / (Pr_{f+1,n} - Pr_{f,n}) > \Delta J_n(x)$

This model is to be used in unrestricted fare structures. It assumes that people will always buy in the lowest open fare classes. In fully-restricted fare structures people book in any fare class that is open even if there is lower open fare classes. Indeed restrictions segment demand and some people agree to pay a higher price to get their desired "fare product". In unrestricted fare structures as there is no restriction between fare classes there is indeed no reason for passengers to buy a ticket in a fare class that is higher than the lowest open one. In this thesis we will refer to this approach as the Gallego-Van Ryzin approach.

3. The Passenger Origin Destination Simulation

3.1. General Presentation

PODS is the **Passenger Origin Destination Simulation**[41], a software package that was developed by Hopperstad at Boeing. It evolved from the **Decision window Model**[42], another software package developed at Boeing. The Decision window Model was essentially taking into account schedules, airline image and aircraft capacities. PODS simulates a market with 2, 3, 4 or more airlines serving various O-D markets on a set of legs, the potential demand on those markets, the revenue management policies applied by the airlines, and the booking decisions made by passengers. The simulator has been described in details by **Zickus**[43], **Gorin**[44], **Carrier**[45] **and Cleaz-Savoyen**[46]. We will only introduce in this thesis the fundamentals that are necessary to understand the results analyzed in the next chapter and the reader should refer to these other theses for further description of the simulator.

3.1.1. The Networks

There are three main environments that have been developed to run tests in PODS. To define an environment we can set the number of airlines, the cities, the legs and the origin-destination markets. We can then choose the schedules and capacities on each leg, the number of fares and the restrictions associated with each fare.

The most basic environment is a single leg market. There are 2 airlines, 2 cities and each airline operates three flights per day in only one direction between those two cities (*Figure 6*). Each airline offers 8 fare classes.

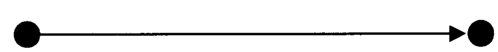

Figure 6: Single market

The intermediate environment, in which most tests have been run, is called Network D. There are 2 airlines. Each one has a hub and they both operate flights between 20

[41] Passenger Origin Destination Simulation, developed at Boeing by Hopperstad, Berge and Filipowski, 1997
[42] *Decision Window Model (DWM),* The Boeing Company, February 1994
[43] ZICKUS, *Forecasting for Airline Network revenue Management; Revenue and Competitive Impacts*, May 1998
[44] GORIN, *Airline Revenue Management: sell-up and Forecasting Algorithms*, June 2000
[45] CARRIER, *Modeling Airline Passenger Choice: Passenger Preference for Schedule in the Passenger Origin-Destination Simulator (PODS)*, May 2003
[46] CLEAZ-SAVOYEN, *Airline Revenue Management Methods for Less Restricted Fare Structures*, May 2005

west-coast spoke cities, the 2 hubs and 20 east-coast spoke cities (*Figures 7 and 8*). Since each airline operates flights in only one direction, there are 482 O-D markets:

- 440 with origin on the west coast and with destination being a hub or a east-coast city.
- 40 with origin being one of the 2 hubs and with destination on the east coast
- 2 between the hubs

On each leg there are 3 flights per day in only one direction, except on the leg between the 2 hubs where both airlines have 3 flights a day in each direction.

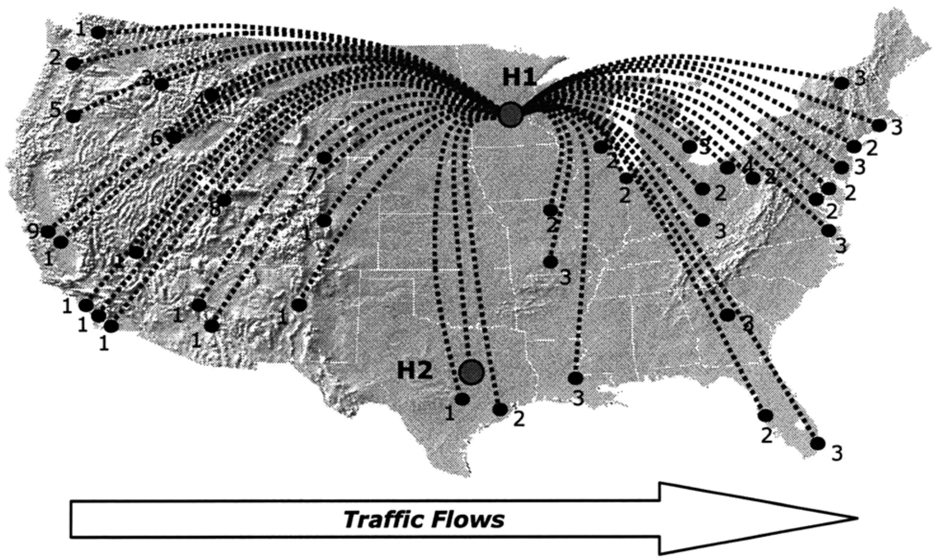

Figure 7: Network of Airline 1

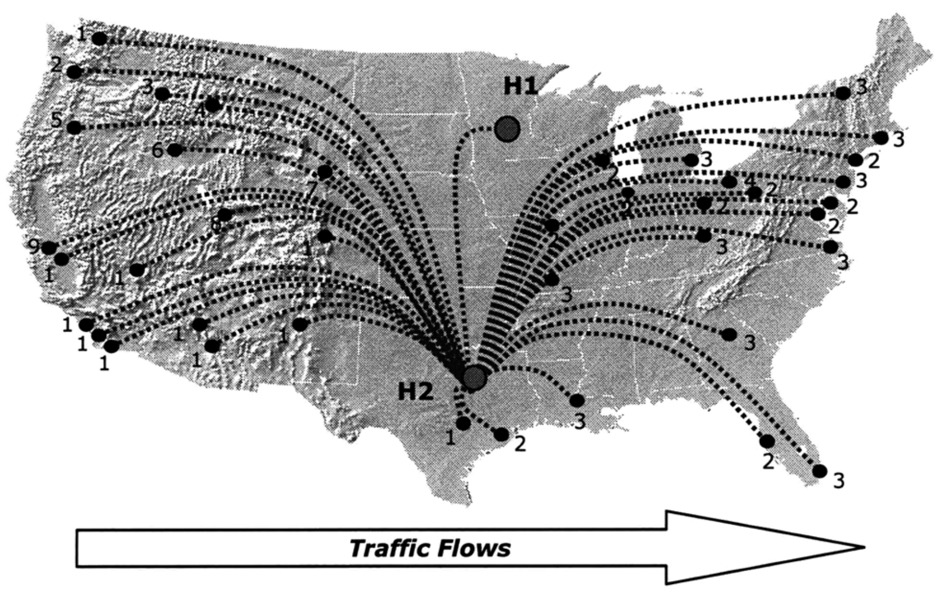

Figure 8: Network of Airline 2

The bigger network is called Network R (*Figure 9*). On this network there are 4 airlines, MSP, ORD, MCI and DFW operating 572 O-D markets.

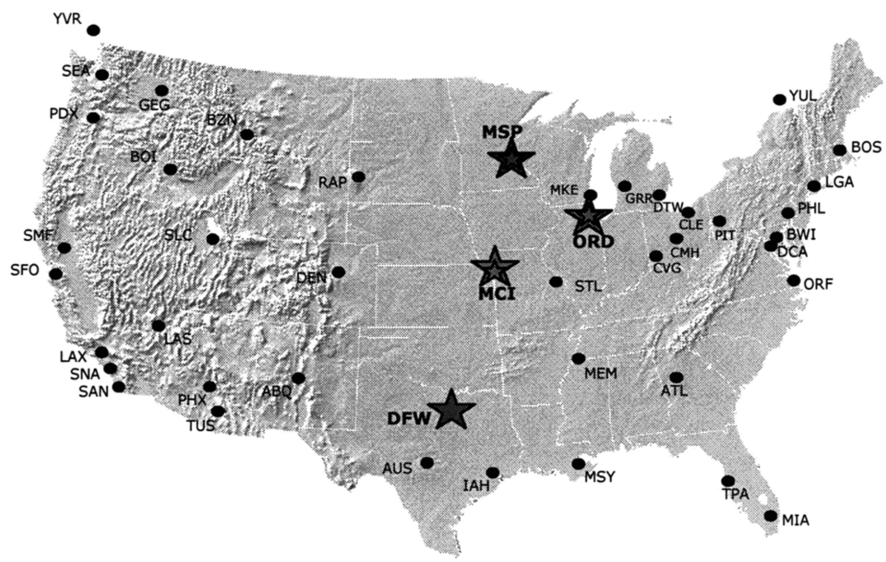

Figure 9: Network R

Each airline has different route networks and serves specific O-D markets:

- MSP has 24 origin cities (the 20 west spokes; MSP; the three other hubs), 24 destination cities (the 20 east spokes; MSP; the three other hubs) and 3 hub bypass services (DEN/LAX/SFO to BOS).
- ORD has 24 origin cities (the 20 west spokes; ORD; the three other hubs), 23 destination cities (19 of the east spokes; ORD; the three other hubs) and 6 hub bypass services (DEN-LGA, SFO-PHL, LAX-ATL/BOS/BWI/PHL).
- MCI has 15 origin cities (11 of the west spokes; MCI; the three other hubs), 20 destination cities (16 of the east spokes; MCI; the three other hubs) and 19 hub bypass services (DEN-ATL/BOS/LGA, LAS/LAX/SFO-PHL, SFO-BOS and 3 services between MSP and IAH/ATL/BOS/STL)
- DFW has 18 origin cities (14 of the west spokes; DFW; the three other hubs), 24 destination cities (the 20 east spokes; MSP; the three other hubs) and 4 hub bypass services (DEN-ATL, LAS-PHL, LAX-ATL/BWI)

3.1.2. The Demand

Demand per day is given as means and standard deviations for O-D markets at a "base fare" such that all this demand would agree to pay this base fare to travel. Passengers are later assigned a maximum willingness-to-pay and sensitivity to restrictions. There are 2 types of passengers, Business and Leisure, and each category requires its own input parameters. Arrival curves define the arrival patterns of passengers by passenger types during the booking process (*Figure 10*).

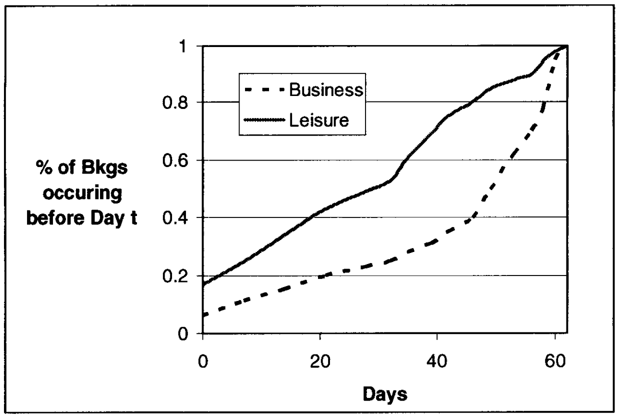

Figure 10: Arrival curves

The willingness-to-pay is determined separately in each category of passenger by the following formula:

$$P(\text{pay at least } f) = \min\left[1, e^{-\frac{\ln(0.5)*(f-\text{basefare})}{(\text{emult}-1)*\text{basefare}}}\right]$$

where the basefare is the fare at which all passengers would travel and emult is an elasticity multiplier parameter such that 50% of the passengers are ready to pay emult*basefare to travel.

The impact of restrictions is taken into consideration by adding to the O-D fares random disutility costs based on input Gaussian densities. There are 3 types of restrictions associated with smaller or stronger disutilities. The costs associated to disutilities are later added to the fare. Each passenger has its own time window in which he has planned to make his trip and there is disutility associated with replanning. Generalized costs are the sum of fare, disutility costs, connecting costs (if the path requires a connection), unfavorite airline costs (each passenger has a favorite airline) and replanning costs (if the path requires departure or arrival outside the time window of the passenger). To travel between his desired origin and destination, the passenger picks the fare that has the lowest generalized costs among all the available fares that are less than his willingness-to-pay. If none is less than his willingness-to-pay the passenger decides not to travel. The various fare structures and fares that are used in PODS are shown in *Tables 1 and 2*.

FC	AP	R1	R2	R3	FC	AP	R1	R2	R3	FC	AP	R1	R2	R3
1	0	0	0	0	**1**	0	0	0	0	**1**	0	0	0	0
2	0	0	0	0	**2**	0	0	1	0	**2**	3	0	1	0
3	0	0	0	0	**3**	7	0	1	1	**3**	7	1	0	0
4	0	0	0	0	**4**	14	0	1	1	**4**	10	1	1	0
5	0	0	0	0	**5**	14	0	1	1	**5**	14	1	1	1
6	0	0	0	0	**6**	21	0	1	1	**6**	21	1	1	1
Unrestricted fare structure					Simplified fare structure					Fully-restricted fare structure				

Table 1: Fare Structures

Class	1	2	3	4	5	6
Average Fare	413	293	179	153	127	101

Table 2: Average fares in Network D

3.1.3. The Runs

To obtain statistically stable results it has been determined that one run of the simulator should consist of 5 trials, each trial corresponding to 600 departures on each leg. Demand is randomly generated according to input parameters and tests are initiated with input values. In each trial the first 200 samples are burnt in order not to have the initiation inputs affecting the results. For each departure there is a reservation period that is divided in 16 time-frames. As shown in *Figure 11* for a given departure the traditional optimizer methods check in each time-frame the actual bookings that already occurred and the historical bookings that had occurred for previous departures and then set a new inventory control policy for the remaining booking period. The choice of passengers among available paths between given origin and destination depends on the availability, the fare, the restrictions as previously described.

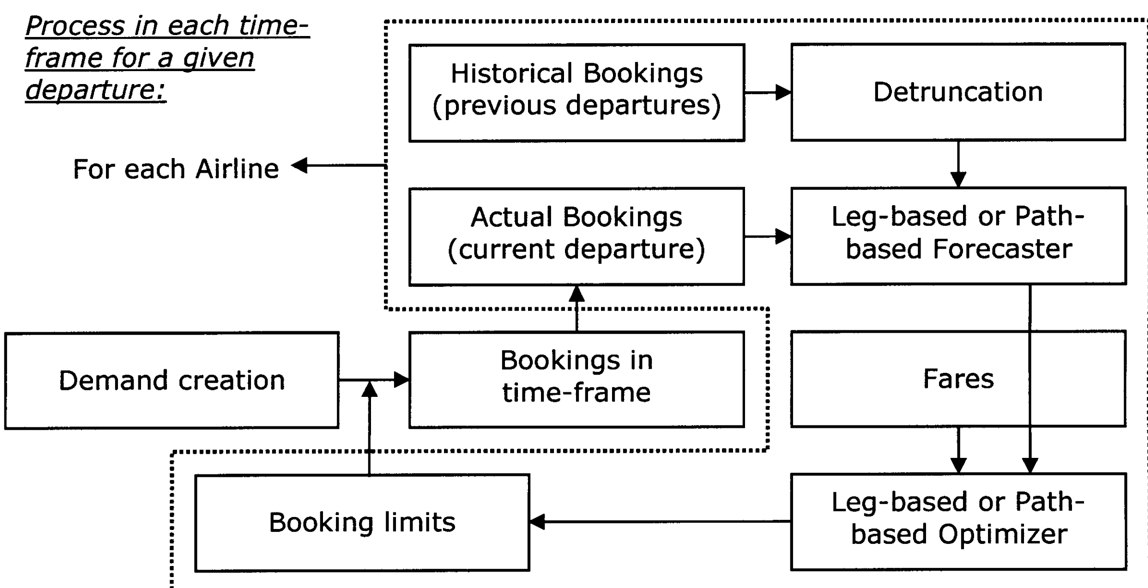

Figure 11: PODS booking process for traditional RM methods

3.2. Detruncation and Forecasting

3.2.1. Traditional Methods

When a fare class is closed down before the end of the booking process actual bookings do not correspond to the total demand for this fare class. Detruncation is then needed to obtain unconstrained demand. The method we will use is booking curve detruncation. The demand curve is obtained by extrapolation after closure of the fare class (*Figure 12*).

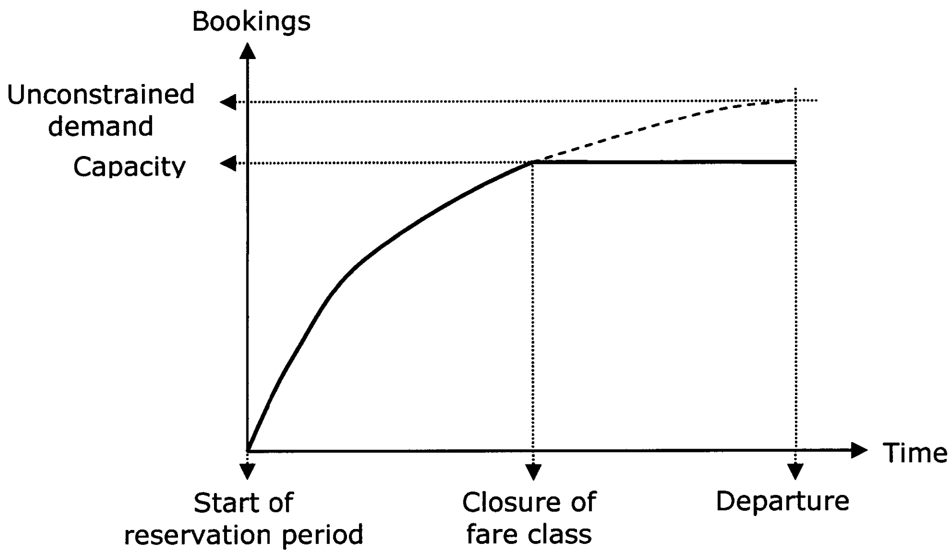

Figure 12: Booking Curve Detruncation

The forecasting method that is mainly used when running tests in fully-restricted environments is Pick-up Moving Average. This method averages the bookings-to-come between the current time-frame and departure and adds this to the bookings that already occurred to estimate the total demand for the departure. It is the method we will refer to as standard forecasting.

To be able to maximize revenue an airline must first be able to adequately forecast current demand patterns on its markets. During many years the revenue management algorithms supposed that demands for the different fare class were independent. But as fare classes become less differentiated it appears that in fact people may frequently switch from one class to another, buying a lower fare or a higher fare. A new representation of demand was necessary. The new assumption for forecast is less restricted fare structures is that people have a given willingness-to-pay, the maximum price that they are willing to accept but will buy a lower price if available.

3.2.2. Q-Forecasting

Traditional forecasting assumes independence of fare classes. Consequently it can average historical bookings by fare class to obtain bookings-to-come in each fare class. However, if we assume that passengers may sell-up or buy-down between fare classes then we can no longer use this traditional forecast. This may especially happens when there are no restrictions between classes, in unrestricted fare structures. In an unrestricted fare structure bookings do no longer depend on the set of fare classes that are open but only on the lowest open fare class as people buy the lowest available fare.

If we use traditional forecasting in unrestricted fare structures then people buying down are counted as demand for lower fare classes. Consequently the RM optimizer decreases the protection for higher fare classes. Then even more people buy down to lower fare classes. This phenomenon is known as "spiral-down" in unrestricted fare structure. Traditional forecasting can no longer be used in unrestricted fare structures. One can not just average actual historical bookings; we need to estimate potential demand for each fare class, demand that will materialize only if the class is the lowest open one.

Belobaba and Hopperstad[47] have developed a way to forecast this potential demand in unrestricted fare structures: "Q-Forecasting" has been designed to predict potential demand by taking into account the possibility of buy-down/sell-up. In each time-frame Q-Forecasting either uses input or estimated probabilities of sell-up to forecast potential demand by fare class as we will describe. In his master thesis **Cleaz-Savoyen**[48] has been testing the performance of Q-Forecasting with traditional RM methods and noticed the improvement in comparison with the use of standard forecasting. He was using input FRAT5s. For the moment different ways to predict probabilities of sell-up are being developed. Yet none of the ones that are currently implemented in PODS was stable enough to be used to get experimental results for this thesis. Consequently in this thesis we will use inputs for probabilities of sell-up.

We will now describe the basic steps of Q-Forecasting. First in each time-frame historical bookings happening in the open fare classes are converted to equivalent Q-bookings, meaning the estimated number of passengers that would have bought a ticket in the lowest fare class, fare class Q, if it had been opened. To do so we use the probabilities of sell-up from class Q to higher fare classes. The number of equivalent Q-bookings is equal to the number of bookings in a fare class divided by the probability of sell-up from fare class Q to this fare class. For each time-frame these computed numbers of equivalent Q-bookings constitute the historical data. Detruncation is applied to those numbers of equivalent Q-bookings.

---

[47] BELOBABA and HOPPERSTAD, *Algorithms for Revenue Management in Unrestricted Fare Markets*, AGIFORS, Auckland, New Zealand, Jan 2004
[48] CLEAZ-SAVOYEN, *Airline Revenue Management Methods for Less Restricted Fare Structures*, May 2005

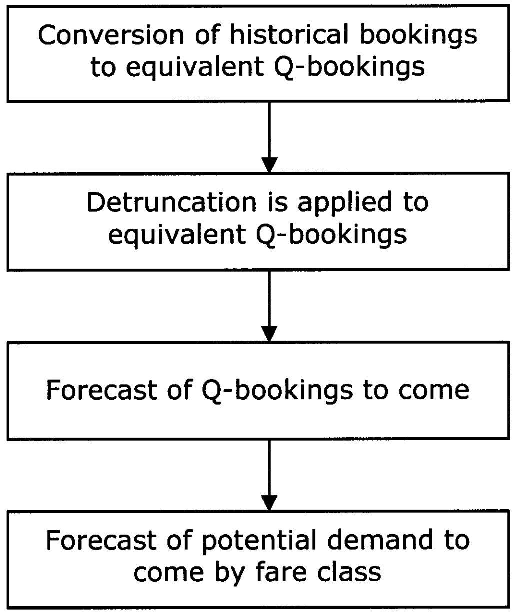

Figure 13: Q-Forecasting Process

The historical data is then used to forecast equivalent Q-bookings to come between the current time-frame and departure. According to these forecasted Q-bookings to come and to probabilities of sell-up we can forecast the potential demand by fare class.

As data is sparse one need to make assumptions on how probabilities of sell-up are related to fare ratios in order to be able to forecast the probabilities of sell-up from Q to any fare. In PODS particular assumptions concerning probabilities of sell-up have been made so that forecasting requires the estimation of a minimum number of parameters. First all the probabilities of sell-up in a given time-frame are supposed to be summarized by the fare ratio at which 50% of the passengers are willing to sell-up, which is called FRAT5. Then the shape of the probability of sell-up according to the fare ratio is supposed to be exponential as shown in *Figure 14*. The formula which is used to link probabilities of sell-up and the FRAT5s is as follows:

$$\text{Psup}_{Q\to f} = e^{\ln\left(\frac{1}{2}\right)\cdot\frac{\left(\frac{faref}{fareQ}-1\right)}{Frat5-1}}$$

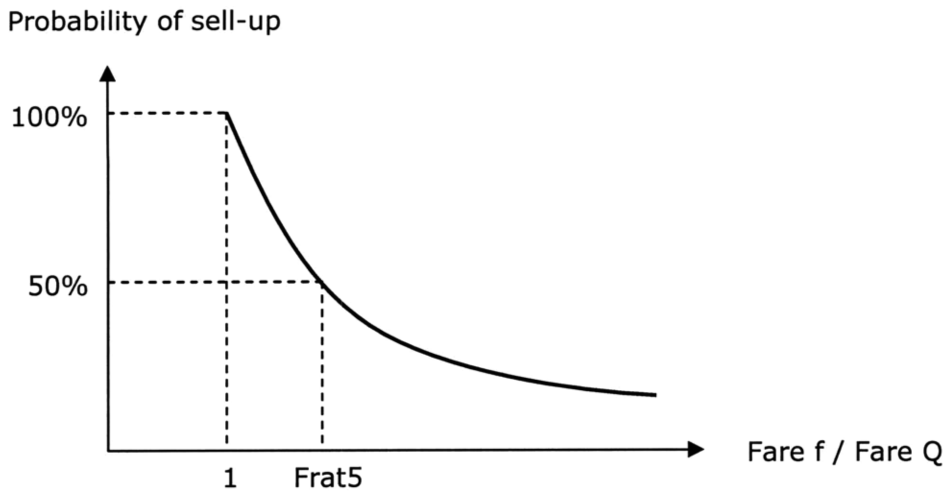

Figure 14: Shape of the probability of sell-up

According to the arrival curves of leisure and business passengers and the various parameters that define the willingness-to-pay of passengers we estimated what the shape of the FRAT5 curve across time-frames is. From this curve we computed other curves representing lower or higher willingness-to-pay. This gave us a set of FRAT5 curves that have been used as input for Q-Forecasting (*Figure 15*). Indeed rather than estimating FRAT5s we used for our simulations a set of pre-determined input FRAT5s. Each set of FRAT5s is designed by a letter or a pair of a letter and a number: E, C, A, A2 or A4 are the sets we will mainly use in this thesis. The rule to name the FRAT5s was: the "higher" the letter the higher the FRAT5s. Thus in each time-frame FRAT5s of set A4 are higher than those of set E.

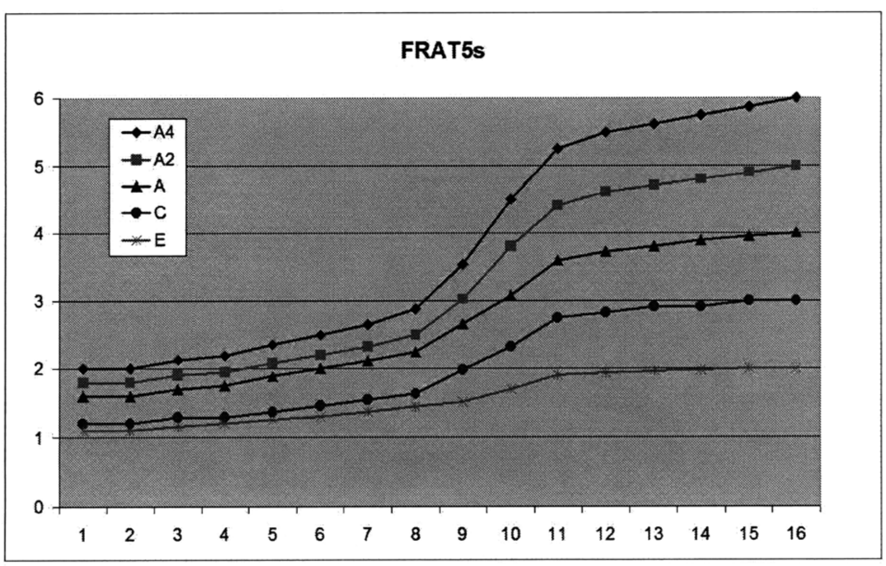

Figure 15: Input FRAT5s

3.2.3. Hybrid Forecasting

In simplified fare structures not all fare classes are differentiated. Only some of them have different restrictions while the remaining fare classes are undifferentiated. In simplified fare structures we can differentiate between two types of behaviors among passengers as introduced by **BOYD and KALLESEN**[49]. If two fare classes do not have the same restrictions some passengers may prefer to pay the higher price because of inconvenience associated with restrictions of the lower fare class. Such demand is called product-oriented. Product-oriented people may not buy the lowest prices because of restrictions associated with them. On the contrary price-oriented people always buy the lowest available price. In fully-restricted fare structures, as all classes are differentiated all demand was supposed to be product-oriented. As demand was thought to be segmented we could use traditional forecasting methods, which suppose independent demand by fare class. In unrestricted fare structures, as there is no difference of restrictions between fare classes, people always book tickets at the lowest available price. All demand is then price-oriented. One then needs to forecast potential demand as traditional forecasting has been shown to lead to spiral-down. Q-Forecasting has been developed to be used in these unrestricted environments.

"Hybrid Forecasting" developed through the PODS consortium by **Hopperstad and Belobaba**[50] has been designed to be used with simplified fare structures. Indeed in semi-restricted environments there is both product-oriented and price-oriented demand. To get a fare with fewer restrictions, some people agree to pay a higher fare than the lowest available one. However, in classes that have the same restrictions, price-oriented people always buy the available one that corresponds to the lowest fare. Hybrid Forecasting separately uses traditional forecast for product-oriented demand and Q-Forecasting for price-oriented demand. Nonetheless the classification of historical bookings as product-oriented demand or price-oriented demand requires some specifications.

We have tested three ways to classify bookings as product-oriented. Two of these methods consider the possibility to buy down to another class available on another path provided by the airline or a competitor and the last one does not. A passenger booking a ticket is classified as product-oriented:

- **HF1:** if the next lower class is available on the same path. (path rule)
- **HF2:** if the next lower class is available on any path provided by the airline that corresponds to the same origin and destination. (airline rule)
- **HF3:** if the next lower class is available on any path in the market. (market rule)

This product-oriented demand is forecasted using traditional forecasters. Then any historical booking which was not classified as product-oriented is classified as price-oriented. This price-oriented demand is forecasted using Q-Forecasting. *Figure 16* represents the scope of the 3 different rules for Hybrid Forecasting: the path rule, the airline rule and the market rule.

---

[49] BOYD and KALLESEN, *The science of revenue management when passengers purchase the lowest available fare*, Journal of Revenue and Pricing Management, Vol. 3, 2004
[50] ZERBIB, *Hybrid-Forecasting*, PODS meeting in Copenhagen, 2004

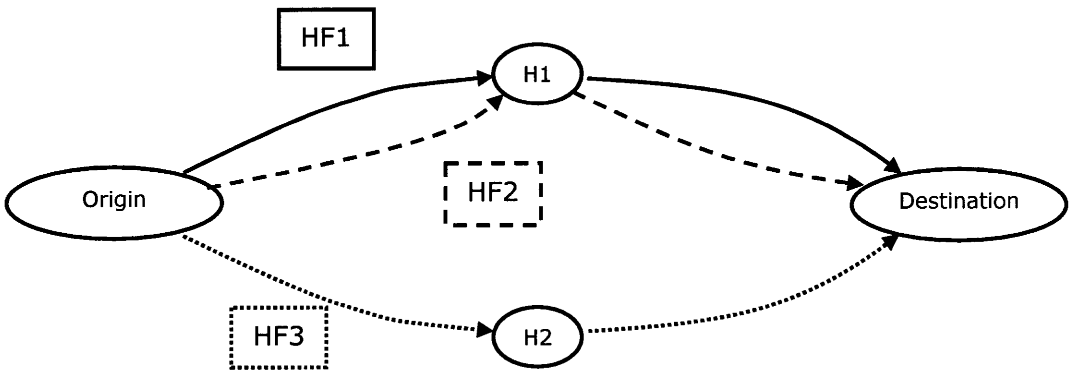

Figure 16: Scope of Hybrid-Forecasting rules

When the next lower class has been closed because of advance-purchase requirements should people buying in the lowest available class be classified as price-oriented or product-oriented? Some of them may have bought down to the next lower class if it was available but they will never do so as this next lower class will always be closed in this time-frame. So one may think those passengers should be classified as product-oriented as they are choosing this fare class because of restrictions. Yet those passengers may at least sell-up in higher classes so one could argue they should be classified as price-oriented. We know that by classifying them as price-oriented we will not account for the possibility of buy-down as the lower fare classes will always be close. We see that both views have arguments. We may consequently even decide to classify only part of those passengers as price-oriented by choosing further differentiation or by picking a percentage.

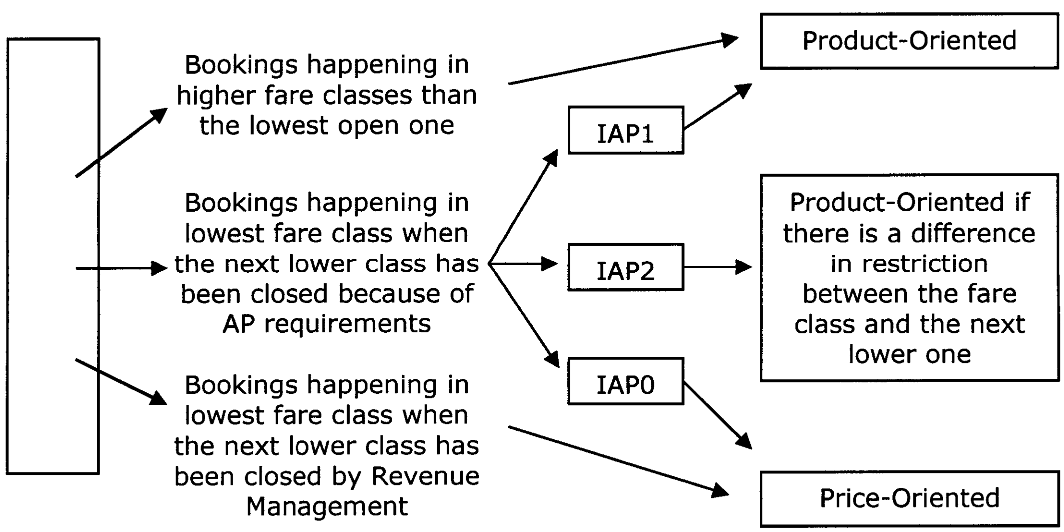

Figure 17: Classification of AP bookings with the path rule of Hybrid Forecasting

When using HF1 one can choose in PODS the way passenger bookings due to AP requirements are classified (*Figure 17*). Here are the 3 ways to classify passengers booking a ticket in a fare class when the next lower one has been closed because of AP requirements:

- **IAP0:** Bookings due to closure because of AP are counted as price-oriented.
- **IAP1:** Bookings due to closure because of AP are counted as product-oriented.
- **IAP2:** Bookings due to closure because of AP are counted as product-oriented only if there is a difference in restrictions between the class and the next lower class.

When using Hybrid Forecasting we have three different classifications of product/price-oriented passengers, three different classifications of AP bookings and several set of FRAT5s. In order not to have to run for each test all the various possible combinations we tried to determine what were the best classifications of product/price-oriented passengers and of AP bookings. First we determine the classification of product/price-oriented passengers that leads to the best revenues. Running several tests with traditional RM methods (EMSRB, DAVN, DP using Hybrid Forecasting and various FRAT5s for hybrid-forecasting, against EMSRb, EMSRb with HF or same method) has shown that in most cases HF1 is the classification of product/price-oriented passengers that leads to the best revenues. We also determined which classification of AP bookings leads to the best results. We run several tests with EMSRB, DAVN, DP-LB using Hybrid Forecasting and various FRAT5s, against EMSRb, EMSRB with HF or same method. In all cases IAP1 leads to weaker results. Consequently we mainly used HF1 with either IAP0 or IAP2 in our tests with hybrid forecasting.

3.3. Traditional RM Methods in PODS

In traditional revenue management the forecasts are used by the optimizer to set booking limits by fare class. We will sometimes use no revenue management to determine what the results would be if the airline let passengers book tickets in their order of arrival in the booking process. When an airline uses no revenue management for a given departure, everybody can book a seat until the plane fills up. Using no revenue management is called using First Come First Served (FCFS). We will now describe the traditional RM methods that we will use in PODS: fixed or adaptive threshold, EMSRb, and DAVN.

3.3.1. Threshold RM Methods

Thresholds methods do not use forecast. They set a threshold percentage in each fare class so that the total number of bookings in this fare class and the lower fare classes do not exceed the booking threshold obtained by multiplying the threshold percentage and the capacity. Once this happens, the fare class is closed. For fixed threshold the threshold percentages are set by fare class at the beginning of the booking process. *Table 3* gives an example of threshold percentages for 6 fare classes. Adaptive threshold is a method that was developed to simulate the use of a basic revenue management method. A load factor target is set at the beginning of the booking process and then at each time-frame the threshold values are computed according to actual bookings. The final load obtained by the airline is equal to the load factor target.

Class	1	2	3	4	5	6
Threshold	1.00	0.80	0.70	0.60	0.50	0.40

Table 3: Threshold percentages

3.3.2. EMSRb

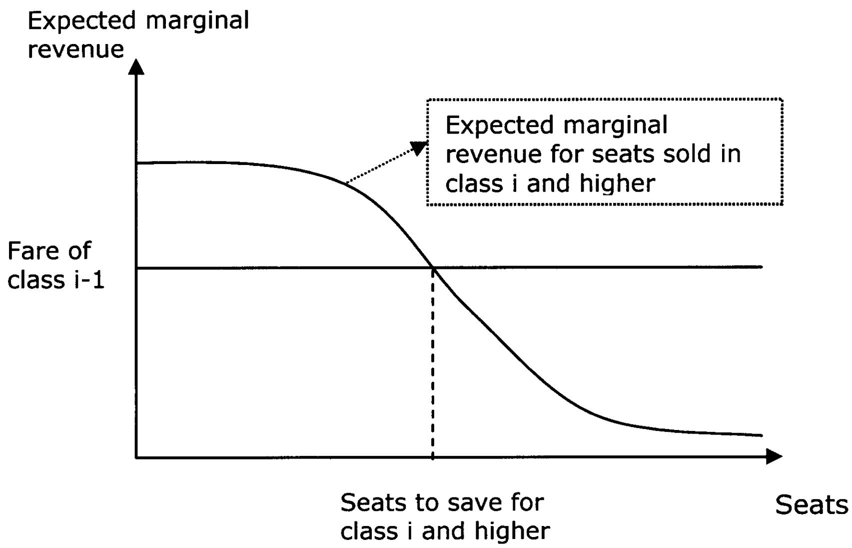

Figure 18: Determination of Booking Limits with EMSRb

EMSRb is the Expected Marginal Seat Revenue method developed by **Belobaba**[51,52]. It uses nested booking inventory control. Demand is supposed to be independent for each fare class, meaning that people are supposed to always buy in the same fare class. Demand-to-come for each fare class is modeled through a Gaussian distribution. The mean and the standard deviation are determined for each class based on detruncated historical data. Then we compute the mean, the standard deviation and the average fare for joined demand, taking into consideration for each fare class total demand in higher fare classes. These statistical parameters are then used to compute the expected marginal revenue. *Figure 18* shows how to determine the number of seats to be saved for fare classes which are higher than a given fare class. We repeat the process for each fare class. The booking limit for all bookings occurring in fare class i and lower fare classes is set equal to the capacity minus the number of seats to be saved for higher classes.

---

[51] BELOBABA, *Air Travel Demand and Airline Seat Inventory Management*, PhD Thesis, MIT, 1987
[52] BELOBABA, *Application of a Probabilistic Decision Model To Airline Seat Inventory Control*, Operations Research, Vol. 37, 1989

3.3.3. DAVN

DAVN is the Displacement Adjusted Virtual Nesting method described by **Williamson**[53]. DAVN requires O-D forecasts as this method enables O-D control. A connecting passenger may displace two local passengers while the total of the 2 local fares are higher than the fare paid by the connecting passenger. DAVN addresses this problem by decreasing connecting fare on each leg from the amounts of the displacement costs associated with the other flights legs of the itinerary. To determine the displacement costs we solve a deterministic linear program and look for the dual solution as introduced by Williamson.

$$Max \left(\sum_i \sum_j p_i^j x_i^j \right)$$

Subject to:

$$x_i^j \; \langle \; f_i^j \qquad \forall i, j$$

$$\sum_i \sum_j x_i^j \delta_i^k \langle C_k \qquad \forall k$$

With:

p_i^j *The fare of fare class j for path i*

x_i^j *The number of passengers in fare class j on path i*

f_i^j *The forecast for fare class j on path i*

C_k *The capacity on leg k*

δ_i^k *= 1 if leg k is part of path i*

= 0 otherwise

For each leg k we can then get the dual solution that corresponds to the marginal revenue of adding an extra seat on the leg. These values are the displacement costs that are subtracted from the connecting fares. On each leg all fares are ordered in virtual buckets according to their new value. The EMSRb optimizer is then applied to virtual buckets on each leg to determine the booking limits by bucket. A passenger request for a connecting booking is accepted only if on any leg of the itinerary the corresponding fare is in a virtual bucket that is still available.

3.4. DP RM Methods in PODS

We now describe the implementation in PODS of the two RM methods based on dynamic programming we study in this thesis: the RM method based on the Lautenbacher approach and the one based on the Gallego – Van Ryzin as described in Chapter 2.

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

3.4.1. DP-Lautenbacher

As introduced in the literature review the formula that is used to compute the expected maximum revenue in decision period n when b bookings already occurred is:

$$R_n(b) = \sum_{f=1}^{K} P_{f,n} \, . \, \max\{ R_{n-1}(b+1) + p_f \, , R_{n-1}(b)\} + P_{0,n} \, R_{n-1}(b)$$

The probabilities $P_{f,n}$ are computed with time-frame leg/class forecasts.

Then $P_{0,n} = 1 - \sum_{f=1}^{K} P_{f,n}$

In fully-restricted environments the traditional forecasting method Pick-up Moving Average is used to produce path/class forecasts of bookings-to-come. To obtain time-frame forecasts, path/class forecasts of bookings-to-come are split out by time-frame using path/class booking curves. Then the obtained time-frame path/class forecasts are rolled up into time-frame leg/class forecasts. The optimizer can then produce the optimal policy to follow during all decision periods remaining before departure. In unrestricted environments the same method is used with forecasts of bookings-to-come that are obtained by partitioning forecasts of equivalent Q-bookings-to-come.

PODS does not process the policy into booking limits; the policy is directly used by the reservation system to control the closure of fare classes.

3.4.2. DP-GVR

The formula used for DP-GVR that we introduced in the literature review of Chapter 2 is as follows:

$$J_n(x) = \max_f\{ \lambda.(Pr_{f,n}.(J_{n-1}(x-1) + p_f) + (1- Pr_{f,n}).(J_{n-1}(x))) + (1 - \lambda).J_{n-1}(x) \}$$

$J_n(x)$ is the maximum expected revenue in time-frame n when there are x seats remaining. DP-GVR does no partition Q-forecasts. It directly uses the leg forecasts of equivalent Q-bookings-to-come. These forecasts are split by time-frame according to leg booking curves. Those forecasts are then used to compute λ for each time-frame. The probabilities of sell-up $Pr_{f,n}$ are obtained for each time frame by using the time-frame FRAT5s:

$$Pr_{f,n} = e^{\ln\left(\frac{1}{2}\right) \cdot \frac{\left(\frac{faref}{fareQ} - 1\right)}{Frat5_n - 1}}$$

The FRAT5s are used twice. They affect both the estimated demand for the lowest fare class and the sell-up probabilities that are directly used by the optimizer. The optimizer determines what the lowest open fare class should be, leading to a daily policy. The daily policy is obtained by averaging the bid prices for the slices associated with each day. The bid prices in time-frame n are $\Delta J_n(x) = J_n(x) - J_n(x-1)$.

While the standard Lautenbacher DP approach (DP-LB) determines which classes should be open for a given time frame, the Gallego-Van Ryzin Model (DP-GVR) determines the lowest open class and it can for this reason be compared to pricing models.

3.4.3. Example based on DP-GVR

To illustrate how the methods based on dynamic programming work we will now introduce a small example based on DP-GVR. We consider a single departure. There are only 3 decision periods, the plane has only 2 seats and there are 3 available fare classes. The fares for those 3 fare classes are \$1600, \$500 and \$100. The arrival rate is 0.5 in each decision period and *Table 4* gives the probabilities of sell-up in each decision period, meaning the probabilities that if someone shows up he will be willing to sell-up to the given lowest open one.

	3	2	1
$Pr_{1,n}$	5%	5%	10%
$Pr_{2,n}$	10%	22%	30%
$Pr_{3,n}$	100%	100%	100%

Table 4: Probabilities of sell-up in each time-frame

The formula to compute the optimal expected revenue is:
$J_n(x) = \max_f\{ \lambda.(Pr_{f,n}.(J_{n-1}(x-1) + p_f) + (1- Pr_{f,n}).(J_{n-1}(x))) + (1 - \lambda).J_{n-1}(x) \}$

		Decision periods		
x	Fare Class	3	2	1
1	1	<u>**163**</u>	118	<u>**80**</u>
	2	145	<u>**126**</u>	75
	3	113	90	50
2	1	175	120	<u>**80**</u>
	2	160	<u>**135**</u>	75
	3	<u>**181**</u>	130	50

Table 5: Computing the expected revenue for DP-GVR

Table 5 gives the computed values of $\lambda.(Pr_{f,n}.(J_{n-1}(x-1) + p_f) + (1- Pr_{f,n}).(J_{n-1}(x))) + (1 - \lambda).J_{n-1}(x)$ in each decision period and the values of $J_n(x)$ are underlined. The table is computed by starting on the right side. Whatever the number of remaining seats is in the last decision period we will try to get from them the maximum expected revenue. If we have one seat remaining which fare class should be the lowest open one in this last decision period? The arrival rate is 0.5, the probability of sell-up to fare class 1 is only 10% but the fare is $1600 which makes an expected revenue of $80 if fare class 1 is the lowest open fare class. By computing the 2 other expected revenues corresponding to fare class 2 or fare class 3 being the lowest open ones we notice that the best choice is to make fare class 1 the lowest open fare class. If we have 2 seats remaining the decision is the same as we do not expect to see more than one request by decision period. We then compute the values for decision periods 2 and 3 by taking into account the previously computed values.

We can deduce in each decision period what the lowest open fare class should be depending on the remaining number of seats. The results appear in *Table 6*. We see that if there is only one seat remaining in decision period 3 we may try to get a passenger in fare class 1 but if we have not sold the seat before decision period 2 we should then reopen fare class 2 during this new decision period. If the seat is still not sold during decision period 2 then we should during the last decision period try again to capture a passenger in fare class 1. If we have 2 seats remaining and we do not manage to sell them we should progressively close fare classes in each decision period.

	Decision periods		
x	3	2	1
1	1	2	1
2	3	2	1

Table 6: Lowest fare class to be open in each decision period

4. Results

This chapter presents the results obtained through the PODS simulation and their analysis. It is divided into two main sections concerning tests run in unrestricted and simplified fare structures. To be able to use traditional RM methods in unrestricted and simplified fare structures we will use Q-Forecasting and Hybrid Forecasting that were described in Chapter 3. In unrestricted fare structures we will use Network D or Network R as described in Chapter 3. Tests will be run against FCFS, against same method, against advanced competitors (EMSRb with Q-Forecasting) and against basic RM methods (AT80). We will present all the obtained results, analyze them and find out the advantages and limitations of the various methods. We will compare results obtained with methods based on dynamic programming to results obtained with traditional RM methods. In simplified fare structures we will compare in Network D the performance of DP-LB with Hybrid-Forecasting to the use of traditional RM methods using Hybrid-Forecasting.

4.1. Unrestricted Fare Structures

In unrestricted fare structures we have the choice to either use the Lautenbacher approach with Q-Forecasting as we would do with EMSRb or DAVN or the Gallego-Van Ryzin approach. The Gallego–Van Ryzin model appears to be more suitable as it assumes that people may sell-up or buy-down between classes. To use this method we need to estimate the probability that a passenger would be willing to book in any fare class if it is the lowest open fare class.

We will report the results obtained with both approaches based on dynamic programming and we will compare those results to those obtained with traditional RM methods: EMSRb, DAVN and the adaptive threshold method AT80.

4.1.1. Base Case in Network D

We first run tests in unrestricted Network D as described in Chapter 3: 2 airlines, 2 hubs and 40 spoke cities. The base case will still have both airlines using EMSRb with standard leg-based forecasting (Probabilistic Detruncation and Pick-up Moving Average Forecasting) which is equivalent to FCFS in unrestricted environments:

- Revenues or the 1st Airline: $747,952
- Load Factor of the 1st airline: 92.09%

The use of regular forecast quickly leads to spiral-down and more and more passengers end up booking in the lowest fare class as both airlines progressively make it the lowest open fare class until departure. We see the final stable consequence of using standard forecasting: bookings exclusively occur in fare class 6 as shown in *Figure 19*. EMSRB behaves as FCFS. As there are no restrictions between fare classes all passengers buy in the lowest open fare class and the fare class is closed only when the plane is full.

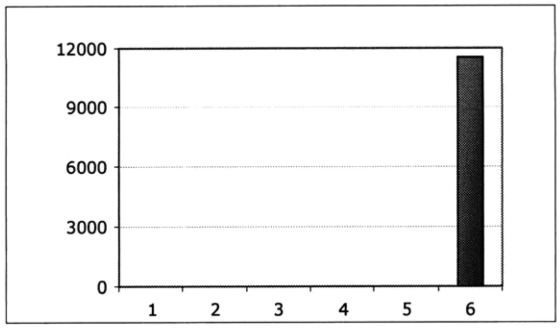

Figure 19: Fare class mix of Airline 1 – Base Case in unrestricted Network D

4.1.2. Against a Traditional RM Method (Network D)

We test the performance of DP-GVR or DP-LB with Q-Forecasting in comparison with other traditional RM methods when the second airline uses EMSRb with standard forecasting (equivalent to First Come First Served): the first airline uses EMSRb with Q-Forecasting, DAVN with Q-Forecasting, DP-LB with Q-Forecasting or DP-GVR while the second airline uses EMSRb with regular forecasting. We will compare the obtained revenues, load factors, closure rates of fare classes and fare class mixes.

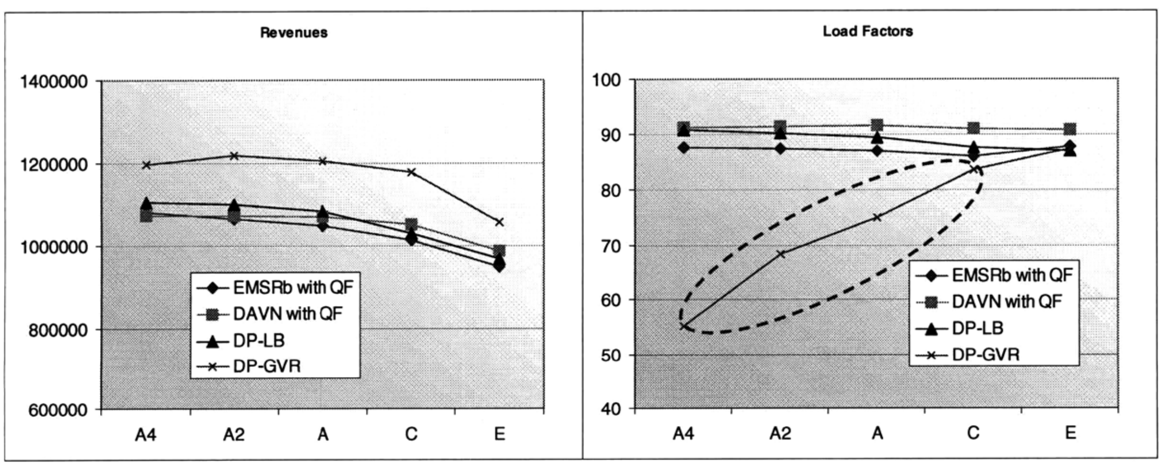

Figure 20: Results obtained against EMSRb with regular forecasting

The first graph of *Figure 20* represents the revenues of the first airline depending on the RM method it uses: EMSRb with QF, DAVN with QF, DP-LB with QF or DP-GVR. For each RM method five different FRAT5 sets have been used: A4, A2, A C and E. The set corresponding to the highest FRAT5s is A4 and the other sets use progressively lower values of FRAT5s. Each line corresponds to the same method tested with different values of FRAT5s used for Q-Forecasting or as input for DP-GVR. The second graph represents the corresponding load factors of Airline 1. The use of DP-LB with Q-Forecasting can lead to small improvements over EMSRb or even DAVN using Q-Forecasting, around 2%. However, DP-GVR gets results that are clearly better than results obtained with any other method using Q-Forecasting. The increase over the results of other RM methods using Q-Forecasting is around 13%, which represents more than 60% improvement over the base case.

Yet DP-GVR is very sensitive to FRAT5 inputs and a small difference in those estimations can result in high differences in load factors and substantial differences in revenues. The effect on results is not especially visible on the revenue graphs but we can see it on the graph showing the load factors: when switching from the set of FRAT5s A4 to the set of FRAT5s C the load factor is increasing from 55% to 83% while the change does not affect the results of other methods as strongly. Yet with a load factor of 55% DP-GVR using FRAT5s A4 still gets revenues that are much higher than those of other methods using Q-Forecasting with the same FRAT5s. This shows that DP-GVR manages to capture more and more demand in high fare classes when using higher, more aggressive, FRAT5s.

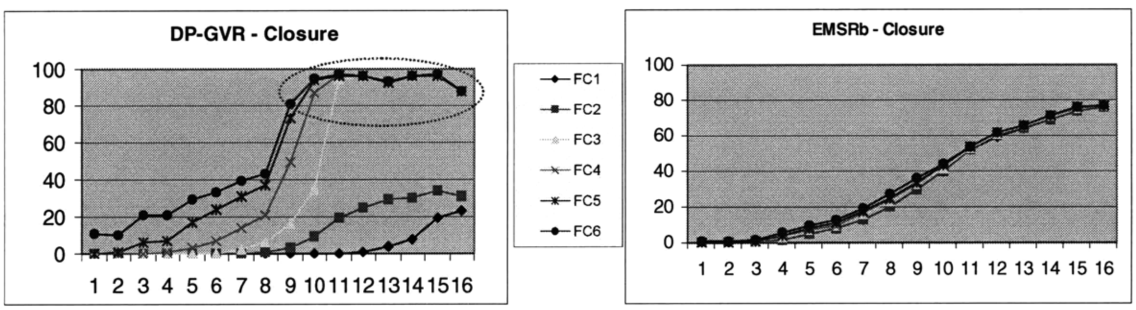

Figure 21: Comparison of the local fare class closures of the 2 airlines

The results shown in *Figure 21* are obtained with FRAT5 C. They represent for each fare class the percentage of local fare classes that are closed on local paths over the network in each time-frame. If for example there were 4 legs and if fare class 2 was closed on only one of them, the closure percentage of fare class 2 would be 25%. On the left graph we see the fare class closure of the Airline 1 - the airline that uses DP-GVR. We can first notice that it closes low fare classes earliest in the first time-frames and gradually close higher and higher classes until it runs out of seats on some legs at the end of the reservation period and consequently has to close some of the fare classes 1. Following the FRAT5 inputs DP-GVR closes low and middle fare classes early in the process in order to save seats for passengers with higher willingness-to-pay who arrive later in the reservation process. The second airline uses EMSRb with regular forecasting. Behaving as FCFS, it lets every person that is willing to travel at the lowest fare book a seat. Consequently all planes fill up and many planes run out of seats early in the reservation period allowing the second airline to capture demand with high willingness-to-pay.

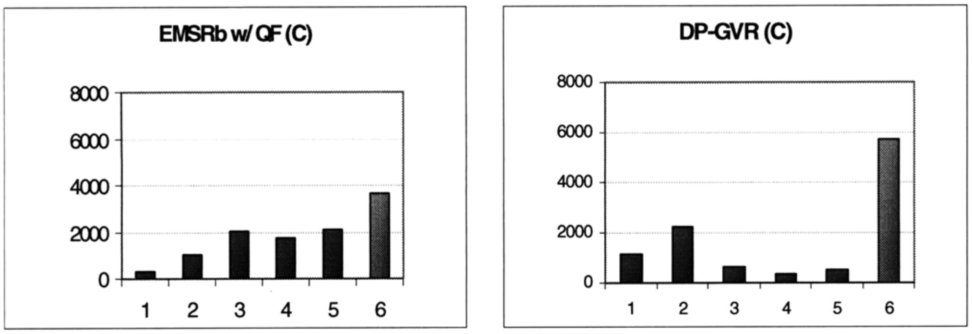

Figure 22: Fare class mixes of Airline 1 using EMSRb with QF or DP-GVR

When looking at the fare class mixes we see that Airline 1 using any RM method with Q-Forecasting is clearly benefiting from the inefficiency of the second airline. The left graph of *Figure 22* shows the fare class mix of Airline 1 using EMSRb with QF and FRAT5s C against EMSRb with regular forecasting. Airline 1 manages to capture demand with high willingness-to-pay as Airline 2 using FCFS runs out of seats. This illustrates how using Q-Forecasting with EMSRb can help prevent the spiral-down effect due to the lack of restrictions and advance-purchase requirements.

Yet DP-GVR whose fare class mix is shown on the right graph is the RM method that best takes advantage of the situation. It gets the best results when compared with EMSRb, DAVN or DP-LB using Q-Forecasting. It generates substantial loads in the highest fare classes. By closing low and middle fare classes early in the reservation process DP-GVR manages to capture the demand with high willingness-to-pay that is turned away by the Airline a which has filled its planes from the beginning of the reservation period with low-fare passengers. This shows that DP-GVR can manage to capture demand with high willingness-to-pay in unrestricted fare structures. Against a non-intelligent competitor it manages to get much better results than other methods. The Lautenbacher approach using Q-Forecasting gets results that are very close to those of traditional methods using Q-Forecasting.

Nonetheless, this case is not completely realistic as it assumes that Airline 2 lets anybody book a seat at the lowest price as long as planes are not full, that Airline 2 is not increasing the fare at all during all the entire reservation period while we could expect that it would at least use some basic threshold revenue management.

4.1.3. Against the Same Method (Network D)

In this section, in contrast with the previous case, both airlines now use the same RM method: same forecaster and same optimizer. Tests are still run in unrestricted Network D with 5 different sets of FRAT5s for EMSRb with Q-Forecasting, DAVN with Q-Forecasting, DP-LB with Q-Forecasting and DP-GVR.

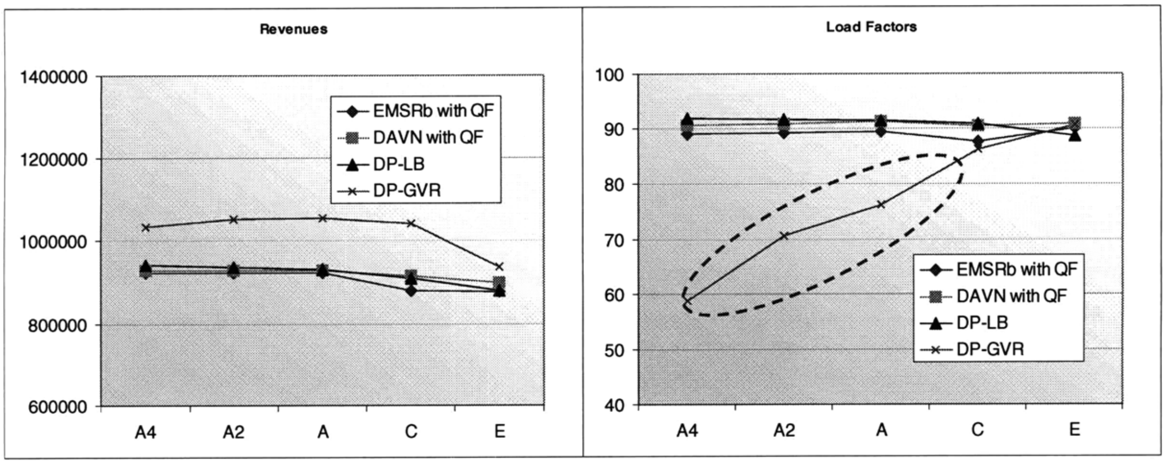

Figure 23: Results obtained against the same method

The results obtained by both airlines when they both use DP-GVR are once again much better than with any other RM method using Q-Forecasting. As shown in *Figure 23* in the best case the increase in revenue over the base case is 40%. This is less than in the case when only Airline 1 uses advanced revenue management but it is still a huge improvement over the use of FCFS. The revenue obtained with DP-GVR is 12% higher than the revenue obtained with any of the three other RM methods using Q-Forecasting. DP-GVR manages to get higher revenues with lower load factors. This shows that when both airlines use DP-GVR they are more efficient in capturing demand with high willingness-to-pay than if they were both using another RM method. We still notice a high sensitivity of DP-GVR to input FRAT5s. When using higher FRAT5s like A2 or A4 which assume a higher willingness-to-pay of the passengers, we get much lower load factors for both airlines.

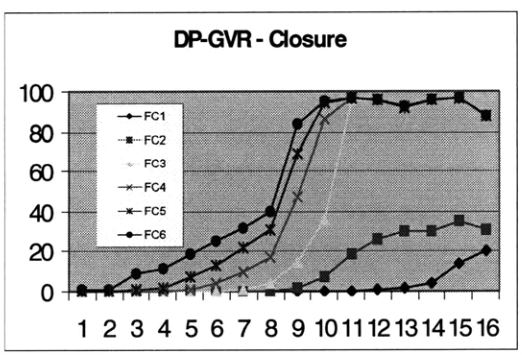

Figure 24: Local fare class closure of both airlines using DP-GVR

When both airlines use DP-GVR they are both closing low and middle fare classes early in the booking process (*Figure 24*). By both closing those fare classes the airlines manage to capture demand with higher willingness-to-pay. If lower fare classes had been open for either of the two airlines most of this demand with high willingness-to-pay would not have materialized as bookings in high fare classes; people would have bought down to the lower open fare classes.

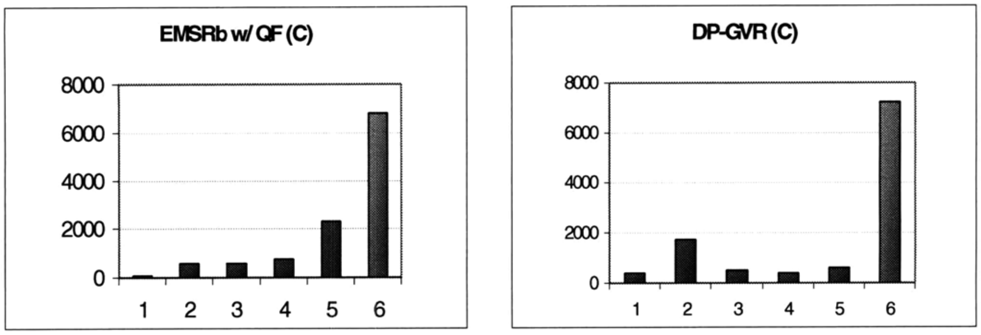

Figure 25: Fare class mixes of both airlines using EMSRb with QF or DP-GVR

Figure 25 shows the fare class mixes obtained by Airline 1 when both airlines use EMSRB with Q-Forecasting with FRAT5s C or DP-GVR with FRAT5s C. In the case when both airlines use EMSRB with Q-Forecasting we see that they capture less demand in higher fare classes than Airline 1 in the previous case when only Airline 1 was using EMSRb with Q-Forecasting and Airline 2 was using FCFS. Indeed at the end of the reservation period the airlines that have both saved seats for demand with high willingness-to-pay have to share this demand while in the previous case the second airline had filled its planes and Airline 1 could capture all the demand with high willingness-to-pay. When both airlines use DP-GVR we get similar results: both airlines get lower loads in high fare classes than in the previous case. Yet they still manage to capture significant loads. Indeed as both airlines close low and middle fare classes demand with high willingness-to-pay that comes at the end of the booking period has no other choice but to sell-up in one of the two high fare classes that are open.

4.1.4. Against EMSRb with Q-Forecasting (Network D)

In this section we now test the performance of EMSRb with Q-Forecasting, DAVN with Q-Forecasting, DP-LB with Q-Forecasting and DP-GVR against more advanced competitors. The competitor uses EMSRb with Q-Forecasting and FRAT5s A4. As we can see on the left revenue graph of *Figure 26* when competing against EMSRb with QF, DP-GVR no longer gets the best results in comparison with other RM methods using Q-Forecasting. It even gets slightly worse results. When using FRAT5s C or E DP-GVR manages to match other RM methods by getting similar revenues but these are not the FRAT5s leading to the best revenues of other RM methods. Other RM methods obtain their best revenues with higher FRAT5s, A2 or A4. However, when using those high FRAT5s with DP-GVR revenues drop as well as load factors. DP-LB with Q-Forecasting manages to get as good results as EMSRB with Q-Forecasting or DAVN with Q-Forecasting.

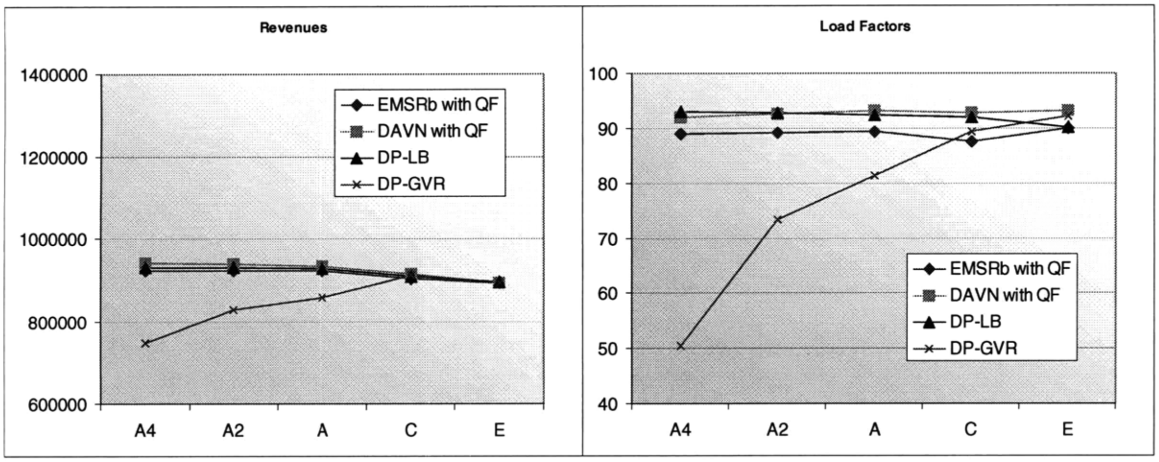

Figure 26: Results obtained against EMSRb with QF

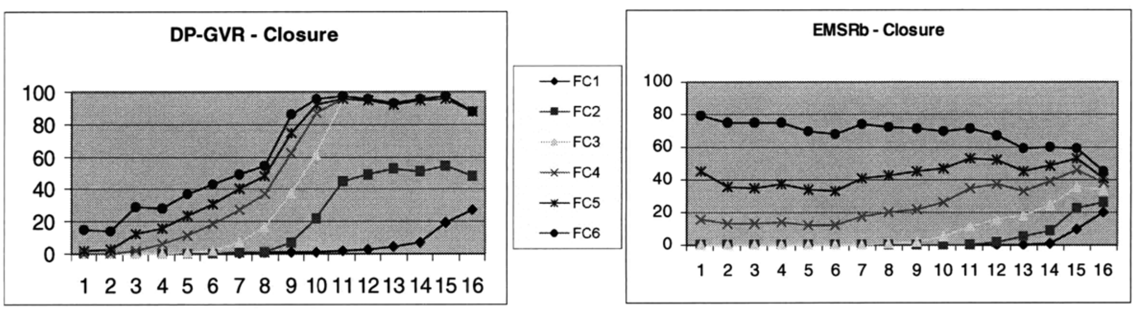

Figure 27: Comparison of the local fare class closures of the 2 airlines

As shown in *Figure 27* DP-GVR still suddenly closes low and middle fare classes in the middle of the reservation period while we can notice that EMSRb with Q-Forecasting closes a portion of them from the early time-frames. The airline using DP-GVR does not save enough seats for high fare classes: it starts closing some fare classes with the highest fare (FC1), meaning that it runs out of seats, earlier than the airline using EMSRb with Q-Forecasting. At the end of the reservation period Airline 2 still has many middle fare classes open while Airline 1 has closed most of them. People with high willingness-to-pay buy down and book seats in the middle fare classes of Airline 2. So DP-GVR using Q-Forecasting with input FRAT5s is getting bad results at both ends of the reservation period. First it does not manage to save enough seats in the first time-frames for demand with high willingness-to-pay that will arrive later and hence fills its planes with low-fare demand, and secondly it does not capture part of the demand arriving at the end of the reservation period as this demand buys down to the middle fare classes of the second airline.

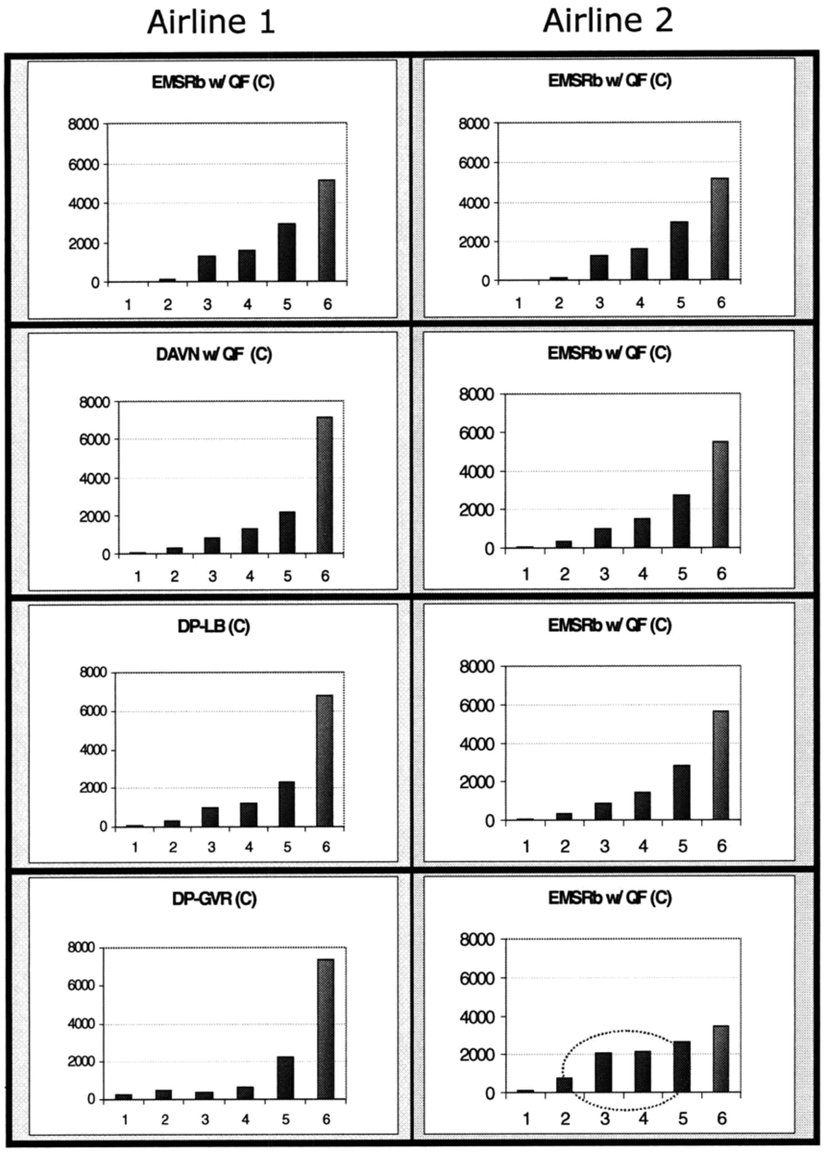

Figure 28: Fare Class Mixes obtained when Airline 2 uses EMSRb with QF

When we look at the fare class mixes shown in *Figure 28* the results obtained with DP-LB using Q-Forecasting are very close to those obtained with EMSRb using Q-Forecasting or DAVN using Q-Forecasting. In fact DP-LB appears to be very similar to EMSRb in its results. This may be explained by the underlying assumptions of both RM methods that are very similar. Both methods assume independence of demand for fare classes. They consider the partitioned Q-forecasts to compute the probabilities of selling a seat in a given fare class on a leg basis. They close a class when according to historical bookings and current time-frame they expect to get more revenue by selling the next seat in one of the higher fare classes. Yet results of both methods are not completely identical. There are several reasons to explain this difference. For example DP-LB uses path/class forecasts that are rolled up into time-frame leg/class forecasts while the forecasts for EMSRb are computed by leg and DP-

LB computes a policy for small decision periods while the bookings limits set by EMSRb are only updated in each time-frame.

Looking at the fare class mixes obtained with DP-GVR strengthens our analysis on the effect of the fare class closures of both airlines. DP-GVR only gets high loads in fare class 6 and substantial loads in fare class 5. But in all higher fare classes DP-GVR can only get very small loads. The inputs are the same. So the RM method is still expecting to get passengers with high willingness-to-pay at the end of the reservation period. But only a few passengers sell-up as most buy down to the second airline. Airline 1 forecasts very small arrival rates at the end of the reservation period as it does not notice that in fact it is the probability of sell-up that is very small. *Figure 29* shows the forecasts of equivalent Q-demand to come obtained with DP-GVR using Q-Forecasting and FRAT5s C against FCFS, EMSRb with Q-Forecasting and against the same method, from left to right. We notice that indeed against EMSRb with Q-Forecasting forecasts of equivalent Q-demand to come are lower at the end of the reservation period than in the other cases. Then the best solution for Airline 1 appears to be capturing demand at the beginning of the booking process when the input probabilities of sell-up are low and consequently it seems more beneficial to fill the planes with passengers in fare class 6.

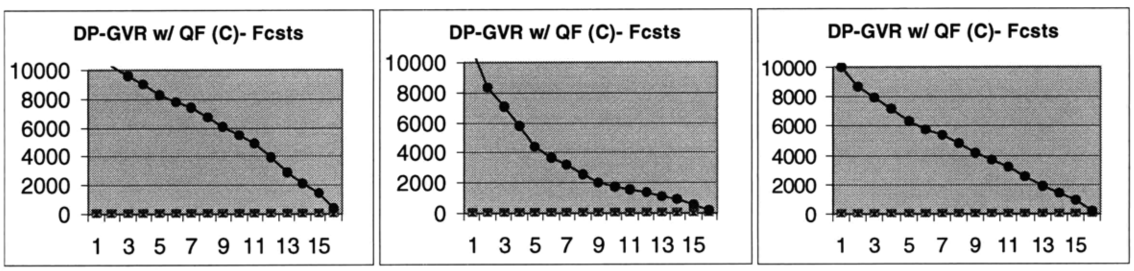

Against FCFS   Against EMSRb with Q-Forecasting   Against same method

Figure 29: Forecasts of equivalent Q-demand to come for DP-GVR with QF

Figure 30 illustrates DP-GVR sensitivity to input FRAT5s. When switching from the set of FRAT5s E to A4, DP-GVR is completely changing its pattern of fare class closure. With FRAT5s E DP-GVR is progressively closing fare classes from 6 to 2 among local paths. With FRAT5s A4 right from the beginning the low fare classes (5 and 6) are closed, and most middle fare classes are closed not long after the beginning of the reservation period. So at the end of the reservation period most planes still have many empty seats. This corresponds to the low observed load factors on *Figure 26*. For EMSRb with Q-Forecasting the change in closure pattern of fare classes is not as important when switching from the set of FRAT5s E to A4. For DP-GVR the decision to open or close a fare class is mainly based on FRAT5s, which are inputs in those simulations. EMSRb with QF is surprisingly more adaptive as its decisions are mainly based on the forecasts of bookings-to-come and as booking class limits are frequently updated: if bookings do not materialize EMSRb with QF adapts the closure of fare classes. Why do EMSRB and DP-GVR react so differently to input FRAT5s?

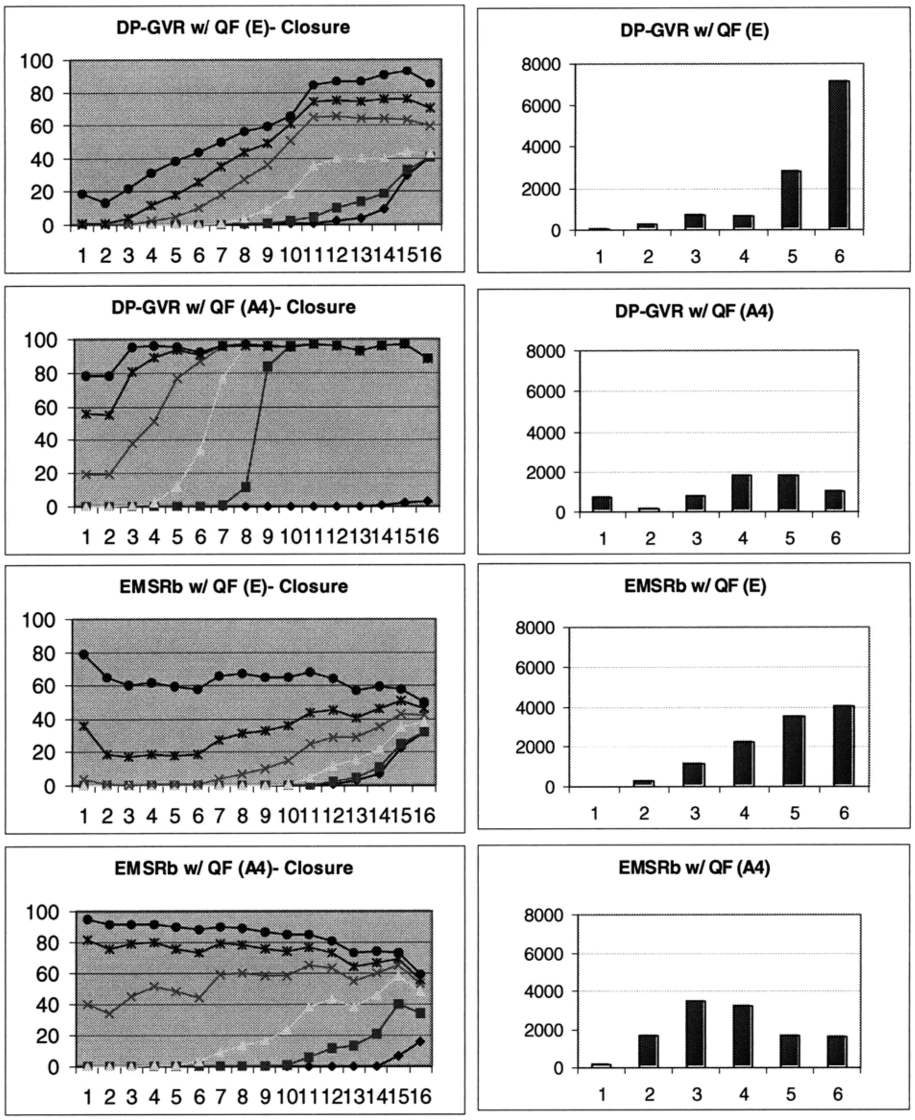

Figure 30: Various fare class closures and their corresponding fare class mixes

4.1.5. Q-Forecasting and Traditional RM Methods

DP-GVR seems to require lower FRAT5s than other RM methods and it seems to be more sensitive to a change in those input FRAT5s. Other methods are designed to be used under restricted fare structures with independent demand so they need boosted forecasts to run properly under unrestricted fare structure; this is the duty of Q-Forecasting. Yet EMSRb, DAVN and DP-LB have been designed to be used under fully-restricted fare structures so using them in unrestricted environments is not appropriate. In this section we will try to understand how they manage to get good results in unrestricted environments with Q-Forecasting and if there are limitations of using such traditional RM methods in comparison with newly developed RM methods that may be more suitable for unrestricted environments.

First we define observed product-demand as the demand materializing in open fare classes whatever the lowest open fare class is and price-oriented demand as the demand materializing only in the lowest open fare class. Those 2 types of demand can be compared to product-oriented and price-oriented demand as presented in Chapter 3 and introduced by **BOYD and KALLESEN**[54]. EMSRb can only use input forecasts of demand that is supposed to materialize whatever the lowest open fare class is: those forecasts are supposed to be forecasts of observed product-oriented demand and it is for such demand that this RM method was developed to set booking limits. Thus those forecasts assume that people will book in the highest fare class even if lower fare classes are open.

However, in unrestricted fare structures we should forecast potential demand rather than averaging demand observed in the past. Indeed at a given point in time demand will materialize only in the fare class that is the lowest open one. We can consequently not consider that "Q-Forecasting" can get rid of the sell-up/buy-down problem encountered by EMSRb with standard forecasting since the RM method itself is not adapted to consider potential demand in such environments. In fully-restricted fare structures forecasts of observed product-oriented demand can be obtained by averaging the observed detruncated demand by fare classes, independent of our open fare classes. But in unrestricted fare structures forecasts of potential price-oriented demand require the estimation of potential demand depending on the lowest open fare class.

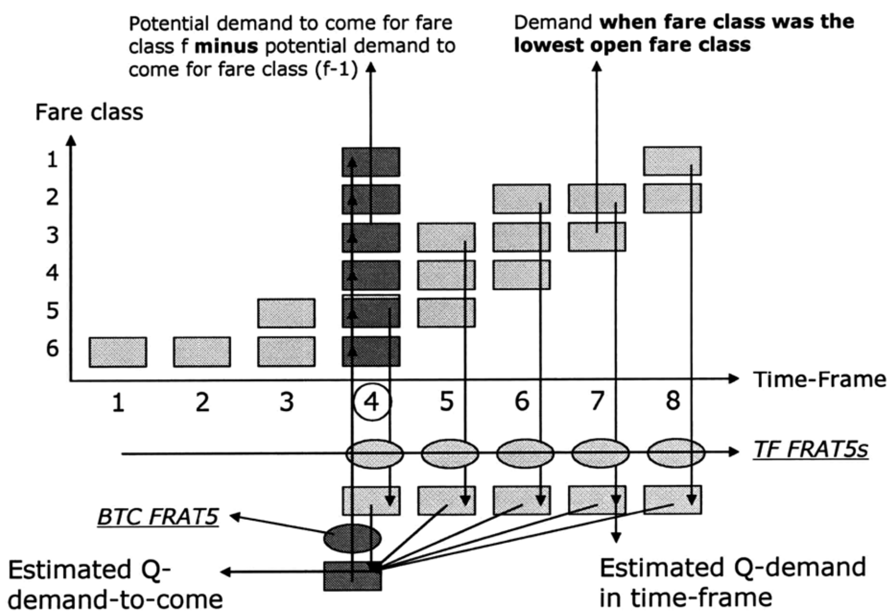

Figure 31: Obtaining equivalent Q-demand for EMSRb in one time-frame

[54] BOYD and KALLESEN, *The science of revenue management when passengers purchase the lowest available fare*, Journal of Revenue and Pricing Management, Vol. 3, 2004

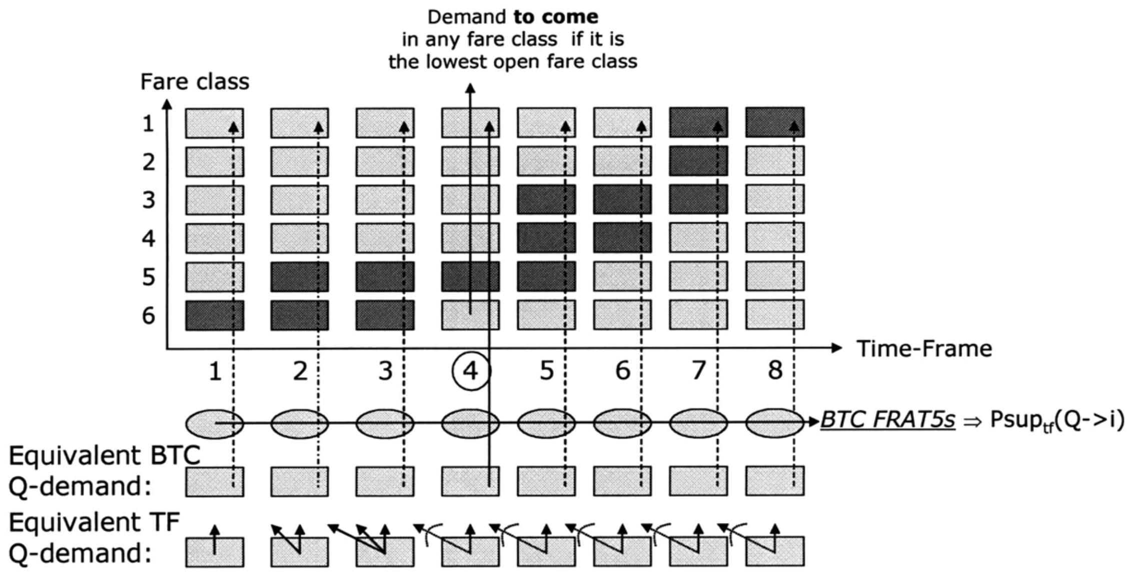

Figure 32: Obtaining all Potential demand from estimated Q-demand

Figure 31 and Figure 32 show how Q-Forecasting obtains the equivalent Q-demand-to-come to be used in EMSRb for a given time-frame when forecasting demand in unrestricted fare structures. On *Figure 31* the current time-frame is time-frame 4 and we see how to estimate the equivalent Q-demand-to-come (darker cells) for this time-frame. The lighter cells correspond to detruncated historical data obtained in each time-frame among several departures when class was the lowest open one. In a given time-frame various fare classes may have been the lowest open ones for different past departures which explains why there sometimes are cells for several fare classes in a time-frame.

The historical observations for future time-frames are converted into equivalent Q-demand by using the time-frame FRAT5s as symbolized by arrows going from top to bottom. For each future time-frame the FRAT5s are used to compute the probabilities of sell-up from class Q to any other class. According to those probabilities the equivalent Q-demand in each time-frame is deduced by dividing the detruncated historical demand for the fare class by the probability of sell-up from Q to this fare class. Then the equivalent Q-demands in each future time-frame are summed up in order to obtain the equivalent Q-demand-to-come.

This equivalent Q-demand-to-come represents the forecasted number of people that would agree to book a ticket at the Q-fare between the current time and departure. By using the FRAT5 to come in the current time-frame this equivalent Q-demand is converted into demand-to-come in any fare class as symbolized by the arrow going from bottom to top in time-frame 4. But if we really consider potential demand then the condition for this demand to materialize is that the fare class has to be the lowest open one until departure. Of course the probability is high that this will not be the case and that part of this demand will not materialize: forecasts are overestimated in comparison with historical data. The method may not be optimal yet it leads to results that are much better than letting everybody book in the lowest fare class. This adaptation of forecasting prevents complete spiral-down from happening. Depending on how the FRAT5s to come are computed we may manage to underestimate the probabilities of sell-up for high fare classes and consequently not to save too many seats for those high fare classes. By giving biased inputs we may finally get results that would be closer to optimality. But if those inputs are estimated in a way that reflects actual potential demand we will overestimate demand for high fare classes and get non-optimal revenues.

Figure 32 shows how we manage to forecast potential demand-to-come (symbolized by lighter cells) for each time-frame and each fare class from estimated equivalent Q-demand that was previously obtained from historical data (darker cells) for only some fare classes in each time-frame. We can then use those forecasts of demand-to-come with traditional optimizers or use them to compute probabilities of sell-up to be used with methods based on dynamic programming.

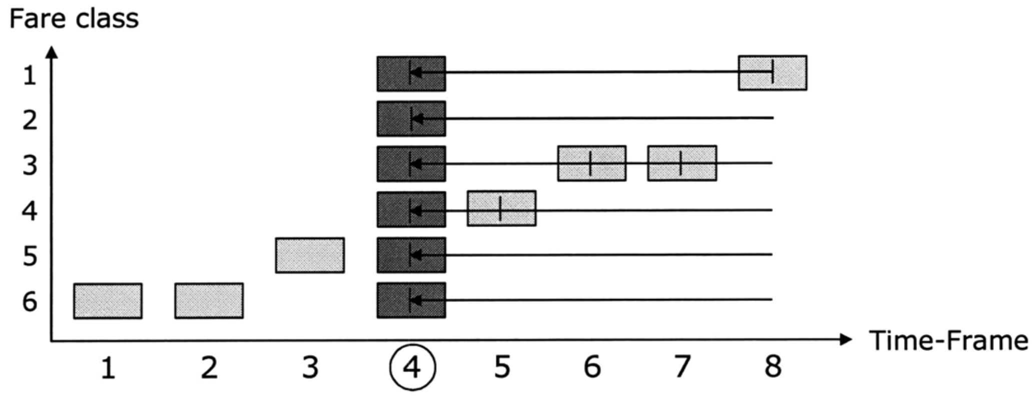

Figure 33: Observed demand-to-come for a given departure

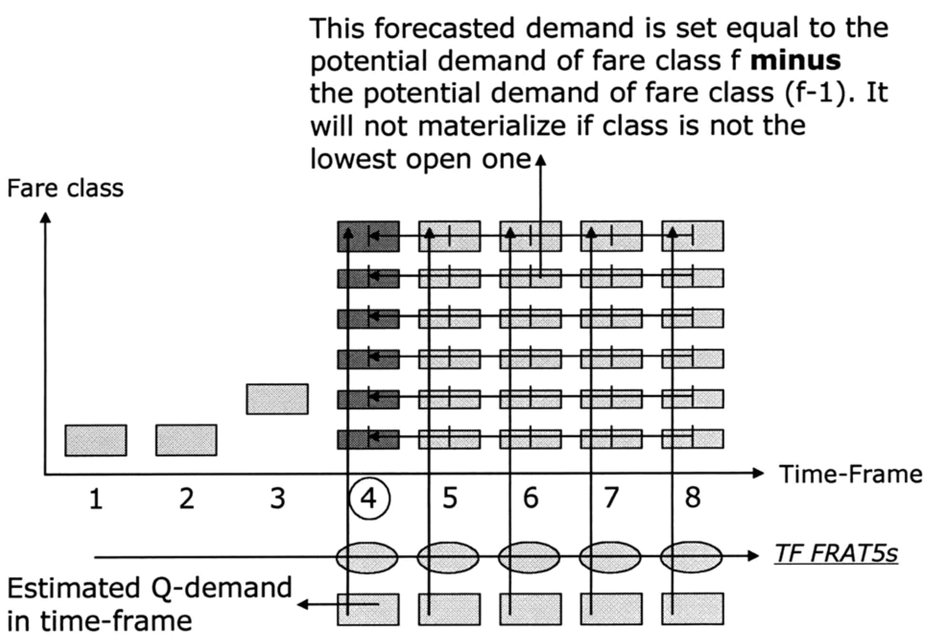

Figure 34: Currently used forecasted demand-to-come with EMSRb

As shown by *Figure 33* and *Figure 34* there is a real difference between the observed demand-to-come for a given departure and the "demand-to-come" that is forecasted by Q-Forecasting in order to be used with traditional RM methods such as EMSRb. *Figure 33* represents the demand that will materialize in each time-frame for a given departure. We see that in each time-frame demand materializes only in one fare class: the one that is the lowest open one in that time-frame. For example to be able to forecast the demand-to-come in time-frame 4 the forecaster needs to be aware of what the lowest open fare class will be in each future time-frame to then add the forecasted demand by fare class for all future time-frame.

Yet those lowest open fare classes are supposed to be determined by the optimizer which is traditionally applied after the forecaster and which consequently uses forecasts of bookings-to-come as inputs. There is a conflict with the traditional representation of revenue management: potential demand will materialize or not according to decisions taken by the optimizer. In unrestricted fare structures the optimizer needs to receive as input potential demand and to set a lowest open in each time-frame before being able to forecast demand-to-come.

Figure 34 shows an equivalent process to the one that is currently used with EMSRb when using Q-Forecasting. We see that demand-to-come is especially overestimated for fare class 1 also this estimation is less important in the time-frames that are close to departure as fare class 1 indeed as more chance to be the lowest open fare class. Subtracting potential demand for the next higher class from potential demand for any class may lead to less overestimation and consequently to a better solution. However, results are arbitrary and consequently may not be optimal. Nonetheless, those heuristic assumptions lead to results that appear to be more stable and less sensitive to the estimation of probabilities of sell-up than those obtained with DP-GVR.

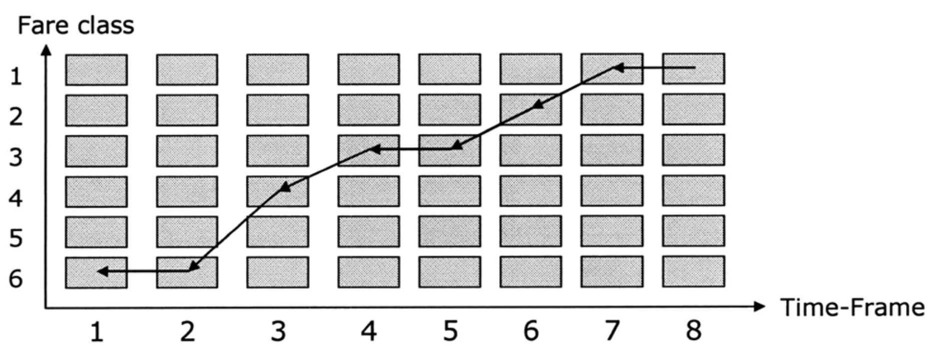

Figure 35: DP-GVR determination of lowest open fare class

DP-GVR determines in each time-frame what the lowest open fare class should be. *Figure 35* symbolizes the process followed by DP-GVR to choose the lowest open fare class in each time-frame in order to lead to optimal revenues. In each time-frame cells symbolize the potential demand that is forecasted to materialize if class is the lowest open one. The arrows symbolize the optimal way found by DP-GVR to choose a fare class closure pattern that leads to maximum revenues according to capacities. In fact decisions are taken over decision periods that are much smaller than a time-frame and during which the probability of seeing more than one passenger request is negligible. But we see that by following this process if potential demand could be perfectly forecasted in each time-frame and for each fare class then we would get a solution that would be close to optimality.

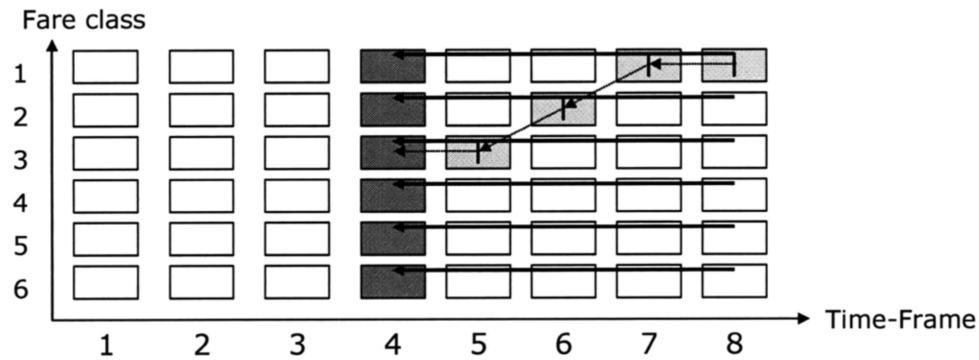

Figure 36: Getting the right bookings-to-come

Figure 36 shows how after having determined what the lowest open fare class will be in each time-frame we would be able to estimate how potential demand will materialize. Then we can compute the right amount of bookings-to-come in time-frame 4. In fact DP-GVR, which appears to be a theoretically more appropriate optimizer to use in unrestricted fare structures, does not even require forecasts of bookings-to-come as it directly uses forecasts of bookings in each decision period. However, this is also the reason why it is difficult to get forecasts that are theoretically required by DP-GVR: forecasts of the probabilities of sell-up and forecast of the arrival rate highly depend on parameters which are difficult to predict for a given future decision period, such as the seat availability of competition.

4.1.6. Against AT80 (Network D)

In this section we will test DP-GVR against AT80, an RM method that does not rely on Q-Forecasting, and compare its results to the results of other RM methods with Q-Forecasting competing against AT80 as well. AT80 is a threshold method that was developed to simulate the basic revenue management of a low-cost airline. We will still run tests in unrestricted Network D and try the same 5 sets of FRAT5s for Q-Forecasting: A4, A2, A, C and E.

When competing against AT80, Airline 1 using any of the 4 methods we have studied gets better results, 15-20% higher, than when performing against EMSRb with QF. As shown *in Figure 37*, by using FRAT5s E DP-GVR manages to get a 40% increase in revenues in comparison with the base case. This is not as good as the maximum increase obtained by other RM methods using Q-Forecasting which manage to get 41% increase in revenues over the base case with high input FRAT5s, A2 and A4. Yet DP-GVR still suffers high sensitivity to FRAT5s. As against EMSRb with Q-Forecasting, the load factors and revenues of the airline using DP-GVR against AT80 drop when trying higher input FRAT5s.

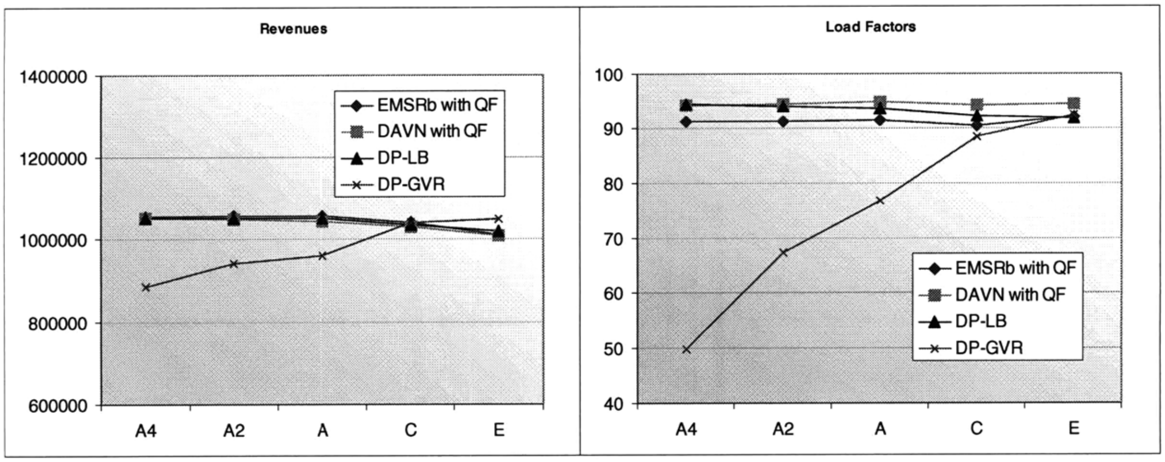

Figure 37: Results obtained against AT80

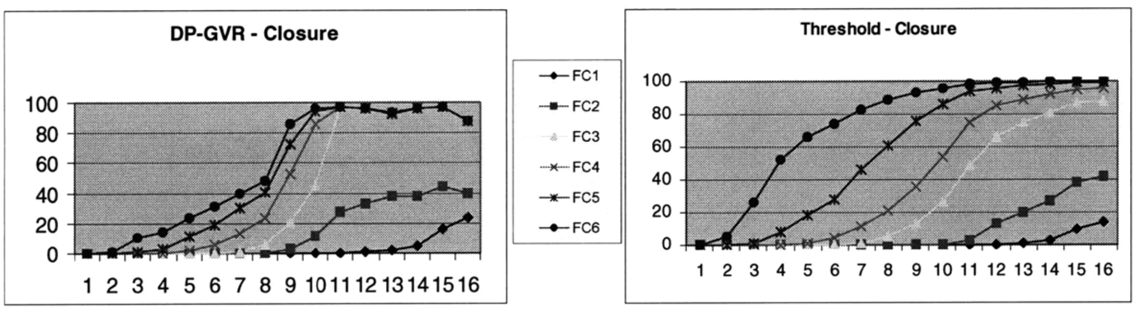

Figure 38: Comparison of the local fare class closures of the 2 airlines

Contrary to the performance of EMSRb with Q-Forecasting in the previous tests, AT80 does not close the low fare classes in the early time-frames and consequently it saves fewer seats for high-fare demand than EMSRb with Q-Forecasting (*Figure 38*). Moreover in the last time-frames the airline using AT80 is closing most of the middle-fare classes. This enables both DP-GVR and AT80 to capture high-fare demand at the end of the booking process while in the previous case demand with high willingness-to-sell-up was buying down in the open middle fare classes of the second airline. DP-GVR gets better loads in fare classes 1 and 2 against AT80 than when competing against EMSRb with QF (*Figure 39*). We can notice that EMSRb with Q-Forecasting is getting higher loads in middle fare classes than when competing against itself and that all the results it obtain, including fare class mixes, are very close to those obtained with the Lautenbacher DP approach.

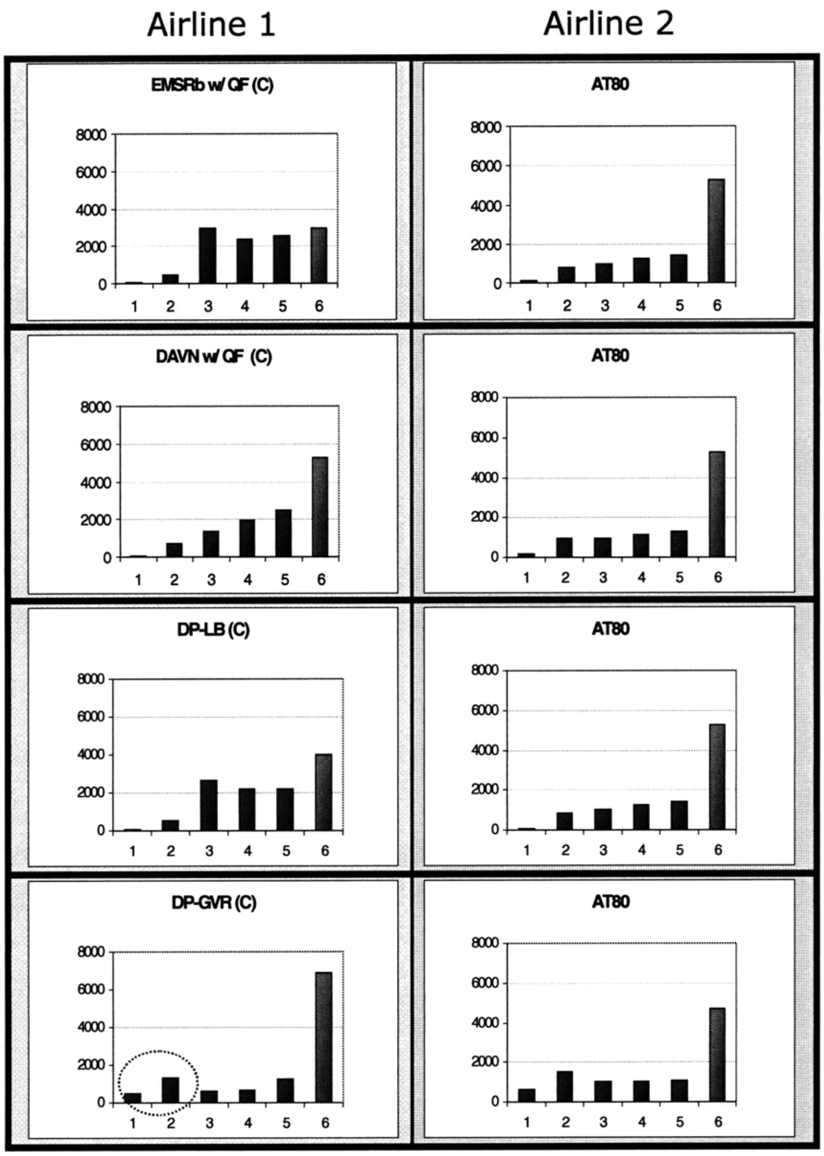

Figure 39: Fare Class Mixes obtained when the 2^{nd} airline uses AT80

4.1.7. Sensitivity of DP-GVR to Input FRAT5s

Figure 40 illustrates the sensitivity of DP-GVR to input FRAT5s. We see that changing input FRAT5s completely changes the closure pattern of fare classes. The top graph represents on the left the closure pattern of fare classes and on the right the bookings obtained by following this policy and is obtained with the set of FRAT5s A4. The bottom graph is obtained with the set of FRAT5s E. We see that when using FRAT5s A4 DP-GVR manages to get a large number of bookings in fare class 6 and in fare class 1 while when using FRAT5s E it still gets most bookings in fare class 6 but many in middle fare classes and only a few in the higher fare class.

DP-GVR (A4) vs FCFS

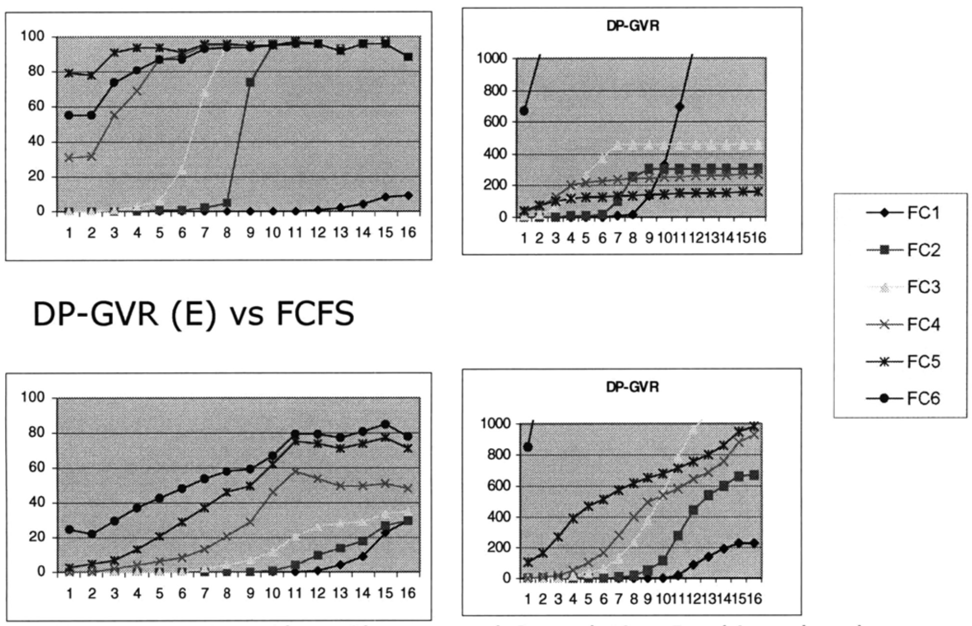

Figure 40: Fare Class Closure and Cumulative Bookings by class

4.1.8. <u>Summary of Results Obtained in Unrestricted Network D</u>

		RM method used by AL 2			
		EMSRb	Same Method	EMSRb with QF	AT80
RM methods used by AL 1	EMSRb with QF	44%	23%	23%	41%
	DAVN with QF	43%	24%	26%	41%
	DP-LB	48%	26%	25%	41%
	DP-GVR	63%	41%	22%	40%

Table 7: Summary of all best revenue increases over the base case in Network D

Table 7 shows a summary of results obtained in unrestricted Network D. Against no revenue management or against the same method, the current implementation of DP-GVR is the RM method leading to the best results while against AT80 or EMSRb with QF the RM method getting the best results for the moment is DAVN with Q-Forecasting. Those disappointing results of DP-GVR against advanced competitors are explained by the difficulty of forecasting probabilities of sell-up. At the moment tests have been run with input probabilities of sell-up which lead DP-GVR to make decisions relying on wrong assumptions and leading to situations where the RM method is not performing as well as it may perform with more accurate information on probabilities of sell-up.

4.2. Results in Unrestricted Network R

In this section we validate the results obtained in Network D in the bigger network with 4 airlines, Network R, as introduced in Chapter 3. The first base case corresponds to all airlines using EMSRb with regular leg-based forecasting (Probabilistic Detruncation and Pick-up Moving Average Forecasting):

- Revenues or the 1st Airline: $1,118,114
- Load Factor of the 1st airline: 93.24%

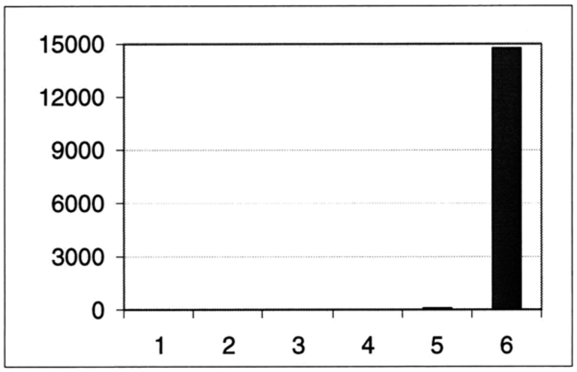

Figure 41: Fare Class Mix of Airline 1

As shown on *Figure 41* Airline 1 gets predominantly passengers in fare class 6 as EMSRb is equivalent to FCFS in unrestricted fare structures. If we were looking at the fare class mixes of the 3 other airlines they would be identical: there would be only passengers in fare class 6 with various loads depending on the airline. *Figure 42* shows the revenue and load factor of each of the 4 airlines. The load factors are identical: each of the 4 airlines is filling its planes with low-fare demand either until departure or until lack of seats. Airline 1 and Airline 2 have a similar situation in the network; they are the dominant airlines. Airline 4 has medium importance and Airline 3 captures only a small part of the total demand of the network.

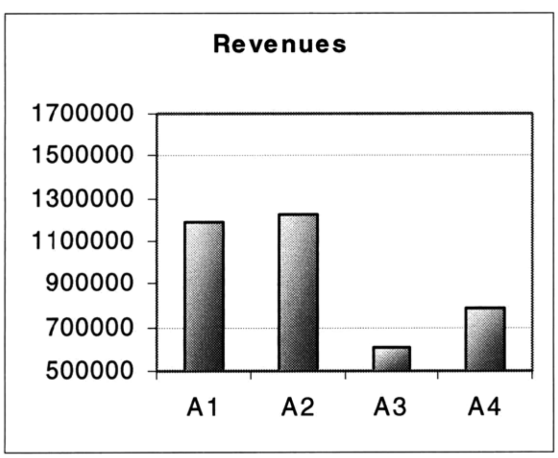

LFs: 93% 93% 94% 94%

Figure 42: Revenues and load factors of all 4 airlines

Figure 43 represents the revenues obtained by each airline when Airline 1 changes its RM methods. The 4 graphs correspond to the 4 optimizers we test for Airline 1: EMSRb with QF, DAVN with QF, DP-LB with QF and DP-GVR. On each graph there are 4 curves, one for the results of each airline. For each optimizer we test 4 sets of FRAT5s: A4, A, C and E. We see than in all cases the best results for Airline 1 are obtained when using FRAT5s A4 or A. Consequently, we will focus on the results obtained with FRAT5s A. *Figure 44* shows the revenues and load factors of all 4 airlines obtained for Airline 1 using any of the 4 tested RM methods with FRAT5s A. In all cases Airline 1 is the carrier getting the best revenues. This result is not surprising as it is the only airline that uses advanced revenue management (all 3 other airlines use FCFS). DP-LB with Q-Forecasting gets better results - approximatively 65% increase in revenues over the base case - than EMSRb with Q-Forecasting or DAVN with Q-Forecasting, which get approximatively 60% increase in revenues over the base case. DP-GVR is the RM method leading to the best results: a 71% increase in revenues over the base case. Yet this case is not realistic: the 3 competitors fill their planes with low-fare demand while Airline 1 is the only one to save seats for demand with higher willingness-to-pay.

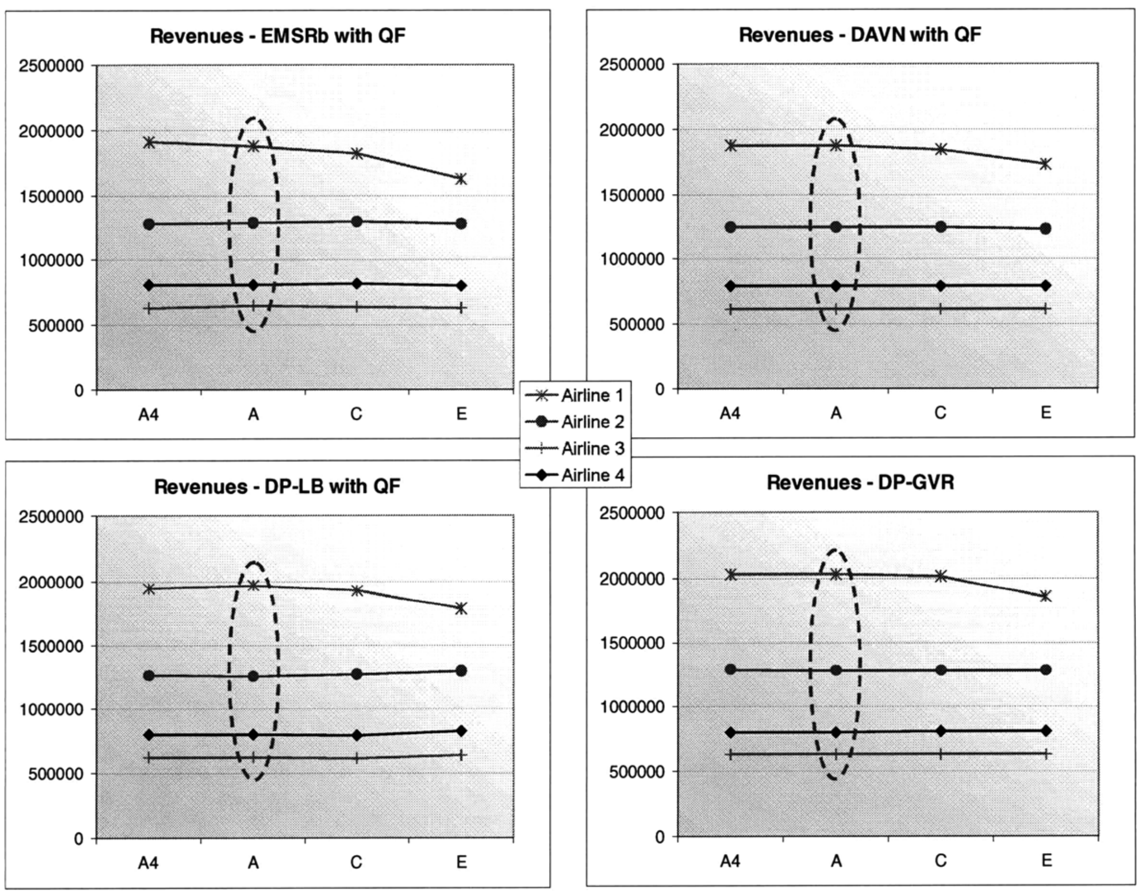

Figure 43: Revenues of all 4 airlines when Airline 1 uses various RM methods

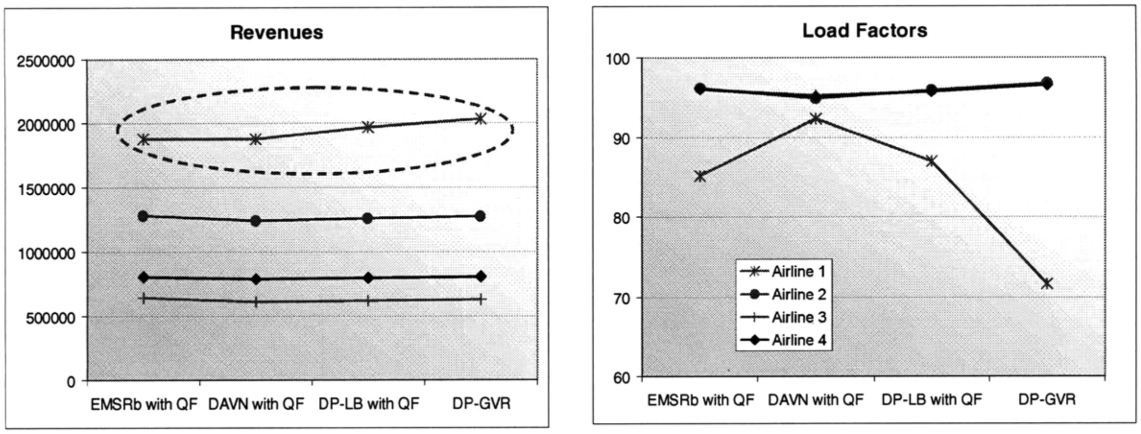

Figure 44: Revenues of all 4 airlines when using various RM methods (FRAT5s A)

Figure 45, *Figure 46* and *Figure 47* represent the fare class mixes of Airline 1 and the cumulative booking curves of each of the 4 airlines during the reservation period. These 3 graphs are obtained respectively for Airline 1 using EMSRb with QF, DP-LB with QF and DP-GVR. We can notice that the cumulative booking curves resulting from the use of EMSRb with QF or DP-LB with QF are very similar although there is a substantial difference in fare class mixes of Airline 1. However, DP-GVR leads to completely different results. We again find the pattern of fare class mix, obtained in Network D. Loads are high in fare class 1 and fare class 6 and low in other fare classes. Airline 1 using DP-GVR saves seats for passengers in fare class 1 according to high input FRAT5s. The other airlines fill their planes up letting all passengers booking in fare class 6. At the end of the booking process passengers have no choice but to sell-up to fare class 1 of Airline 1.

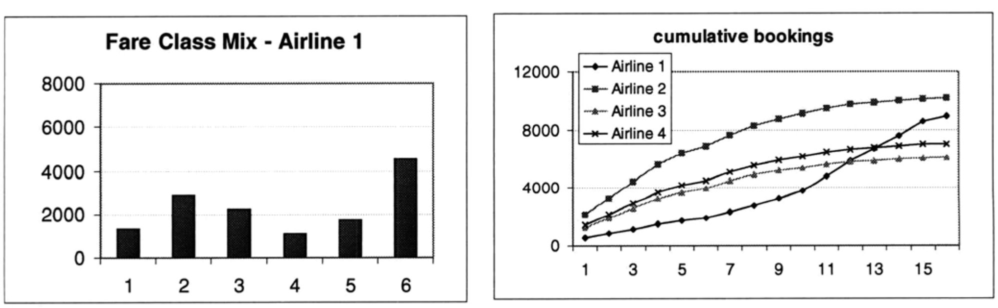

Figure 45: Results of Airline 1 - EMSRb with QF

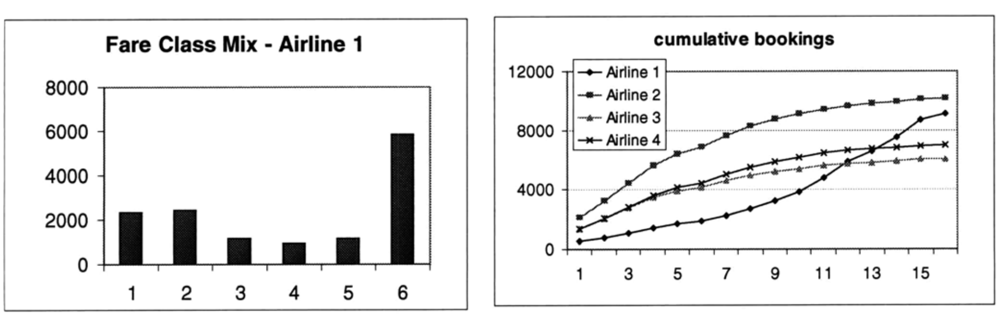

Figure 46: Results of Airline 1 - DP-LB with QF

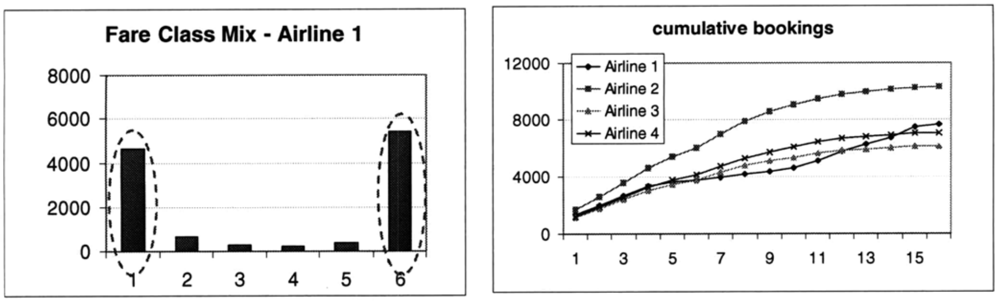

Figure 47: Results of Airline 1 - DP-GVR

We now consider a second base case corresponding to the 3 other airlines using advanced revenue management methods while Airline 1 will still be using one of the 4 RM methods we want to compare: EMSRb with Q-Forecasting, DAVN with Q-Forecasting, DP-LB with Q-Forecasting or DP-GVR. Airline 2 and 4 use EMSRb with Q-Forecasting and Airline 3 uses AT80. The results obtained by the Airline 1 when it uses EMSRb with standard forecasting are as follows:

- Revenues or the 1st Airline: $1,242,975
- Load Factor of the 1st airline: 96.17%

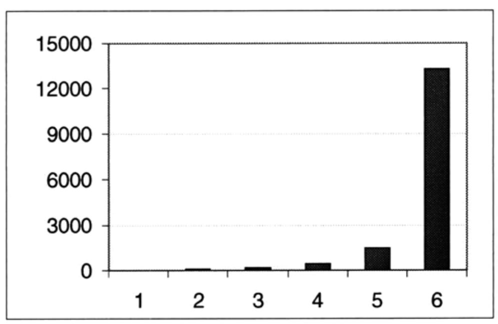

Figure 48: Fare Class Mix of Airline 1

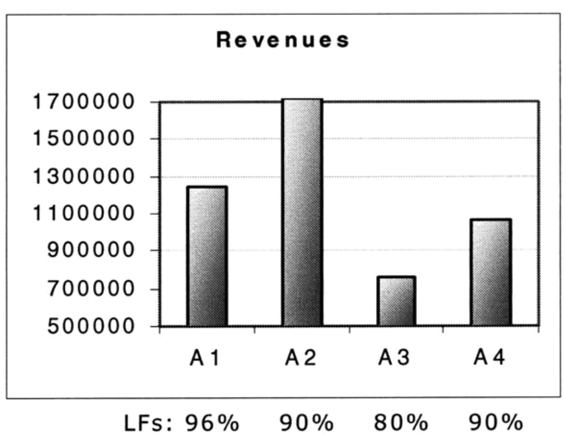

Figure 49: Revenues and load factors of all 4 airlines

Figure 48 shows the fare class mix of Airline 1 when using EMSRb with standard forecasting. *Figure 49* shows the revenues and load factors of the 3 airlines. Airline 2 uses EMSRB with Q-Forecasting. In comparison with the previous base case Airline 2 is now performing much better than Airline 1 when Airline 1 uses EMSRb with standard forecasting. Airline 4 also uses EMSRb with Q-Forecasting and almost doubled its revenues in comparison with the previous base case. Also, Airline 3 benefits from the use of AT80, though it may get even better results if it used a more advanced revenue management method. When using AT80, which is designed to aim at a load factor of 80%, Airline 3 gets much lower load factors than other airlines. Airline 1 does not lose revenues in comparison with the previous base case because there is enough demand to allow Airline 1 to keep filling its planes with low-fare passengers.

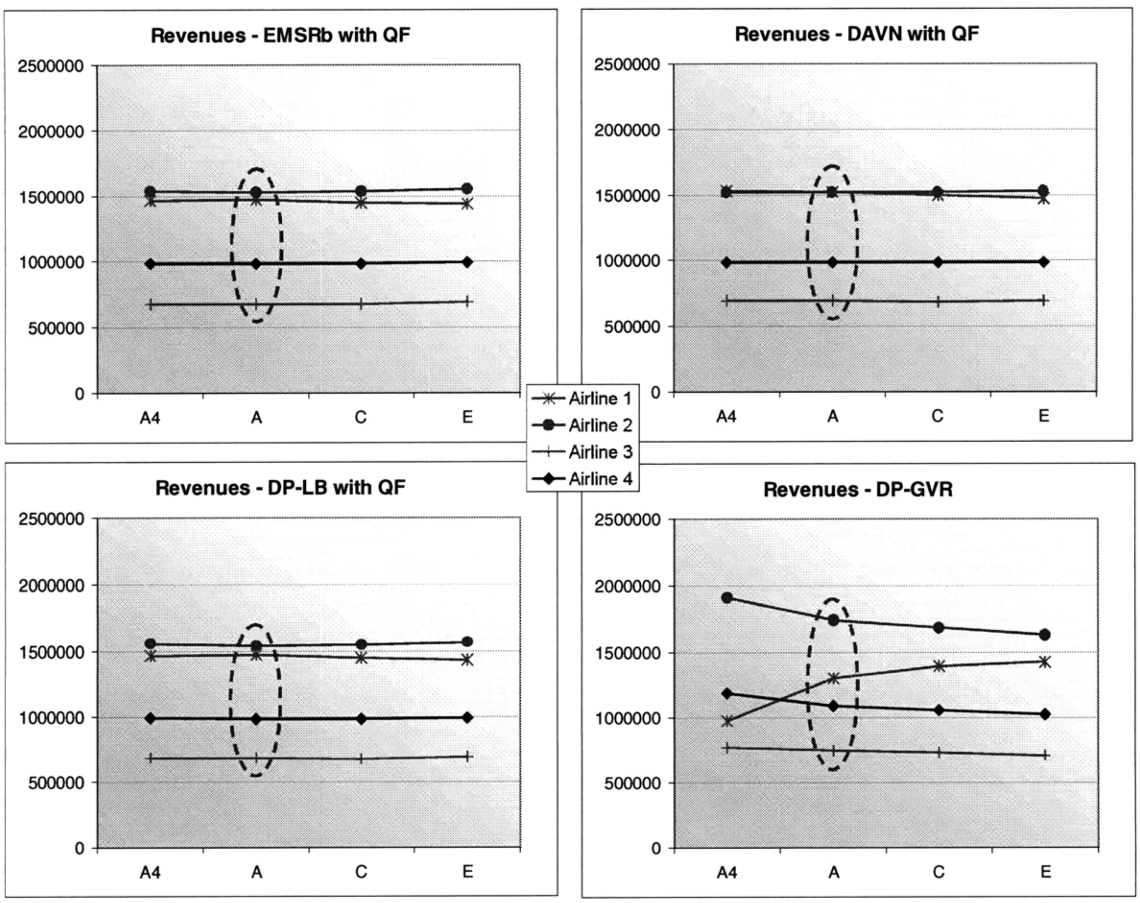

Figure 50: Revenues of all 4 airlines when using various RM methods

Figure 50 shows the revenues obtained by each airline when Airline 1 changes its RM methods: EMSRb with QF, DAVN with QF, DP-LB with QF or DP-GVR. Now when Airline 1 uses EMSRb with Q-Forecasting, DAVN with Q-Forecasting or DP-LB with Q-Forecasting it performs just as well as Airline 2. But when it uses DP-GVR with input FRAT5s it does not match the results of Airline 2 whatever the input set of FRAT5s is among the 4 sets of FRAT5s tested. Airline 1 manages to get revenue increases over the base case that reach 18% for both EMSRb with Q-Forecasting and DP-LB with Q-

Forecasting, 23% for DAVN with Q-Forecasting and 14% for DP-GVR. In the case of DP-GVR all 3 other airlines clearly benefit from the inefficiency of Airline 1 in capturing high-fare demand. *Figure 51* presents the results obtained specifically with FRAT5s A. In the case of DP-GVR there is both a drop in revenues and in load factor.

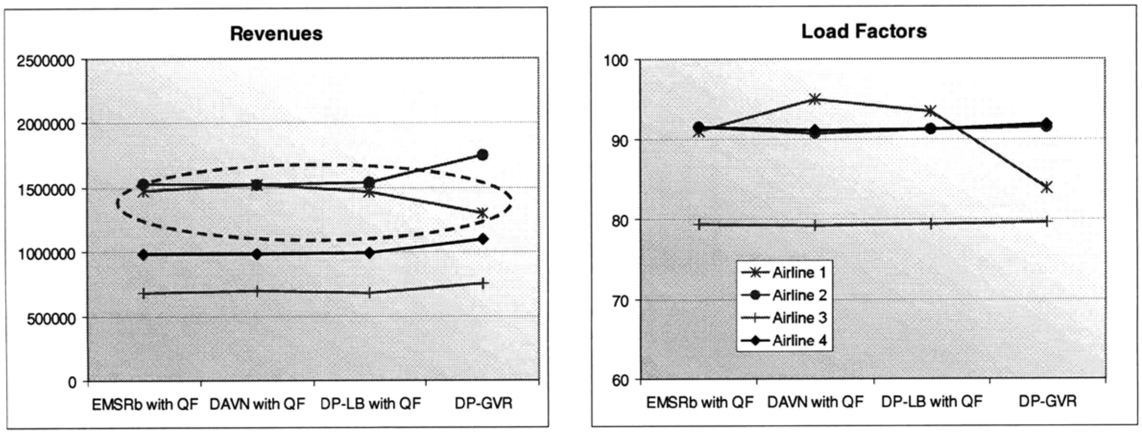

Figure 51: Revenues of all 4 airlines when using various RM methods (FRAT5s A)

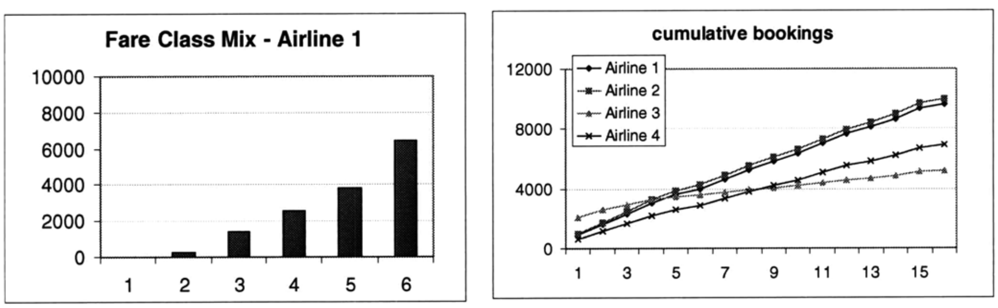

Figure 52: EMSRb with QF against Eb-QF/AT80/Eb-QF

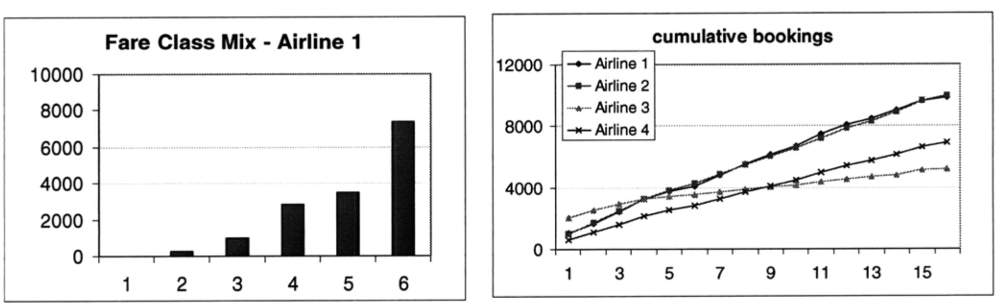

Figure 53: DP-LB with QF against Eb-QF/AT80/Eb-QF

Figure 52, *Figure 53* and *Figure 54* show the fare class mixes and the cumulative booking curves corresponding to the case when Airline 1 uses EMSRb with Q-Forecasting, DP-LB with Q-Forecasting or DP-GVR. Once again the fare class mixes and the cumulative booking curves obtained with EMSRb with Q-Forecasting and with

DP-LB with Q-Forecasting are close while the results obtained with DP-GVR are very different. DP-GVR captures low-fare demand at the beginning of the booking process and then behaves as if there was no demand arriving. At the end of the booking process input FRAT5s are higher because business passengers are expected to represent a more important part of the demand arriving. But if other airlines have lower open fare classes then this demand with high willingness-to-pay buys down and DP-GVR records a low arrival rate. Consequently it then focuses on the beginning of the reservation period where demand has a low willingness-to-pay. At the end of the reservation period it keeps trying to capture demand with high willingness-to-pay that will never materialize.

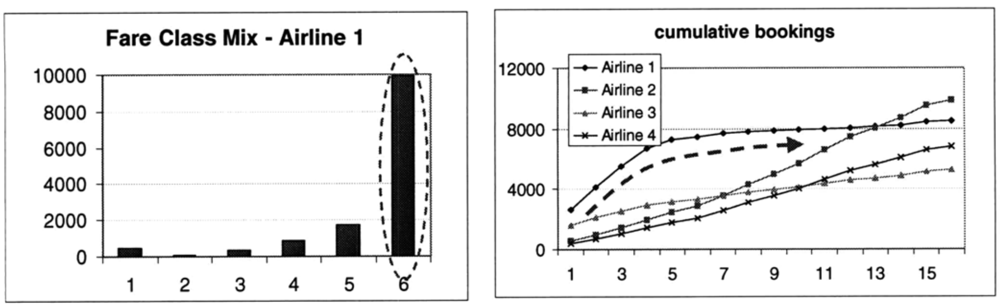

Figure 54: DP-GVR against Eb-QF/AT80/Eb-QF

The results obtained in a bigger network are strengthening the results we had obtained in Network D. DP-LB with QF gets results that are similar to those obtained with EMSRb with QF so its use does not represent an improvement in unrestricted fare structures. DP-GVR gets patterns of fare class closure that are very different from those obtained with EMSRb. Yet those patterns rely on the probabilities of sell-up that are either input or estimates. In the implementation we simulated those probabilities of sell-up were inputs and we see the limitations of having a method that is not adaptive: DP-GVR does not adapt its closure of fare classes to be more competitive when demand buys down as its decision rely on inputs, even if those inputs do not correspond to reality. Against competitors using no revenue management or against a competitor using the same RM method, DP-GVR shows very good results. But against a competitor that uses advanced RM management the tested implementation of DP-GVR does not manage to adapt its closure of fare class as it relies on forecasts that are not adaptive and it then gets worse results than other RM methods using Q-Forecasting.

In unrestricted environments to capture demand for highest fare classes airlines have to close all lower fare classes because passengers only book tickets in the lowest open fare class. We will now study the performance of DP-Lautenbacher in simplified fare structures that are intermediate fare structures between fully-restricted fare structures and unrestricted fare structures. In simplified fare structures there are fewer restrictions than in fully-restricted fare structures: some classes have different restrictions but some classes are still undifferentiated.

4.3. Simplified Fare Structures

The second part of our results was obtained for simplified fare structures in Network D. In semi-restricted environments we can not anymore assume independence of fare classes. Some classes have the same restrictions: they are undifferentiated. If the lowest open fare class is one of those undifferentiated fare classes, some people may have agreed to book in a higher fare class that has the same restrictions as the lowest open one but they will buy down to get the lowest available price. Thus *Table 8* shows that fare classes 3, 4, 5 and 6 are undifferentiated. If all fare classes higher than 5 are open people will only book in fare classes 1, 2 and 5. Indeed Fare classes 1 and 2 have fewer restrictions than lower fare classes so some passengers are ready to pay a higher price to get a ticket in those fare classes. But fare classes 3 and 4 have the same restrictions as fare class 5 so passengers buy down to fare class 5 to pay a lower price. Hybrid Forecasting was developed to forecast demand in such simplified fare structures as explained in Chapter 3. We will analyze the results obtained with Hybrid Forecasting used with EMSRB, DAVN and the Lautenbacher DP approach. DP-GVR is for the moment designed to be used exclusively in unrestricted fare structures. DP-GVR has not yet been adapted to simplified fare structures so we will focus on the performance of DP-LB.

FC	AP	R1	R2	R3
1	0	0	0	0
2	0	0	1	0
3	7	0	1	1
4	14	0	1	1
5	14	0	1	1
6	21	0	1	1

Table 8: Simplified Fare Structure

4.3.1. Base Case in Network D

The base case still corresponds to both airlines using EMSRb with standard leg-based forecasting (Probabilistic Detruncation and Pick-up Moving Average Forecasting):

- Revenues or the 1st Airline: $1,029,823
- Load Factor of the 1st airline: 83.6%

As represented on *Figure 55* even in the base case there are now bookings occurring in all fare classes. Indeed the phenomenon of complete spiral-down that occurs in unrestricted fare structures does not occur in simplified fare structures. In simplified fare structures people still book in highest fare classes that have different restrictions from other fare classes. In the reported fare class mix we differentiate between bookings that were classified as product-oriented and the remaining booking that were classified as price-oriented. Product-oriented bookings are bookings that were made in a fare class while lower fare classes were still open. Price-oriented bookings are the remaining bookings. There are two kinds of price-oriented bookings: AP bookings or RM bookings. AP bookings are the bookings made in a fare class when the next lower class was closed because of AP requirements and RM bookings are the bookings made when the next lower class was closed by revenue management. In semi-restricted environments even EMSRb with standard leg-based forecasting manages to get good results and to capture demand in high fare classes since some product-oriented passengers are ready to pay a higher price to get a ticket with fewer restrictions.

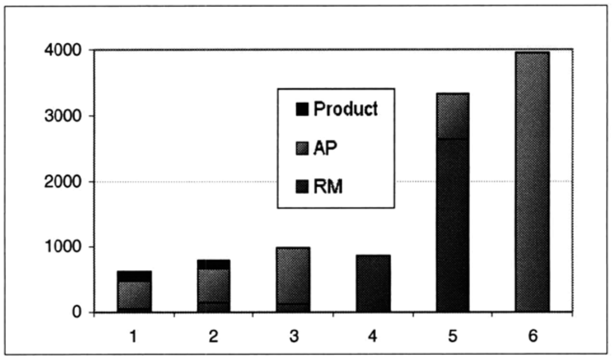

Figure 55: Fare class Mix of Airline 1 – Base Case in semi-restricted Network D

4.3.2. Benefit of Hybrid Forecasting for EMSRb and DAVN

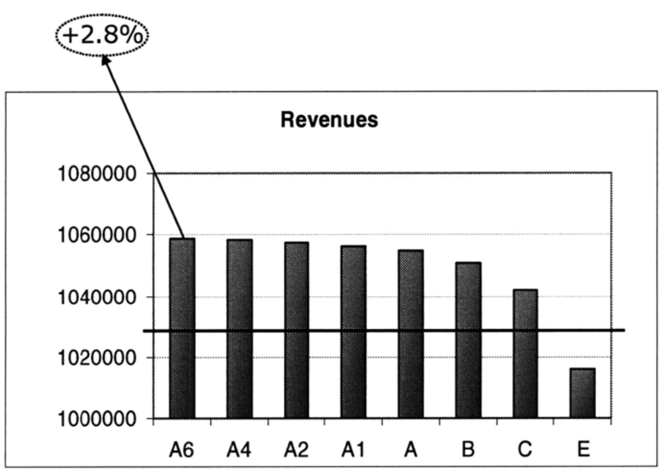

Figure 56: Revenues of Airline 1 using EMSRb with HF against FCFS

In this section we test the benefit of Hybrid Forecasting when used with EMSRb and DAVN, before looking in the next section to results obtained with the Lautenbacher dynamic programming approach. *Figures 56* show the revenues of Airline 1 using EMSRb with Hybrid Forecasting and different sets of FRAT5s when the second airline uses EMSRb with standard forecasting. For Hybrid Forecasting we used HF1 and IAP0. Airline 1 manages to get a 2.8% increase in revenues in comparison with the base case. Yet if low FRAT5s are used, like E, then revenues fall under those of the base case. *Figure 57* show similar results for DAVN. The dotted line represents the revenues of Airline 1 using DAVN against EMSRb with standard forecasting. Airline 1 using DAVN with Hybrid Forecasting gets higher revenues and manages to get a 4.5% increase in revenues.

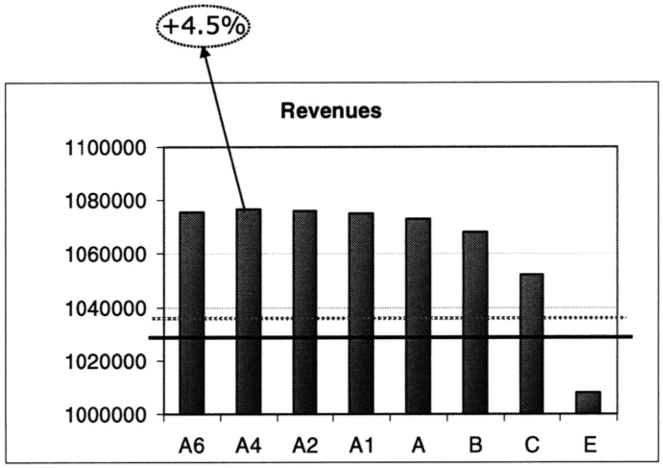

Figure 57: Revenues of Airline 1 using DAVN with HF against FCFS

Figure 58 shows the fare class mixes of Airline 1 using EMSRb or DAVN with Hybrid Forecasting (HF1, IAP0 and FRAT5s A4). The percentage in the left top corner is the load factor. We see that DAVN manages to get higher revenues by capturing more demand but not especially in high fare classes. DAVN captures as much demand in high fare classes but also gets higher loads in fare class 6. Those higher revenues are in part obtained through higher load factors.

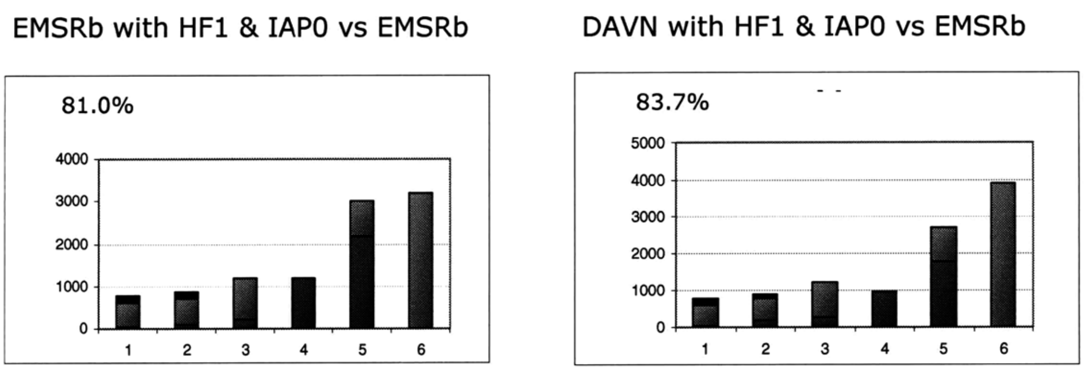

Figure 58: Fare class mixes of Airline 1 using HF (FRAT5 A4)

4.3.3. Use of Hybrid Forecasting with DP

We now test the performance of DP-LB with Hybrid Forecasting in semi-restricted Network D. As exposed previously the current approach of Gallego–Van Ryzin is for unrestricted fare structures and has not yet been adapted to simplified fare structures. First we test Airline 1 using DP with path-based standard forecasting against EMSRb with standard forecasting:

- Revenues of the 1st Airline: $1,045,199 (1.5% revenue increase over the base case Eb vs Eb both with standard forecasting)
- Load Factor of the 1st airline: 79.7%

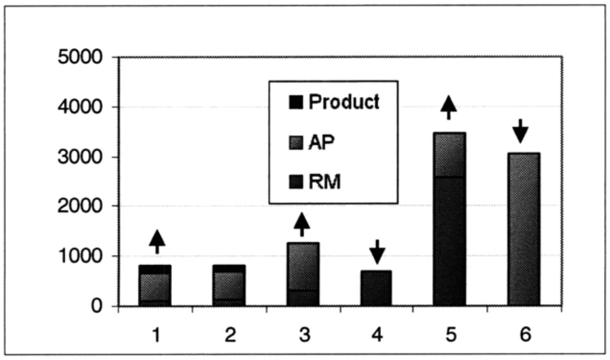

Figure 59: Fare class Mix of Airline 1 – DP Base Case in semi-restricted Network D

Figure 59 shows the fare class mix of Airline 1 using DP-LB with standard forecasting against EMSRb with standard forecasting. We see that DP-LB alone manages to get good loads in high fare classes. In comparison with loads obtained by EMSRb with standard forecasting in the base case, DP-LB with standard forecasting manages to get higher loads in fare classes 1 and 3 by reducing the number of passengers accepted in fare classes 4 and 6. The arrows show the evolution of loads by fare class in comparison with loads obtained in the base case.

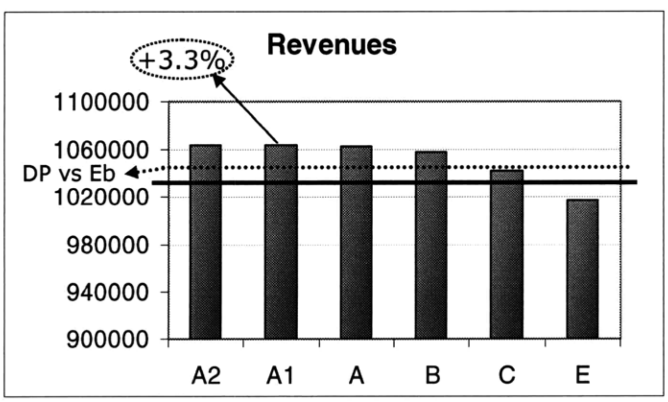

Figure 60: Revenues of Airline 1 using DP with HF against Eb with std. Fcst.

Figure 60 shows that DP-LB using hybrid-forecasting with the right set of FRAT5s manages to get a 3.3% increase in revenues over the revenues of the base case. This represents a 1.7% increase over the revenues obtained with DP-LB and standard forecasting. This is better than the results obtained with EMSRb with Hybrid Forecasting and less then those obtained with DAVN using Hybrid Forecasting.

We now test DP-LB with Hybrid Forecasting competing against the same method. For Hybrid Forecasting we use HF1, IAP0 and various FRAT5s. *Figure 61* shows the increases in revenues Airline 1 manages to get over the base case. The optimal increase in revenues over the base case is 2.4% which is more than when EMSRB with Hybrid Forecasting competes against itself (1.9% increase in revenues over the base case) and just slightly less than when DAVN with Hybrid Forecasting competes against itself (2.5% increase in revenues). Once again we see that DP-LB that is not a real O-D optimizer even if it uses O-D forecasts manages to get results that are in between those obtained with EMSRb and those obtained with DAVN.

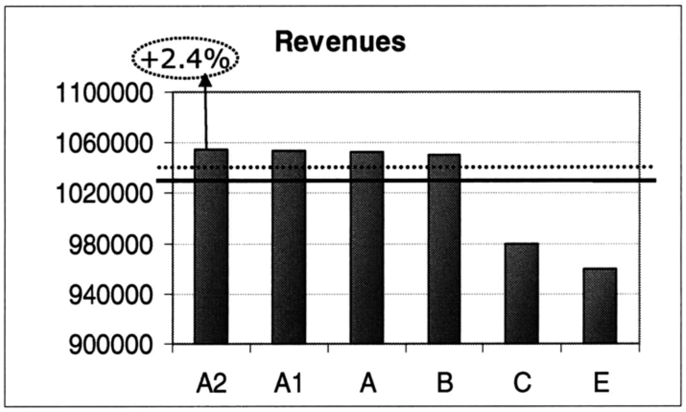

Figure 61: DP/HF1 (IAP0) vs DP/HF1 (IAP0)

4.3.4. Summary of Results in Semi-restricted Network D

- Base case: EMSRb vs EMSRb -	/ Base Case
EMSRb/HF vs EMSRb	+2.8%
EMSRb/HF vs EMSRb/HF	+1.9%
DAVN vs EMSRb	+0.7%
DAVN/HF vs EMSRb	+4.5%
DAVN/HF vs DAVN/HF	+2.5%
DP/HF vs EMSRb	+3.3%
DP/HF vs DP/HF	+2.4%

Table 9: Results obtained in semi-restricted network D

Table 9 presents a summary of the results we obtained in semi-restricted Network D. Those results underline what we had been foreseeing in unrestricted Network D: DP-LB get results that are a little better but very similar to those obtained with EMSRb. In PODS the booking limits of EMSRb are updated 16 times and this number of iterations is sufficient to make the RM method become adaptive to realized demand when this demand differs from forecasts. If the booking limits of EMSRb were not updated at all during the reservation period then we can expect EMSRb may get worse results. Yet thanks to those updates EMSRb adapts in each time-frame to the future arrival pattern of demand among fare classes. So the theoretical advantage of DP-LB does not lead in experimental results to important difference in revenues. As detailed for unrestricted fare structures most underlying assumptions of DP-LB and EMSRb are identical. Both methods assume independence of demand between different fare classes, the demand being supposed to be segmented. They use partitioned Q-Forecasts and close a fare class when forecasts show that it would be more profitable to try to sell the next seat in one of the higher fare classes than in the current lowest open one. Consequently it is not surprising that both methods get results that are so close. As DAVN is an O-D-based optimizer it manages to get better results than DP-LB.

5. Conclusion

5.1. Summary of Findings

The objective of this thesis was to study the performance of methods based on dynamic programming in unrestricted and simplified fare structures and to focus especially on two approaches: the Gallego–Van Ryzin approach and the Lautenbacher approach. We explained the mechanisms of each method, we used the implementation of traditional RM methods and DP-based methods in the Passenger Origin Destination Simulation to get experimental results and we made a performance analysis of the obtained results.

As explained throughout this thesis traditional RM methods are outdated in unrestricted and simplified fare structures, where demand is no longer segmented. Q-Forecasting and Hybrid Forecasting, as described in Chapter 3, were developed to adapt traditional methods, which assume independence of demand, to these less-restricted fare structures. This thesis describes the benefit gained through Q-Forecasting and Hybrid Forecasting used with traditional RM methods. We showed that the use of Q-Forecasting and Hybrid-Forecasting can prevent spiral-down and lead to revenues that are higher than those obtained with basic revenue management methods similar to those used by low-cost carriers. However, we underlined that obtained revenues may still be non-optimal and that new optimizers may be required to reach optimality.

In unrestricted fare structures the question is no longer "How many seats should we save for each fare class?" but rather "What is the lowest class that should be open in a given time-frame?". Indeed we may sell a seat in a fare class only if it is the *lowest open one*. Traditional RM methods consider that passengers buy in all open fare classes whatever the lowest open fare class is, while in unrestricted fare structures they only buy in the lowest available one. Consequently it may be difficult to reach optimality with traditional RM methods in less-restricted fare structures. To get optimal revenues we need a new method to determine what the lowest open class should be at each time of the reservation period by considering potential demand that may materialize in a fare class if it is the lowest open one.

Methods based on dynamic programming consider small decision periods and corresponding forecasts rather than considering the whole period until departure and forecasts of bookings-to-come. The two DP methods we consider in this thesis are very different. The dynamic programming approach of Gallego–Van Ryzin, DP-GVR, takes into consideration the fact that passengers may sell-up or buy down between fare classes and is designed to be used in unrestricted fare structures while the approach of Lautenbacher, called in this thesis DP-LB, assumes independence of fare classes as traditional RM methods. At any time of the booking process and for any number of remaining empty seats, the Gallego–Van Ryzin approach determines which fare class should be the lowest open in order to reach optimal revenues. We will now summarize separately results obtained with both methods.

According to tests run in unrestricted fare structures, when competing against an airline using no revenue management, the choice of RM method leading to the best revenues is DP-GVR. Against more advanced RM methods this is not true anymore.

As we used input probabilities of sell-up, DP-GVR was not adapting to the competition and the method got worse revenues than other RM methods using Q-Forecasting. However, when testing the performance of RM methods against themselves the method getting the best results was DP-GVR again. We showed that this variability in the performance of DP-GVR is related to its sensitivity to forecasted probabilities of sell-up and arrival rates. This data is difficult to forecast since it depends on parameters that are difficult to predict, such as the competition's fares and seat availability in future decision periods. In our simulation the use of input probabilities of sell-up that did not reflect these dynamics lead to situations in which DP-GVR did not adapt to competition.

One of the main challenges is that observed willingness-to-pay of demand varies with competition. When the competition used no revenue management, it filled its planes with low fare demand and lacked seats at the end of the reservation period allowing the airline using DP-GVR to capture demand with high willingness-to-pay. When all airlines used DP-GVR they closed low fare classes and managed to make people sell-up in high fare classes. However, when at least one airline had low fare classes or middle fare classes open at the end of the reservation periods passengers bought down to this fare class, capturing demand with high willingness-to-pay. DP-GVR expected to capture this demand because of high input probabilities of sell-up reflecting the arrival of business passengers at the end of the booking process. Nevertheless the current implementation of DP-GVR was shown not to adapt to this situation as it is not designed to re-estimate probabilities of sell-up. In all tests run in the simulation we assumed the same shapes for probabilities of sell-up while DP-GVR seems to require adaptive forecasting of probabilities of sell-up to perform effectively.

When DP-GVR competes against carriers using advanced revenue management, we have shown in this thesis that its results are very sensitive to willingness-to-pay inputs. The revenues obtained with Q-Forecasting and traditional RM methods are not as sensitive to input probabilities of sell-up. Much work has been done on booking forecasts but we showed in this thesis that DP-GVR also requires more advanced forecast of probabilities of sell-up. The input or estimated probabilities of sell-up are used directly by the RM method. However, these are forecasts that are difficult to estimate. The implementation of DP-GVR we have simulated uses input parameters and does not take into consideration the actual bookings to forecast probabilities of sell-up. If actual passengers do not want to sell-up but input probabilities of sell-up are high, DP-GVR keeps trying to get people in the high fare classes and loses many high-fare paying passengers that buy-down and are captured by the other airlines. If people do want to sell-up but input probabilities of sell-up are low DP-GVR will not close classes and thus lose much demand for high-fare classes. In the next section, we will propose avenues for further research concerning this RM method.

Both in less restricted fare structures and unrestricted fare structures DP-LB gets revenues that are higher than those obtained with EMSRb but smaller than those obtained with DAVN (all methods using Q-Forecasting for unrestricted fare structures and Hybrid Forecasting for simplified fare structures). DP-LB gets results that are very close to those obtained with EMSRb.

Overall, the performance of DP-LB in simplified and unrestricted fare structure is limited: this method requires longer computation times than other methods assuming independence of fare classes for results that are on average just slightly better than those obtained with EMSRb and worse than those obtained with DAVN. DP-GVR appears to be more promising but the experimental results we obtained in this thesis do not prove whether this method has a real potential against advanced competitors in less-restricted fare structures. We have shown that DP-GVR is very sensitive to forecasts of probabilities of sell-up and one will have to check whether more advanced forecast of probabilities of sell-up could lead to stable advantage of DP-GVR over other RM methods.

5.2. Further Studies

We now discuss further studies, which, in our opinion, should be done in the field of revenue management. We will specifically present some thoughts on the estimation of probabilities of sell-up in unrestricted fare structures since we have shown that the next logical step for further study on DP-GVR would be testing its performance with more advanced estimation of probabilities of sell-up.

Forecasting and optimizing are usually treated sequentially while there may be unexpected feedback effects. Indeed the forecasters are fed with actual bookings obtained after using a given optimizer but this optimizer itself based its decisions on forecasts obtained with historical data. Detruncation was developed to feed the optimizer with actual demand rather than observed one. However, in less-restricted fare structures it is difficult to deduce potential demand from observed demand as people buy-down or sell-up more easily between fare classes and many parameters have to be taken into consideration to forecast the probabilities of sell-up. This may lead to situations where an airline underestimates potential demand because all this demand buys down to fare classes of competition. In less-restricted fare structures an optimal revenue management method would be a method that would converge both to more accurate forecasts of potential demand and to optimal revenues according to input forecasts. To improve actual methods the interactions between forecasting and optimizing should be studied in detail.

Improving the forecasting of probabilities of sell-up is one of the main points on which airlines will have to focus if they want to get better revenue management methods in less-restricted fare structures. This improvement may enable methods based on dynamic programming to get better results in unrestricted fare structures against advanced competitors. It would improve both the forecast of arrival rates and the forecast of potential demand by fare class. These are the two crucial forecasts on which DP optimizers base their decisions in order to reach an optimal policy. The results we obtained with DP-GVR when facing advanced competitors were not as good as we may have expected when studying its theoretical advantages. But in this thesis we only used input probabilities of sell-up. The next logical step for further studies is to incorporate estimation of probabilities of sell-up with DP-GVR.

First we should decide which probabilities of sell-up we want to forecast. In this thesis we considered probabilities of sell-up between the Q class of an airline and any other class of the same airline. Yet we may consider probabilities of sell-up that are defined in a different way. The considered population can be the whole demand for the market at a given base fare or the average demand for the airline at a given base fare. The second case is fine if other competitors do not often change their closure policy. If they do then the average base demand for the airline may vary significantly from one departure to another, and then we need to take into consideration for each departure the total demand for the market and the current policy of the competitors.

Before thinking about estimating probabilities of sell-up some thinking should be done on all parameters that may influence demand on a particular path in a given time-frame. We should try to formulate demand in time-frame t for departure i with as few assumptions as possible. For example here is one attempt:

$d(F,i,t)=D_i*pi(t)*g(F,f_1,...,f_n,t)$

Where:

F: current fare of the airline

d(F,i,t): demand in time-frame t for departure i

D_i: total demand for departure i

$p_i(t)$: effect on total demand of closure policies used on other paths (served by the airline and its competitors for the same O-D market) before t

$f_1,...,f_n$: current fares available on other paths

$D_i*p_i(t)$ is the total demand for the O-D market in time-frame t (it depends on the past closure policy of all airlines for the considered departure). D_i may be assumed to be constant for all departures. There may be other parameters influencing demand in a given time-frame yet we just want here to underline that forecasters make many assumptions, which may not have biased forecasts in fully-restricted fare structures but which need to be reconsidered in unrestricted fare structures. Indeed in fully-restricted fare structures demand in time-frame is often assumed to be $D_i*a(t)*g(F)$ where a is a function which value depends only on the time-frame and g now depends only on F, the current fare of the airline.

Notice the importance of $p_i(t)$ and $f_1,...,f_n$ in the proposed formula in unrestricted fare structures: to estimate the probabilities of sell-up in a given time-frame an airline should theoretically consider the closure state of the various fare classes it offers on the same O-D market but on other paths and the closure state of the fare classes on all competitor paths. Today many passengers are not sensitive to the airline they are traveling with. They are just looking for the cheapest ticket they can find to travel from a given origin to a given destination at a specific time. So to perform better, RM methods should take into consideration all those market conditions. However, it is difficult to incorporate the impact of the competitors on the demand that an airline may be able to capture. Probabilities of sell-up depend on all other available fares offered in the same O-D market, on the time and number of connections of the various fares, and on convenience or restrictions that may be associated with them. All those parameters may be taken into consideration when estimating probabilities of sell-up.

We can not take into account all those parameters when estimating the probabilities of sell-up as historical data is sparse and we have to make assumptions in order to determine a formula that depends on a small number of parameters, as few as possible. For example we have to decide whether or not we can make the assumption that some parameters are constant across time-frames. Parameters should be such that by adjusting them the formula can fit historical data while at the same time converging as fast as possible to fit real unobserved probabilities of sell-up (for fare classes that have not yet been set as the lowest open fare class in the past).

How can we get information on probabilities of sell-up? *Figure 62* shows historical data obtained for 8 time-frames and 6 fare classes. We notice than in some time-frames, for example time-frame 4, we have no information on sell-up as it was always the same fare class that was open in the past. We can not make assumption on sell-up by looking at timeframe 4 only. In time frame 6 we had two different fare classes that were open in the past. We then get some information on sell-up between fare class 4 and fare class3. In time-frame 5 we had 3 different fare classes. This time-frame gives much more information on sell-up and we may more accurately forecast potential demand for other fare classes. To get information on sell-up we may also compare information that comes just before and just after the closure of a fare class for a given departure.

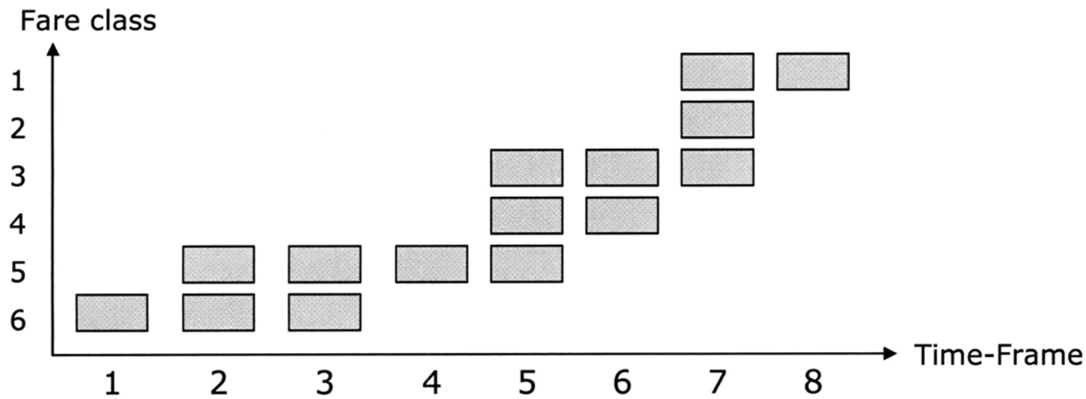

Figure 62: Historical data

DP-GVR was developed to be used in unrestricted environments. Yet in fully-restricted environments some people may sell-up as well, so the choice of the lowest open fare classes is important too. Traditional methods tend to keep the old closure pattern of fare classes as the input forecasts are obtained in cases when this pattern was used and there is a feedback effect. DP-GVR does not make assumptions on the shape of the optimal closure patterns and it looks for the one that leads to optimal results according to estimated arrival rates and probabilities of sell-up.

If this optimal pattern is very different for various departures DP-GVR is supposed to be more adaptive defining optimal policies depending on the number of remaining seats at a given point in time. If the arrival pattern does not change for different departures, if competitors do not change their closure patterns from one departure to another and if we know that the closure pattern we have used in the past is quite optimal then we do not really need to account for the possibility of sell-up and we may keep on using it. But when the arrival pattern changes from one departure to another or when the closure pattern we have been following is not at all optimal then the effect of sell-up may be important. By using traditional forecasting we can only average the effects of sell-up while DP-GVR can take into consideration forecasted potential demand.

How can we adapt DP-GVR to fare structures with restrictions? We first have to adapt the way to forecast potential demand for each fare class. Traditional forecasts should be applied to all fare classes that are not the lowest open one. This would give us the minimum number of people buying in a fare class when it is not the lowest open one. On top of that we should account for the fact that people buying in the lowest class may sell-up/buy-down. The effect of price-oriented demand is the difference between the average number of bookings when i was the lowest open fare class and the average number of bookings when it was higher than the lowest open fare class. (or 0 if negative). The problem is to estimate what the additional demand for each fare class is if we set this fare class as the lowest open one.

Finally we will have to find a way to take into account the network effects with methods based on dynamic programming. The DP methods studied in this thesis where setting policies by leg while we may get better results if we manage to consider network effects and set policies for O-D paths. Airlines may also choose to apply methods based on dynamic programming to only part of the network or as a mean to find new optimal patterns of fare class closure on given flight legs or O-D markets.

From an experimental point of view, it may be interesting to study what happens when airlines use different fare structures because when running tests for this thesis we always used the same fare structures for all airlines. In reality however, airlines often use the same fare structure.

As a conclusion, the approaches based on dynamic programming, similar to DP-GVR, appear to be theoretically very promising for airlines looking for new revenue management methods in unrestricted fare structures. However, we will still need to demonstrate experimentally whether through some improvements they can get good results in any situation, by adapting to competition. The results of our simulations show that the use of DP-GVR can lead to reasonable fare class closure but that those policies are very sensitive to the probabilities of sell-up used by the optimizer. DP-GVR obtained worse results than other some RM methods against competitors using advanced RM methods.

The analysis of the corresponding results show that the airline using DP-GVR lost revenues because of the use of inadequate probabilities of sell-up (and consequently arrival rates). This was not a problem for other traditional RM methods using Q-Forecasting, the forecasting method developed to forecast potential demand in unrestricted fare structures. Advanced forecasting would require considering parameters which are difficult to predict, such as competition's fares and availability in future decision periods.

We can expect that the ability to more accurately estimate and manage the willingness to sell-up of demand would enable carriers to get higher revenues. This requires new ways to forecast demand but apparently also new optimizer methods in order to reach optimality.