

A Competitive Approach to Airline Revenue Management

by

Olivier d’Huart

Ingénieur des Arts et Manufactures
Ecole Centrale Paris, 2009

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 2010

©2010 Massachusetts Institute of Technology.
All rights reserved.

Signature of Author: _____

Department of Civil and Environmental Engineering
May 10th, 2010

Certified by: _____

Peter Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Accepted by: _____

Daniele Veneziano
Chairman, Departmental Committee for Graduate Students

A Competitive Approach to Airline Revenue Management

by

Olivier d'Huart

Submitted on May 10th 2010 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

ABSTRACT

Since the 1980s, the airline industry has seen two major changes: Deregulation, which led to an increase in competition, and the development of revenue management systems. Paradoxically the revenue management models used have not incorporated many competitive considerations. In this thesis we study the interactions between existing airline revenue management systems in competitive markets. We use the Passenger Origin Destination Simulator (PODS) for simulations.

After a review of the past research on such interactions, we develop our own model and use simulations to identify and measure the extent to which revenue management systems of competing airlines affect each other. The model introduced highlights the importance of spill-of-demand between airlines. We show that with current revenue management practice, a legacy carrier should be less sensitive than a low-cost carrier to revenue management competitive interactions. As compared to an equivalent monopoly, an airline oligopoly tends to allocate more seats to high-fare passengers and fewer seats to low-fare passengers. With steady demand distributions, an airline's expected revenues are a decreasing function of the seat capacity allocated by its competitors to high-fare passengers. Existing revenue management systems react to competitor moves automatically only if a change in the seat allocation rule by an airline occurs over a large enough number of successive departures to be detected by forecasters.

We then suggest how to improve revenue management based on the interactions identified. With steady demands, if a competitor increases (respectively decreases) its seat allocation for high-fare passengers, the best response to optimize revenues on the short-term is to decrease (respectively increase) seat allocation for high-fare passengers. We also show that the use of EMSRb-optimization by competitors results in a near-optimum competitive equilibrium, and a near-optimum cooperative equilibrium if airlines do not share revenues. Rarely can competitive interactions justify an airline to override the EMSRb seat-allocation rule to optimize revenues. Last, we introduce LOCO-based Forecast Multiplication, a heuristic forecast adjustment made in response to the current seat availability of the competitors that can increase an airline's revenues substantially.

Key Words: *Airline Revenue Management, Competition*

Thesis Supervisor: Dr. Peter Belobaba

Title: Principal Research Scientist of Aeronautics and Astronautics, MIT

ACKNOWLEDGEMENTS

First I would naturally like to thank my thesis and research advisor Dr Peter Belobaba. He is a great teacher, and knows how to share his enthusiasm for his work. His knowledge and experience of airline revenue management have been praised by airlines all around the world, and have been very valuable to me from my first day at MIT up to his many reviews of this thesis. I consider myself very fortunate to have worked with him.

I would also like to thank the airline members of the PODS Consortium for sponsoring my research. It has always been a pleasure to share my ideas with them at the PODS meetings worldwide. Special thanks go to Craig Hopperstad, the programmer of the PODS Simulator, without whom I could not have done the simulation part of my thesis. Also special thanks to my fellow PODS colleagues: Claire for her incredible availability that led to so many endless discussions, Chris for the great trip in Patagonia, Sarvee and Himanshu for answering all my questions.

I am grateful to Ecole Centrale Paris for offering the opportunity to pursue a simultaneous double degree at American Universities, and especially at MIT. The Institute is such a unique and vibrant place to learn about our immensely complex world and to develop oneself. These last two years have with no hesitation been the best of my life (until now!). I have tremendously matured here, and MIT has perfectly equipped me to go forward in the future. I understood the real meaning and challenges of globalization, and I am confident that the Institute will remain as much committed to its mission of generating, disseminating, and preserving knowledge to bear on the world's great challenges and for the betterment of mankind.

I would like to thank all the people that made my stay here unforgettable and exciting. Particular thanks to my fellow ICAT students for having endured my presence in the lab. For the rest, I will just thank more particularly Renée, Sarah, FX, Stéphane, Thomas, the French, the Greek, the Lebanese, the Italian, the MIT Rugby Team, and the 3H.

I also would like to thank my parents for having given their best for twenty years to permit me to become what I am.

Last but not least, I wish all the best to “Le Crou”, also known as la Dude, la Dune, la Mike et la Sime. May we remain the same, pursue our adventures and live our dreams. « Quand on est jeune, c’est pour la vie. » -*Georges Clémenceau* -

TABLE OF CONTENTS

ABSTRACT	3
ACKNOWLEDGEMENTS	5
TABLE OF CONTENTS	7
LIST OF FIGURES	9
CHAPTER 1: INTRODUCTION	15
1.1 Background Introduction to Airline Revenue Management	15
1.2 Airline Revenue Management Lacks a Competitive Approach	21
1.3 Objectives of the Thesis	23
1.4 Structure of the Thesis	23
1.5 Experimental Tool: PODS Airline Revenue Management Simulator.....	23
CHAPTER 2: LITERATURE REVIEW	29
2.1 Former PODS Studies of Competitive Environments	29
2.1.1 Wilson’s Results on Competitive Impacts of Airline Revenue Management.....	29
2.1.2 Other PODS Simulations in Competitive Environments	30
2.1.3 Conclusions	33
2.2 Theoretical Approaches.....	34
2.2.1 Models with Independent Demands for Fare Classes	34
2.2.2 Models with Dependent Demands for Fare Classes	39
2.2.3 Conclusions	40

CHAPTER 3: IMPACT OF COMPETITION ON CURRENT AIRLINE REVENUE MANAGEMENT SYSTEMS	41
3.1 Competition vs. Monopoly	41
3.1.1 Preliminary Theoretical Results	41
3.1.2 Experimental Methodology	56
3.1.3 Simulation Results	59
3.1.4 Conclusions	76
3.2 Competitive Interactions with Current Airline Revenue Management ..	77
3.2.1 Forecasting	77
3.2.2 Optimization of the Booking Limits	78
3.2.3 Conclusions	88
 CHAPTER 4: PROPOSED COMPETITIVE APPROACHES	 89
4.1 LOCO-based Forecast Multiplication	89
4.1.1 Concept	89
4.1.2 Experimental Methodology	96
4.1.3 Simulation Results	98
4.1.4 Conclusions	122
4.2 Deliberate Competitive Overrides	123
4.2.1 Deliberate active reactions to competitor moves	123
4.2.2 Deliberate initial competitive moves	127
4.2.3 Conclusions	130
 CHAPTER 5: CONCLUSION	 131
5.1 Research Findings and Contribution	131
5.2 Future Research Directions	136
 REFERENCES / BIBLIOGRAPHY	 137

LIST OF FIGURES

CHAPTER 1

Fig 1.1.1	Airline fare product differentiation	16
Fig 1.1.2	A simplified example of an airline fully-restricted fare structure	17
Fig 1.1.3	A simplified example of an airline semi-restricted fare structure	17
Fig 1.1.4	A simplified example of an airline unrestricted fare structure	17
Fig 1.1.5	Principle of nested fare class booking limits and availabilities.....	18
Fig 1.1.6	Airline Revenue Management Process Flow. Source: Talluri and van Ryzin [33]	20
Fig 1.2.1	Traditional airline revenue management model	22
Fig 1.2.2	More realistic model of two competing airlines.....	22
Fig 1.5.1	PODS simulator structure as used in the framework of this thesis	24
Fig 1.5.2	Passenger choice module of the PODS simulator as used in this thesis	25
Fig 1.5.3	Passenger cumulative arrivals pattern	26
Fig 1.5.4	Time-of-day percentage demand for a market where difference in local time at destination upon arrival and local time at origin upon departure is 3hrs.....	26
Fig 1.5.5	Decomposition of the booking period before the flight departure	27

CHAPTER 2

Fig 2.1.1	Normal form of typical results by Wilson [37]	30
Fig 2.1.2	Normal form of results by Dar [8] in a network with an unrestricted fare structure.....	31
Fig 2.1.3	Normal form of results by Lee [22] in a network with a restricted fare structure.....	31
Fig 2.2.1	Parallel between the Newsboy Problem and Littlewood’s Revenue Management Problem.....	35
Fig 2.2.2	The model of Netessine and Shumsky [31].....	36
Fig 2.2.3	Bipartite graph of the two airline, independent classes, competitive game. Source: Talluri and van Ryzin [33]	38
Fig 2.2.4	Results of Isler and Imhof [18] simulation.....	39

CHAPTER 3

Fig 3.1.1	Decomposition of non-constrained demand for a fare class at an airline.....	42
Fig 3.1.2	Decomposition of non-constrained demand for a fare class at an airline.....	42
Fig 3.1.3	Competitive approach, duopoly situation, for a set of classes (1..i).....	43
Fig 3.1.4	Decomposition of non-constrained demand with sell-up	44
Fig 3.1.5	Competitive approach, duopoly with sell-up for a set of classes (1..i)..	45
Fig 3.1.6	Determination of a nested protection level, example of monopoly situation.....	46
Fig 3.1.6bis	Determination of a nested protection level, example of monopoly situation with passenger sell-up	47
Fig 3.1.7	Determination of a nested protection level, example of a duopoly situation where passengers do not spill-in between airlines	48
Fig 3.1.8	Determination of a nested protection level, example of a duopoly situation with passenger spill-in between airlines.....	49
Fig 3.1.8bis	Determination of a nested protection level, example of a duopoly situation with passenger sell-up within each airline and spill-in between airlines.	50

Fig 3.1.9	Monopoly scenario vs. competitive scenarios in a single market	57
Fig 3.1.10	Fully-restricted fare structure considered.....	58
Fig 3.1.11	Unrestricted fare structure considered.....	58
Fig 3.1.12	Average sum of airline forecasts, 8am flight, 24 days prior departure	59
Fig 3.1.13	Average total market booking limits, 8am flight, 24 days prior departure	60
Fig 3.1.14	Average total bookings by fare class for a day of departures.....	60
Fig 3.1.15	Average total daily market revenues, passengers and load factors	60
Fig 3.1.16	FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure...	61
Fig 3.1.17	FCOC - Average total market booking limits, 8am flight, 24 days prior departure.....	62
Fig 3.1.18	FCOC - Average total bookings by fare class for a day of departures.....	62
Fig 3.1.19	FCOC - Average total daily market revenues, passengers and load factors.....	62
Fig 3.1.20	Average sum of airline forecasts, 8am flight, 24 days before departure	63
Fig 3.1.21	Average total market booking limits, 8am flight, 24 days prior departure	64
Fig 3.1.22	Average total bookings by fare class for a day of departures.....	64
Fig 3.1.23	Average total daily market revenues, passengers and load factors	64
Fig 3.1.24	Average sum of airline forecasts, 8am flight, 24 days prior departure	65
Fig 3.1.25	Average total market booking limits, 8am flight, 24 days prior departure	66
Fig 3.1.26	Average total bookings by fare class for a day of departures.....	66
Fig 3.1.27	Average total daily market revenues, passengers and load factors	66
Fig 3.1.28	FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure...	67
Fig 3.1.29	FCOC - Average total market booking limits, 8am flight, 24 days prior departure.....	68
Fig 3.1.30	FCOC - Average total bookings by fare class for a day of departures.....	68
Fig 3.1.31	FCOC - Average total daily market revenues, passengers and load factors.....	68
Fig 3.1.32	Average sum of airline forecasts, 8am flight, 24 days prior departure	69
Fig 3.1.33	Average total market booking limits, 8am flight, 24 days prior departure	70
Fig 3.1.34	Average total bookings by fare class for a day of departure	70
Fig 3.1.35	Average total daily market revenues, passengers and load factors	70
Fig 3.1.36	Average sum of airline forecasts, 8am flight, 24 days prior departure	71
Fig 3.1.37	Average total market booking limits, 8am flight, 24 days prior departure	72
Fig 3.1.38	Average total bookings by fare class for a day of departures.....	72
Fig 3.1.39	Average total daily market revenues, passengers and load factors	72
Fig 3.1.40	FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure...	73
Fig 3.1.41	FCOC - Average total market booking limits, 8am flight, 24 days prior departure.....	73
Fig 3.1.42	FCOC - Average total daily market revenues, passengers and load factors.....	73
Fig 3.1.43	Average sum of airline forecasts, 8am flight, 24 days prior departure	74
Fig 3.1.44	Average total market booking limits, 8am flight, 24 days prior departure	74
Fig 3.1.45	Average total bookings by fare class for a day of departures.....	75
Fig 3.1.46	Average total daily market revenues, passengers and load factors	75
Fig 3.2.1	Passive competitive reaction of Airline A's EMSRb revenue management system when Airline B decreases its booking limit on the long term	82
Fig 3.2.2	Passive competitive reaction of Airline A's EMSRb revenue management system when Airline B increases its booking limit on the long term	82
Fig 3.2.3	Revenue implications of the passive interactions of revenue management systems when an airline implements a better form of revenue management	83
Fig 3.2.4	Graphical representation of the case of two competing airlines using EMSRb..	84
Fig 3.2.5	Competitive interaction, Airline B decreases a booking limit on the short-term..	86
Fig 3.2.6	Competitive interaction, Airline B increases a booking limit on the short-term..	86

CHAPTER 4

Fig 4.1.1	Single market duopoly scenario studied.....	97
Fig 4.1.2	Unrestricted fare structure considered.....	97
Fig 4.1.3	Fully-restricted fare structure considered.....	97
Fig 4.1.4	Initial thresholds set by Airline 2	99
Fig 4.1.5	Average proportional change in revenues, load factor, and yield for each airline when Airline 1 uses FM.....	100
Fig 4.1.6	Comparison of Airline 1 average daily bookings by fare class when it is using FM1 and when it is not.....	100
Fig 4.1.7	Statistics for Airline 1 using FM1 over 200 successive days of departures.....	101
Fig 4.1.8	Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 over 200 successive departure dates.....	101
Fig 4.1.9	Distribution over 200 successive departure dates of the percentage change in revenues for Airline 2 when Airline 1 uses FM1	102
Fig 4.1.10	Distribution over 200 successive departure dates of Airline 1 base case revenues when FM1 increases revenues and when it decreases revenues	103
Fig 4.1.11	Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1.....	103
Fig 4.1.12	Examples of values of FM1 over sample flights at different departure dates	104
Fig 4.1.13	Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates.....	106
Fig 4.1.14	Average proportional change in revenues, load factor, and yield for each airline when Airline 1 uses FM.....	107
Fig 4.1.15	Comparison of Airline 1 average bookings by fare class for its 3 daily flights when it is using FM1 and when it is not.....	108
Fig 4.1.16	Statistics for Airline 1 using FM1 over 200 successive days of departures.....	108
Fig 4.1.17	Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 over 200 successive departure dates.....	109
Fig 4.1.18	Distribution over 200 successive departure dates of the percentage change in revenues for Airline 2 when Airline 1 uses FM1	109
Fig 4.1.19	Distribution over 200 successive departure dates of Airline 1 base case revenues when FM1 increases revenues and when it decreases revenues	110
Fig 4.1.20	Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1.....	110
Fig 4.1.21	Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates.....	113
Fig 4.1.22	Average proportional change in revenues, load factor, and yield when both airlines use FM.....	115
Fig 4.1.23	Comparison of Airline 1 average bookings by fare class when both airlines are using FM1 and when they are not	115
Fig 4.1.24	Comparison of Airline 2 average bookings by fare class when both airlines are using FM1 and when they are not	115
Fig 4.1.25	Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 and Airline 2 over 200 successive departure dates.....	116
Fig 4.1.26	Distribution of the percentage increase in total market revenues that FM1 would generate over 200 successive departure dates	116
Fig 4.1.27	Distribution over 200 successive departure dates of Airline 1 base case revenues when FM1 increases or decreases Airline 1 revenues.....	117

Fig 4.1.28	Distribution over 200 successive departure dates of Airline 2 base case revenues when FM1 increases or decreases Airline 2 revenues.....	117
Fig 4.1.29	Distribution over 200 successive departure dates of base case total market revenues when FM1 increases or decreases total market revenues.....	118
Fig 4.1.30	Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1 by both airlines.....	118
Fig 4.1.31	Distribution over 200 successive departure dates of Airline 2 revenues without and with the use of FM1 by both airlines.....	119
Fig 4.1.32	Distribution over 200 successive departure dates of total market revenues without and with the use of FM1 by both airlines.....	119
Fig 4.1.33	Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates.....	121
Fig 4.2.1	Decision diagram. RM response to a competitor unusually open	124
Fig 4.2.2	Decision diagram. RM response to a competitor unusually closed.....	126
Fig 4.2.3	Impacts of an override from the EMSRb rule on an airline's revenues.	128
Fig 4.2.4	Decision diagram. Override from an EMSRb near-optimum protection level.....	129

CHAPTER 1.

INTRODUCTION

In this chapter, we first introduce the background field of this thesis: airline revenue management. We then define the specific problem studied by this thesis within this field: the reciprocal interactions between the revenue management practices of competing airlines. These interactions have been mostly overlooked until now, but could have a non-negligible effect on the performance of revenue management systems. After delimiting this problem, we define the objectives of the thesis, and the thesis structure that has been chosen to best communicate on the results of our research. We close this chapter by introducing the experimental tool used throughout the thesis, the Passenger Origin Destination Simulator (PODS), a software simulator of airline revenue management.

1.1 Background Introduction to Airline Revenue Management

The work of this thesis is set in the general field of **airline revenue management**. Also known as **airline yield management**, it started developing in the 1980s, following US airline deregulation. Since then, instead of merely selling the right to travel on a seat of the aircraft, airlines have been able to offer different products to their passengers, by associating the simple right to be on an aircraft with purchase restrictions such as a refundability of the ticket, a Saturday night stay at destination, the number of days in advance the ticket has to be bought etc. A set of purchasing restrictions, associated with a price, define what is called a “**fare product**”. Offering a variety of fare products, each with their own price and level of service, allows airlines to extract more revenues from passengers, through a process of demand differentiation. This process cannot be considered strictly as price discrimination as each fare product corresponds to a different service. The objective of airline revenue management is to sell the right fare product to the right passenger at the right time, in order to optimize total revenues.

Airline revenue management generally assumes that the marginal cost of boarding one additional passenger can be neglected compared to the incremental revenue associated with the sale of a ticket. Therefore even if it merely focuses on the revenue side of the airline profit equation, the practice of revenue management is often viewed as increasing profit through increasing revenues with costs remaining constant.

Fig 1.1.1 shows how airlines can extract more revenues from such fare product differentiation instead of selling only one product at a given price P with its demand Q . Total revenues for the airline are represented as the surface of the rectangles under the demand curve. Ideally, an airline would try to get the entire area below the curve by charging each individual passenger his willingness to pay. Using more than one fare product better covers demand and extracts more revenues.

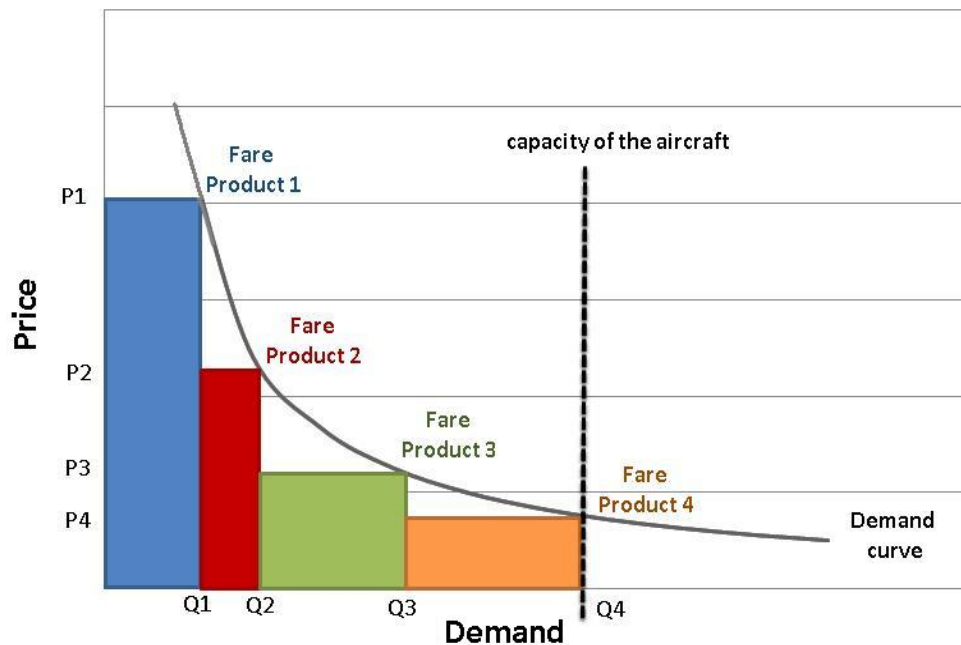


Fig 1.1.1 Airline fare product differentiation

In order to optimize revenues, airline distribution, inventory control and reservation systems manipulate “fare classes” or “booking classes”. A fare class groups one or more similar fare products. In this thesis we designate the cheap restricted classes intended for the flexible leisure passengers as “low fare classes” and the expensive fare classes for the inflexible business passengers as “high fare classes”. For each airline origin-and-destination market, marketing departments design fare products and fare classes which together define a “fare structure”.

In the beginning, and for effective revenue management, airlines aimed at segmenting demand into independent distinct subsets each looking for a specific fare product. They implemented fare structures where no fare classes have the same restrictions, where each fare class is characterized by its unique set of restrictions. Such fare structures are called **fully-restricted fare structures**. Later low-cost carriers adopted in their business model fare structures where all the classes have the same set of loose restrictions and differ only by their fare. Called **unrestricted fare structure**, such simplified fare structures are more popular among customers and thus allow low-cost carriers to encroach significantly on the market shares of incumbent legacy carriers. Yet unrestricted fare structures decrease the efficiency of revenue management systems as they make segmentation much more difficult. In order to keep their market shares and at the expense of short term revenues, legacy carriers have sometimes matched the move made by low-cost carriers towards simplified fare structures. They have removed some restrictions, in particular the Saturday-night-stay, so that some classes ended up with the same set of restrictions. Such fare structures are called **semi-restricted fare structures**.

Fig 1.1.2 shows an example of a fully-restricted airline fare structure. The Saturday night stay restriction is supposed to prevent business travellers with high willingness-to-pay from purchasing low-fare classes. Similarly, no possibility to change flights or obtain a refund and the necessity to purchase the ticket a long time in advance prevents non flexible travellers with a high willingness-to-pay from purchasing the cheapest fares. Fig 1.1.3 shows an example of semi-restricted fare structure and Fig 1.1.4 an example of unrestricted fare structure.

Fare Class	Fare	Advance purchase requirement	Saturday night stay required at destination	Change fee	Refundable
Y	\$600	0 days	no	no	yes
B	\$500	7 days	no	no	no
M	\$300	14 days	yes	no	no
Q	\$200	21 days	yes	yes	no

Fig 1.1.2 A simplified example of an airline fully-restricted fare structure

Fare Class	Fare	Advance purchase requirement	Saturday night stay required at destination	Change fee	Refundable
Y	\$600	0 days	no	no	yes
B	\$500	7 days	no	no	no
M	\$300	14 days	no	yes	no
Q	\$200	21 days	no	yes	no

Fig 1.1.3 A simplified example of an airline semi-restricted fare structure

Fare Class	Fare	Advance purchase requirement	Saturday night stay required at destination	Change fee	Refundable
Y	\$600	0 days	no	yes	no
B	\$500	0 days	no	yes	no
M	\$300	0 days	no	yes	no
Q	\$200	0 days	no	yes	no

Fig 1.1.4 A simplified example of an airline unrestricted fare structure

Airline revenue management has combined the practices of forecasting, overbooking, fare-class mix control and origin-destination network control to determine how many seats to protect for each fare class from the fare structure at any given time over the booking period preceding a flight.

The objective of **forecasting** is to estimate future quantities such as passenger demand for fare classes or passenger price sensitivity based solely on observations from the past, which are essentially bookings on previous flights. These forecasted quantities are then used for the optimization processes of overbooking, fare class mix control and origin-destination network control.

The objective of **overbooking** is to determine the total number of bookings to accept for a flight, larger than the actual capacity, to make up for cancellations and no-shows from passengers. Overbooking algorithms balance the loss of revenue associated with an empty seat taking-off with the cost of denying boarding to a passenger.

The objective of **fare class mix control** is to determine how many seats to protect for each fare class on an aircraft. Fare class mix control aims at saving seats for the most expensive fare classes designed for late-booking business passengers, and make them available at the end of the booking period shortly before departure. The number of seats an airline chooses to save for specific classes is referred to as a **protection level**. In practice this is achieved through limiting the availability of seats in low fare-classes for passengers with a low willingness-to-pay (typically leisure passengers) who tend to be able to book earlier before flight departure. An airline sets **booking limits** for these classes. One speaks of **nesting** protection levels when protection levels overlap for successive hierarchically imbricated sets of classes rather than for separate individual classes. A **nested protection level** is the minimum seat capacity that an airline wants to protect for a given class and its higher classes, rather than just for the given class. A **nested booking limit** is the maximum number of seats that the airline is willing to sell to a given class and its lower classes, rather than just for a given class. Once the total number of bookings for a fare class and its lower fare classes has reached the nested booking limit, the fare class is not available anymore, and passengers can only purchase higher fare classes. The practice of nesting protection levels is more robust and is widespread in the industry. With nested protection levels/booking limits, a high class cannot close before a lower class does. Making fare classes **available** is also known as **opening or closing the fare classes**. Fig 1.1.5 shows an example of nested booking limits, current bookings, and fare class availabilities for a flight with a capacity of 100 with no overbooking. If there was overbooking, actual aircraft capacity of 100 would be replaced by a larger virtual capacity in the table below.

Fare Class	Fare	Nested booking limit (aircraft capacity 100)	Current bookings	Fare class available/open?
Y	\$600	100	1	yes
B	\$500	90	5	yes
M	\$300	65	25	no
Q	\$200	40	40	no

Fig 1.1.5 Principle of nested fare class booking limits and availabilities

Origin-destination network control is a more recent significant development of airline revenue management. Fare class structures in a network are defined for each specific origin-destination market. Airlines can make seats available to different fare classes based on individual passengers' revenue contribution at the level of the airline's entire network rather than at the level of a single flight leg. Origin-destination network control balances the availability (the booking limits) between local and connecting passengers on each flight leg. In simple terms, it chooses between 3 options: blocking two local seats on two different legs to protect for an eventual connecting booking, protecting the local seats from a connecting booking by blocking the latter, or making both the connecting and the local bookings possible. Such network level optimization has been implemented by the major network airlines most advanced in the field of revenue management.

In this thesis we focus on the problem of competition for the fare class mix control. More detailed information about the historical development of airline revenue management, as well as a comprehensive description of revenue management economic justification, practice and algorithms can be found in Belobaba[4], McGill and van Ryzin [30] or Talluri and van Ryzin [33].

Figure 1.1.6 is extracted from Talluri and van Ryzin [33] and represents the structure of a typical airline revenue management process flow. On the sales side, the different distribution channels (call-center, global distribution system, web server etc.) are linked to a reservation system which centralizes the information on the availabilities of fare classes. The booking limits and thus the availabilities of fare classes are decided by the revenue management system, composed of a forecasting module and an optimization module. This revenue management system needs to know what the current fare structures are (that is product specifications and prices) and uses information extracted from historical booking databases. The revenue management system is operated by **revenue management analysts** who can **override** the automatic optimization/forecasting and choose booking limits manually to handle unexpected situations or implement specific strategies.

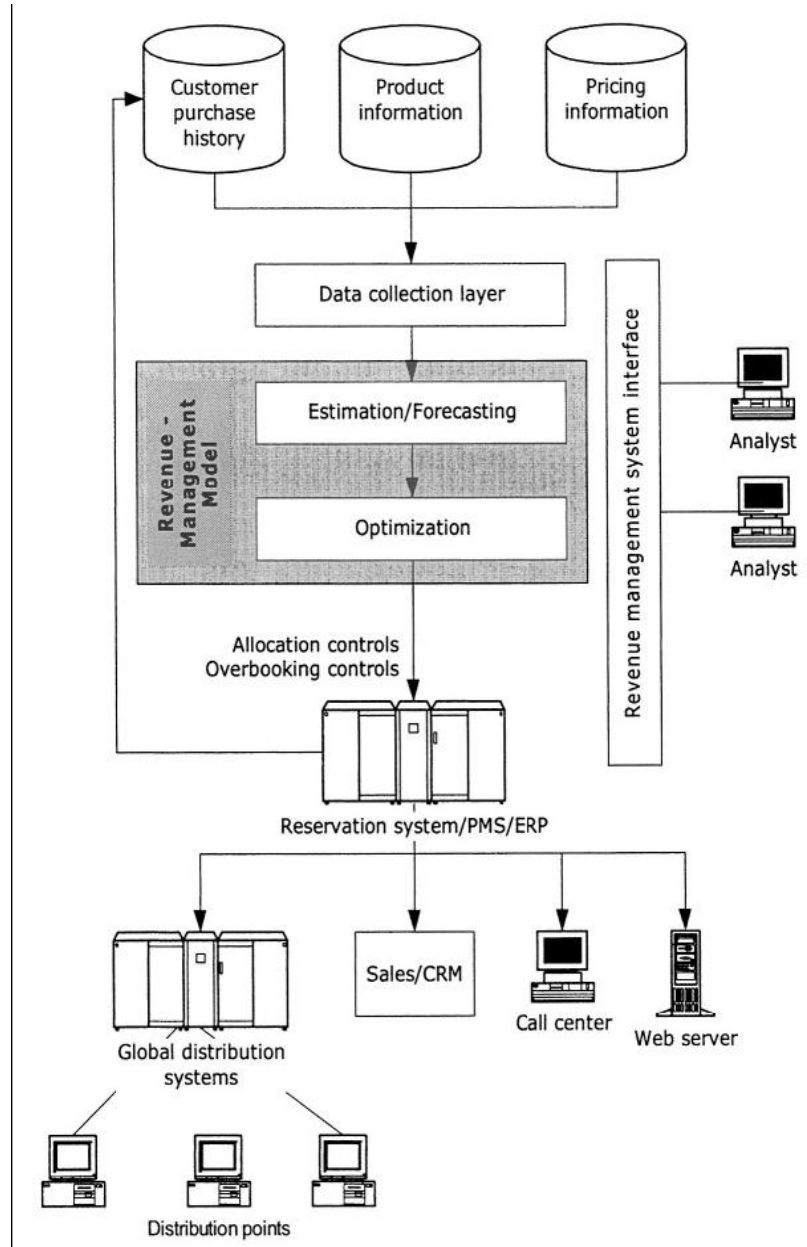


Fig 1.1.6 Airline Revenue Management Process Flow. Source: Talluri and van Ryzin [33]

1.2 Airline Revenue Management Lacks a Competitive Approach

Within the field of airline revenue management, this thesis tackles a specific issue: the reciprocal interactions between the revenue management practices of competing airlines. Competition has been an understudied topic in the field of stochastic inventory theory, the field to which airline revenue management belongs. Current airline revenue management systems consider the airline revenue optimization problem as isolated from its competitive environment. Both research and practice have assumed that agents consider most parameters such as price or demand as exogenous parameters. But in the real oligopolistic world of airline competition, these parameters are endogenous. They are determined by the simultaneous action of all agents. The seat allocation decisions of an airline affect the passenger demands for seats on other airlines.

An airline seat inventory optimizer determines fare class booking limits setting marginal revenues equal to marginal costs, where marginal cost is the opportunity cost of reserving a seat for a high fare class rather than selling it in a lower fare class. The first limit of this approach is that it uses a biased forecast. As of today, an airline forecasts its demand using solely its own past and current booking data, and extrapolates it without considering that these past demands as well as its demand-to-come are strongly dependent on competitor moves. The second limit of this approach is that it can be considered as only “first order optimization”. The allocation of seat inventory among fare classes by one airline affects the demand and optimal seat allocation of the other airlines. It might be useful for an airline to understand and evaluate how its own booking limits affect its competitors, and how in turn the competitor’s behaviour will have an effect on the initial airline’s demand.

Therefore, current decisions made by competing airlines might be suboptimal as airlines apply a monopolistic model to a competitive environment. The following two schemes Fig.1.2.1 and Fig.1.2.2 illustrate our point that traditional airline revenue management lacks a competitive approach. Figure 1.2.1 represents the model used by traditional revenue management systems. Demand is assumed by the optimization and forecasting modules to be an externality independent of competitor moves. Forecasting of the future demand is made using only the knowledge of past booking limits and past bookings. Bookings for a fare class at each time frame correspond to the minimum of the booking limit and the demand for that fare class. Figure 1.2.2 shows what could be a more realistic approach. By “**spill-of-demand**”, we mean demand that exceeds booking limits, and that is therefore diverted to the competitor. In such a competitive model, demand would not be an externality anymore. It would depend on competitor moves, which generate spill that affects our demand. Similarly, our own moves generate spill which affect the optimization process of the competitor. Optimization should take into account the reciprocity between the moves of the two revenue management systems, and/or the forecaster should take into account the competitive environment to evaluate future demand to come.

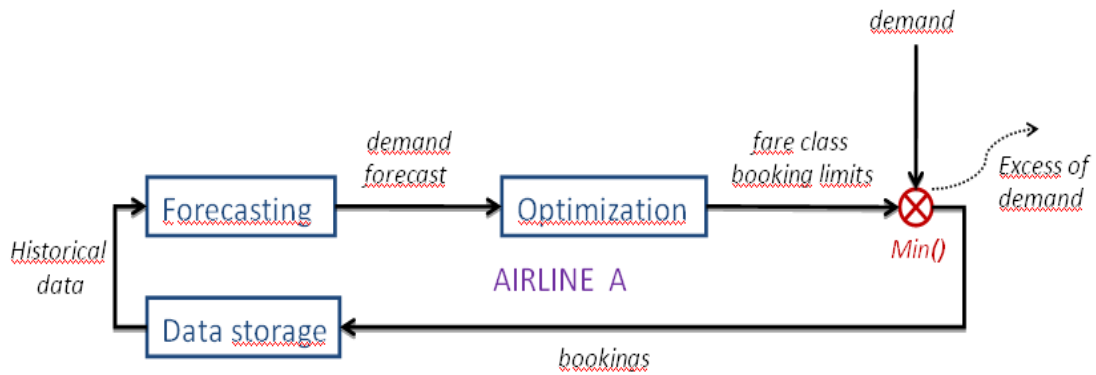


Fig 1.2.1 Traditional airline revenue management model
 $Min()$ refers to a Minimum function of the two inputs providing the output

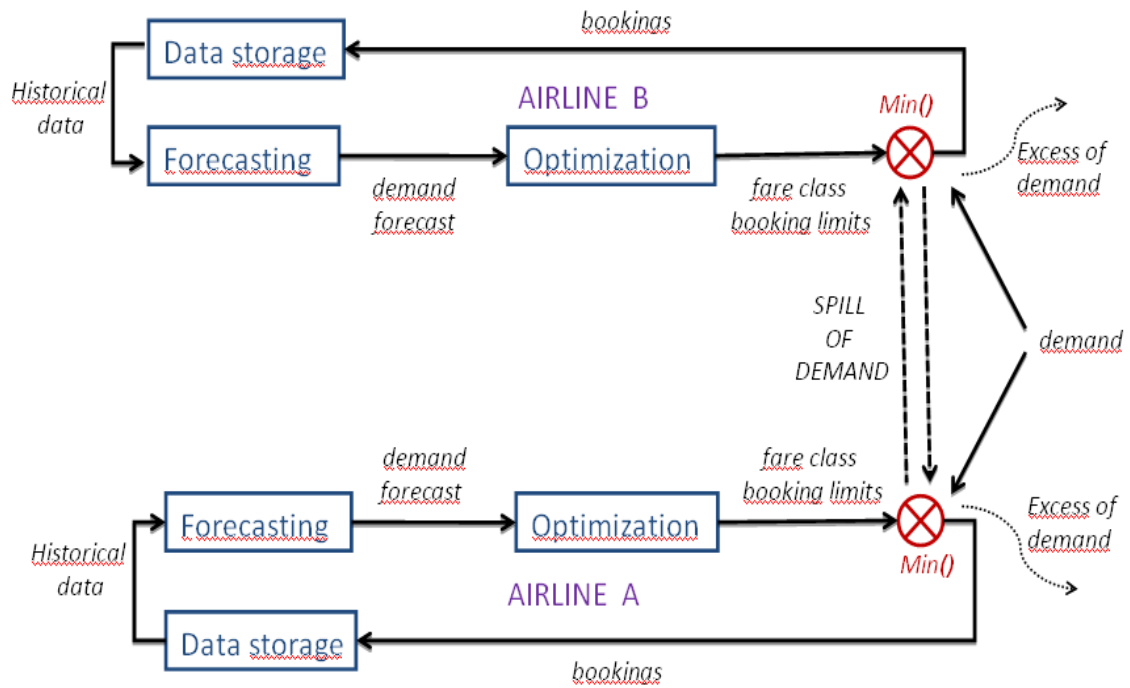


Fig 1.2.2 More realistic model of two competing airlines
 $Min()$ refers to a Minimum function of the two inputs (demand and the booking limit) providing the output (actual bookings)

1.3 Objectives of the Thesis

The first objective of this thesis is to provide a review of the state of the research on the relationship between competition and revenue management.

The second objective of this thesis is to identify how the revenue management systems of competing airlines currently interact between each other, and what the consequences of these interactions are in terms of revenues and bookings.

The third and last objective is to develop approaches of revenue management that would take into account competition as identified in the first and second part in order to improve the efficiency of revenue management systems.

Our constant concern is to provide an analysis that is directly applicable to the industry. We keep focused on **current** revenue management systems, how they are affected by competition and what changes could be made to include competitive effects. We will consider revenue management models where the price of fare classes is an exogenous variable, as these are the ones currently in use in the industry. We will not tackle joint pricing and inventory control as this topic is still at the state of research even in a monopolistic environment.

1.4 Structure of the Thesis

The structure of this thesis follows the pattern of the three objectives stated above. The first part consists of this introduction that defines the background, the problem, the objectives and the methodology of our research. The second part consists of a literature review of the past research on competitive analysis of airline revenue management, first at the macroscopic simulation level, and then at the more theoretical microscopic scale. The third part of the thesis consists of an analysis of the current interactions between competing revenue management systems to identify and evaluate them. In the last and fourth part we propose new approaches of revenue management that take into account the interactions identified in the fourth part.

1.5 Experimental Tool: PODS Airline Revenue Management Simulator

Throughout this thesis, our research methodology includes testing airline revenue management scenarios with a software simulator, the Passenger Origin-Destination Simulator (PODS). We provide here a short description of this simulator. For the reader that would like a thorough description of PODS, we strongly recommend Hopperstad Consulting[17], Belobaba, Hopperstad and Tam [3], or Lua [28]. We will only provide here a summary description of the PODS simulator which does not cover all of its modelling complexity.

PODS is a simulator whose main purpose is to test and analyze the performance of all kinds of revenue management methods. It is currently used at MIT and within a consortium of nine world-leading airlines, and we use it in this thesis for all the simulations performed. PODS aims at simulating the real world decisions of passengers and airline decision models to evaluate the latter and to assess any new paper-developed airline revenue management technique.

It is composed of two major modules: a passenger choice module and a revenue management module. Fig 1.5.1 shows the structure of the PODS simulator as used for the simulations described in this thesis.

The **passenger choice model** generates passengers with their choice specificities: an origin and a destination, a passenger type, a willingness to pay, an airline preference, a path preference, a time-of-day preference, and restriction preferences. The passenger choice module interacts with the revenue management module through a passenger choice model that aims at recreating the real world choice made by passengers when they book their tickets among fare classes made available by airline revenue management systems. Fig 1.5.2 describes the process of passenger generation and choice as used in the simulations of this thesis.

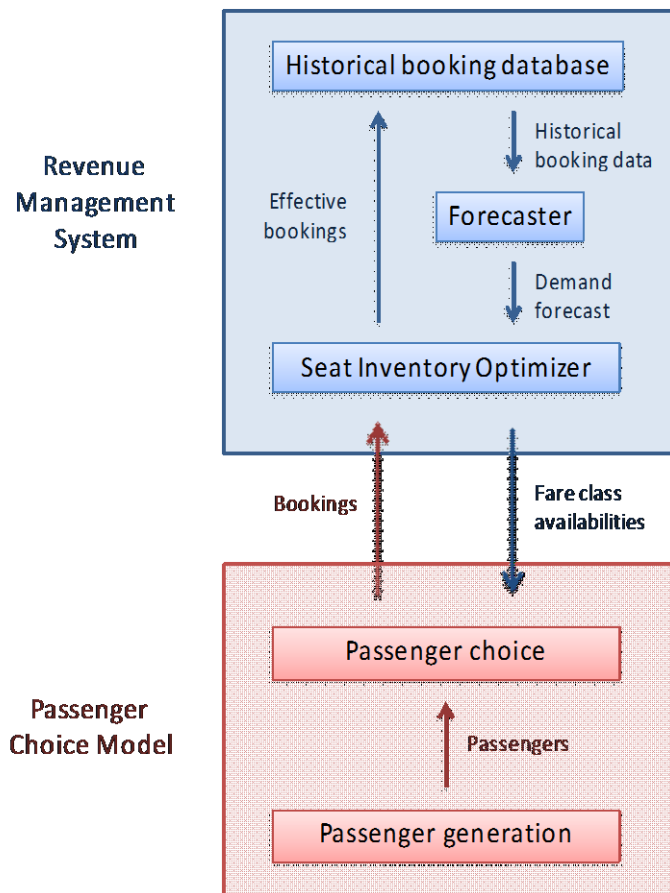


Fig 1.5.1 PODS simulator structure as used in the framework of this thesis.

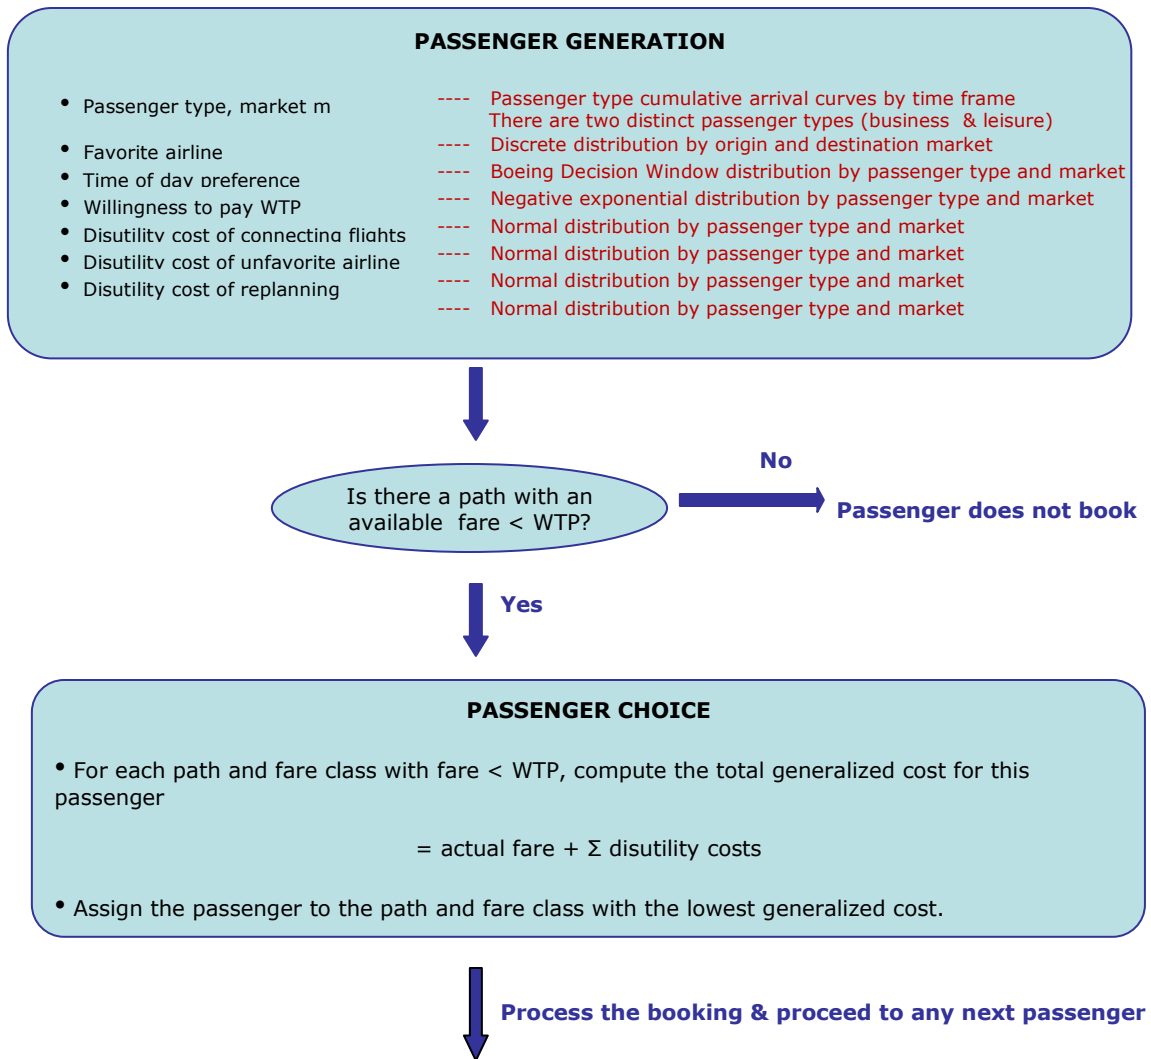


Fig 1.5.2 Passenger choice module of the PODS simulator as used in this thesis

Throughout the PODS simulations in this thesis, there are two types of passengers: leisure and business passengers. The arrival of these passengers follows the arrival pattern given in Fig 1.5.3. The time-of-day preference of these passengers follows the Boeing Time Decision Window Model [5]. A typical time-of-day demand curve is provided in Fig 1.5.4. The disutility costs used are those computed by Lee [23]. Passengers' willingness-to-pay and total market demands are then calibrated so as to reach load factors comparable to recent industry standards in North America, that is 80% to 85%.

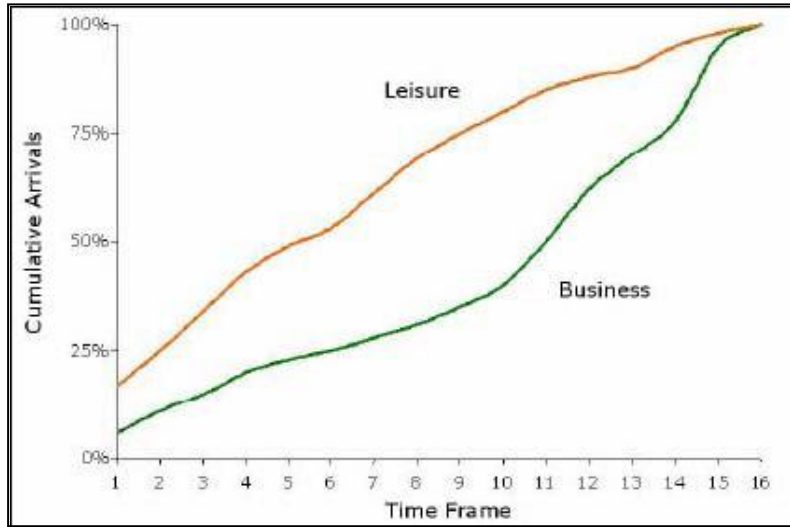


Fig 1.5.3 Passenger cumulative arrivals pattern




Fig 1.5.4 Time-of-day percentage demand for a market where the difference in local times at destination upon arrival and the local time at origin upon departure is 3 hours.

The **revenue management system** in the simulations aims at reproducing airline revenue management systems. It consists of a historical booking database, a forecaster, and a seat allocation optimizer. The seat allocation optimizer determines which fare classes to make available at each time frame, using the demand forecasts and the bookings in hand for a specific departure. Competing airlines within a network have a wide variety of choices for their own fare structures, for their forecasting techniques and seat allocation optimization methodologies. For the list of all these techniques and for their quantitative descriptions, the reader should refer to Hopperstad Consulting[17], Belobaba, Hopperstad and Tam [3] or Lua [28]. The revenue management module interacts with the passenger choice module by communicating which fare classes are available for booking when a passenger is generated.

The revenue management system decomposes the booking period before a flight departure into time frames as shown in Fig 1.5.5. At the beginning of each time frame, forecasts are updated and booking limits are re-optimized. During the time-frame passengers arrive and interact with the revenue management system. The closer in time the departure, the more frequent the revenue management system update and re-optimization.

Time Frame	Days until Departure	Time Frame Duration (days)
1	63	7
2	56	7
3	49	7
4	42	7
5	35	4
6	31	3
7	28	4
8	24	3
9	21	4
10	17	3
11	14	4
12	10	3
13	7	2
14	5	2
15	3	2
16	1	1

Fig 1.5.5 Decomposition of the booking period before the flight departure.

Throughout this thesis, a single simulation typically consists of 5 independent trials, each of which corresponding to an iterative result of 600 departure days on each leg. The first sample is initiated by user-defined inputs, which are gradually updated with new computed data for the next sample. In each trial, the first 200 samples are discarded to eliminate the effects of initial conditions since each sample has some degree of correlation to the next sample. To ensure statistical significance of simulation results, the overall result for each simulated airline is thus obtained by averaging the results of the last 400 samples for all of the 5 trials that add up to a total of 2000 daily simulations.

CHAPTER 2.

LITERATURE REVIEW

In this chapter, we provide a review of the state of the research on the relationship between competition and airline revenue management.

In the first part we introduce the results of previous PODS studies of airline revenue management scenarios in different competitive configurations. The thesis by Wilson [37] is the only PODS study to have exclusively focused on the competitive challenges of revenue management. The main purpose of the other presented studies was not to address the competitive interactions between revenue management systems, but we use their results to examine that particular aspect. We extract some general conclusions on the impact of competition for airlines implementing revenue management.

In the second part we provide a review of the more theoretical approaches of competition that have been developed in the field of airline revenue management. We only reference models where price for fare classes is a fixed exogenous variable. We do not consider inquiries of joint pricing and inventory control, as even in a non-competitive setting such approaches are still at the stage of research and not implemented at the vast majority of airlines.

2.1 Former PODS Studies of Competitive Environments

In this first part, we provide a review of previous simulations of airline revenue management in competitive environments made using the PODS simulator introduced in 1.5. PODS has been of great help in simulating the performance of airline revenue management techniques in large competitive networks and with an extensive consumer choice model. Simulations have permitted the identification of second order drawbacks/benefits of the proposed methodologies. Competitive feedbacks are among the strong identified limitations of revenue management mechanisms that are difficult to anticipate in theory.

2.1.1 Wilson's Results on Competition and Airline Revenue Management

The first known study of the impact of revenue management under competitive market conditions was made by Wilson [37] and was in fact using the PODS simulation. A short synthesis of his work was published as Belobaba and Wilson [2]. Wilson points out that the benefits of a revenue management system to an airline are a function of the revenue management capabilities of its competitors. In a single non-stop market, an airline that will innovate with its revenue management system will see its revenues increase, while its competitors' revenues will decrease. By better protecting seats for high-fare passengers, an airline using accurate revenue management will spill more low-fare passengers and less high-fare passengers to its competitors. The latter will experience a shift of their demand towards low-fare classes, and will protect less for the high-fare classes, resulting in an

overall decrease of their revenues. Still, Wilson [37] shows that competitive revenue management is not a “zero-sum game”. As soon as one competitor innovates, total market revenues increase. The revenue gains of an innovating carrier are not entirely at the expense of its competitors. And if both airlines decide to apply systematic revenue management, both carriers see their revenues increase substantially. A carrier without revenue management can catch up with others.

The results of Wilson can be displayed in the game-theoretical normal form provided in Fig 2.1.1, where payoffs correspond to a proportional increase in revenues as compared to when no airline implements revenue management. EMSRa and EMSRb refer to two versions of the leg-based EMSR fare class mix control method by Belobaba [1] described by Wilson. In this game there are two stable beneficial equilibriums (in grey) consisting in both airlines implementing revenue management. All airlines have an incentive to reach one of them. This explains why a majority of airlines throughout the world have implemented some form of revenue management.

	Airline B: no RM		Airline B: EMSRa		Airline B: EMSRb	
	<i>Airline A revenues</i>	<i>Airline B revenues</i>	<i>Airline A revenues</i>	<i>Airline B revenues</i>	<i>Airline A revenues</i>	<i>Airline B revenues</i>
Airline A: no RM	+ 0 %	+ 0 %				
Airline A: EMSRa	+ 8.2 %	- 1.4 %	+ 8.2%	+ 8.4 %		
Airline A: EMSRb	+ 10.4 %	- 1.4 %	+ 7.5%	+ 7.4 %	+ 10.3%	+ 11.0 %

Fig 2.1.1 Normal form of typical results by Wilson [37]. Payoffs are % revenue gains as compared to when no airline uses revenue management.

2.1.2 Other PODS Simulations in Competitive Environments

Although the main purpose of most of the PODS simulations mentioned below was not to study competition between revenue management systems, our objective is to regroup their results and reach conclusions on that particular point.

Most of ulterior PODS simulations have confirmed the results by Wilson [37]. Dar [8], Lee [22] and Kayser [20] allow us to extend to airline networks the conclusions made by Wilson in a single market. Improvements of revenue management systems should lead to the same kind of non-zero-sum discrete game, with a stable optimal equilibrium, and where the more players that implement good revenue management, the better the overall industry revenues. Fig 2.1.2 displays proportional increases in revenues for both competitors as compared to when both competitors use merely a leg-based fare-class mix control EMSRb

system with standard pick-up forecasting. The improvements examined are Hybrid Forecasting (HF), or origin-destination network control techniques such as Displacement Adjusted Virtual Nesting (DAVN), Probabilistic Bid-Price (ProBP) or Heuristic Bid-Price (HBP). Technical specification of these methods can be found in Dar [8]. Five improvements out of seven benefit the airline making the move forward while its competitor is hurt. The two other possible improvements benefit both airlines, but their efficiency can be questioned as they increase more the passive airline's revenues. Fig 2.1.3 is more complete and also displays proportional increases in revenues for competitors if both improve their current EMSRb system. The same pattern of results is found. If only one player makes a move, he will increase his revenues and decrease its competitor's revenues to a lesser extent. If both players make a move, both will increase their revenues. A stable equilibrium exists, and it is revenue positive for both carriers.

		Airline B : EMSRb	
		Airline A revenues	Airline B revenues
Airline A : EMSRb		+ 0 %	+ 0 %
Airline A : EMSRb + HF		+ 2.6 %	- 0.4 %
Airline A : DAVN		+ 0.7 %	+ 1.2 %
Airline A : DAVN + HF		+ 4.1 %	- 1.1 %
Airline A : ProBP		+ 0.4 %	+ 1.0 %
Airline A : ProBP + HF		+ 3.3 %	- 1.0 %
Airline A : HBP		+ 1.9 %	- 0.6 %
Airline A : HBP + HF		+ 3.7 %	- 0.7 %

Fig 2.1.2 Normal form of results by Dar [8] in a network with an unrestricted fare structure: Payoffs are % revenue gains as compared to when both airlines use EMSRb . Dar [8] finds the same pattern of results with a restricted fare structure.

	Airline B: EMSRb		Airline B: DAVN		Airline B: HBP	
	Airline A revenues	Airline B revenues	Airline A revenues	Airline B revenues	Airline A revenues	Airline B revenues
Airline A: EMSRb	+ 0 %	+ 0 %				
Airline A: DAVN	+ 2.4 %	- 1.7 %	+ 0.8 %	+ 0.7 %		
Airline A: HBP	+ 2.5 %	- 1.6 %	+ 0.9 %	+ 1.2 %	+ 0.6 %	+ 0.8 %

Fig 2.1.3 Normal form of results by Lee [22] in a network with a restricted fare structure Payoffs are % revenue gains as compared to when both airlines use EMSRb .

Zickus [38] explores different innovating forecasting techniques and finds the same pattern of pay-offs for the game of their implementation by air carriers.

Tam [34] and Vanhaverbeke [36] use PODS to evaluate the performance of different dynamic programming formulations in single markets and network environments. They show that the performance of these advanced algorithms depends strongly on the method used by the competitor. They cannot identify an unconditional winner among the formulations they study. Determining which formulation is best to implement depends on the revenue management system used by the competitor. To remedy to this problem, Tam suggests that information about the competitor be incorporated in the algorithms.

Guo [16] concentrates his thesis on the estimation of sell-up for unrestricted environments. **Sell-up** is the process by which a passenger whose first choice for a fare class is not available decides to pay more for a higher available fare class, so long as the price remains below his maximum willingness to pay. Generating passenger sell-up increases airline revenues. One of Guo's main conclusions is that the performance of sell-up estimators decreases with the sophistication of the competitors' revenue management system. The more sophisticated an airline revenue management system, the better it will capture passengers with a high willingness-to-pay, so the better it will compete with methods generating sell-up.

Gorin [15] examines average fares before and after entrance of low-cost carriers in new markets using PODS simulations. His objective is to show that competitive revenue management behaviours rather than predatory pricing can explain an overall decrease in average fares. Gorin shows how important the interactions with the revenue management system of the new competitor are. Yet the diversity of scenarios and results mentioned in his thesis do not allow for a more precise general conclusion.

Lua [28] followed by d'Huart and Belobaba [10] conduct an analysis of the competitive practice of matching the availabilities of competitor fare classes. This tactic is intended to sustain market shares and load factors in the face of low-cost competition and has been widespread in the industry. Matching the competitor's availabilities significantly overrides previously optimized booking limits, and can result in an increase in revenues for the competitor and a decrease in revenues for the proponent. Such evolution of revenues is due to the change of demand experienced by the competitor and initiated by the matching airline. For instance, if airline A matches a competitor's availabilities as soon as this competitor is more open than airline A (therefore more competitive in search engines), airline A will be more competitive fare-wise and capture more low fare demand. This low fare demand captured from the competitor frees space for high fare passengers for the competitor, thus reducing the spill of high fare passengers from the competitor to airline A. As a result, the practice of matching by airline A moves its demand from high fare classes to low fare classes, and can decrease its revenues.

Last but not least in two independent theses, Ferea [11] and Darot [9] tackle the issue of revenue management for airline alliances, and provide us with results on cooperation as compared to monopoly and competition. Ferea and Darot both show that when two network airlines optimize their seat allocation by cooperating and openly sharing their bid-prices, revenues are increased as compared to a strictly competitive situation. More surprising are Ferea's and Darot's common result that two cooperating airlines make more total revenues than a single equivalent monopoly. Darot shows that this behaviour is due to passenger spills. To do so, he configures the PODS simulator so that passengers can not be spilled to other airlines/flights/fare classes. They only book their first choice. In that configuration, the monopoly revenues are higher than two cooperating airlines.

2.1.3. Conclusions

All these results using the PODS simulator to test revenue management in competitive environments suggest the following conclusions:

a - One airline improving its revenue management system leads to an increase of its revenues at the expense of its competitors. Yet one cannot speak of a zero-sum game as the increase in revenues for one airline is higher than the decrease in revenue of its competitors. It better matches class availabilities with demand, spills less high-fare class and more low-fare class passengers to competitors. Not only does this improve the leading airline's fare class mix and revenues, it also shifts the demand experienced by the competitor towards low fare classes. The latter ends with a worse fare class mix and worse revenues. Conversely, when a tested revenue management technique happens to decrease the innovating airline's revenues, it usually increases the passive airline's revenues by having improved its mix of demand through spill.

b - When both airlines improve their revenue management system, both benefit. There is no long-term advantage of being the first or the last to move towards the improved revenue management system.

c - Two airlines cooperating improve their revenues as compared to two airlines competing. More surprising is the fact that two airlines cooperating can generate more revenues than one single monopoly with the same total capacity.

2.2 Theoretical Approaches

In this second part, we review the published theoretical models for airline revenue management that take into account the competitive environment. We only consider models where price for fare classes is a fixed exogenous variable. We do not consider papers on joint pricing and inventory control, as such approaches are still at the state of research even in non-competitive settings, and are not implemented at the vast majority of airlines.

Economic theories on price discrimination, monopoly pricing, oligopoly pricing, perfect competition pricing, pricing under capacity restrictions or production with fixed prices can at first glance provide fundamental insight into revenue management. Yet the complexities that characterize the problem of airline revenue management, such as a multitude of different products, stochastic dependent demands arriving over different time periods or competition make it impossible to apply directly any existing economic theory. As Talluri and van Ryzin [33] point out, “in most real-world contexts there are many economic forces at work operating at different levels and different time scales”. They provide a good review of the economic theories best related (but inapplicable directly) to the problem of airline revenue management.

2.2.1 Models with Independent Demands for Classes

There are only a few theoretical studies of competition that apply directly to airline revenue management. Parlar [32] followed by Karjalainen [19], Lippman and McCardle [26] as well as Mahajan and van Ryzin [29] develops a competitive version of the classic Newsboy Problem, in which a firm’s strategy is to choose an inventory level for a single perishable good before a one period stochastic demand comes. The Newsboy Problem happens to be linked to the traditional airline fare class mix control problem formulated by Littlewood [27]. An airline sells two fare classes (a high fare class and a low fare class) subject to stochastic demands, in one market only and at one moment before departure. The demand for the low fare class is assumed to be sufficient to always reach its booking limit. Littlewood’s problem is equivalent to the Newsboy Problem: The airline chooses an inventory level for the high fare class at the expense of the low fare class. If a seat protected for the high fare class is not sold, this will generate a loss equal to the price of the low fare class. On the contrary if such a protected seat is sold to a high fare class, this will generate a profit equal to the fare premium between the high fare class and the low fare class. The parallel between the Newsboy Problem and Littlewood’s problem is drawn in Fig 2.2.1.

To model competitive interactions of inventory level decisions, Parlar [32] decomposes the allocation of demand between airlines and fare classes into two successive processes: an **initial allocation** and a **reallocation** that corresponds to an **overflow**. The initial allocation of demand between airlines and fare classes is an externality corresponding to mere passenger preference for a specific airline and a specific class, and does not depend on the booking limits that airlines choose. It is in the reallocation process that an airline’s booking limits impact its competitor. In the case of two airlines, reallocated/spilled demand from airline B to airline A for the low fare class corresponds to the excess of initial demand for the low fare class of airline B as compared to its booking limit. For the high fare class, reallocated demand from airline B to airline A corresponds to the excess of initial demand for the high fare class of airline B as compared to its remaining capacity. Fig. 2.2.2 shows a graphical representation of initial allocation and reallocation/overflow.

Parlar[32] assumes that the initial allocation consists of independent demands for each of the two firms, whereas Lippman and McCardle [26] assume that these demands are correlated and Mahajan and van Ryzin [29] use a more complex dynamic consumer choice model. Parlar proves the existence and uniqueness of a Nash equilibrium to his game and gives lower and upper bounds for this equilibrium. He also shows that when one of the players acts irrationally to damage the other one, the maximin defensive solution for the latter reduces to the solution of the classic non-competitive problem. Parlar eventually argues that perfect cooperation increases the well-being of both players as compared to competition. Lippman and McCardle as well as Mahajan and van Ryzin find some similar results, but restrict themselves to the case where the unit salvage value of a non sold product and the unit opportunity cost of a shortage in inventory are equal to 0, which is not the case of the airline fare class mix control problem. Karjalainen[19] extends the case of Parlar to n-firms.

Netessine and Shumsky [31] draw inspiration from Parlar to develop a pertinent theoretical study of competitive airline revenue management that we review in more detail. They consider a duopoly fare class mix control problem with two independent classes by airline, represented in figure 2.2.2. Two competing airlines, airline 1 and airline 2, each offer capacity C_i , offer a low fare class L at price p_{Li} and a high fare class H at price p_{Hi} . There are initial random exogenous demands by fare class and by airline d_{Hi} and d_{Li} , with nonnegative support and with differentiable cumulative distributions. These four initial demands are known by both airlines. Demands are independent across fare classes but not across airlines. Reallocation of demand happens among airlines but not among fare classes. Demand for low-fare classes occurs before demand for high fare classes. Booking limits are static; they are fixed once before demand comes. Customer cancellations are ignored.

NEWSBOY PROBLEM	BASIC AIRLINE FARE CLASS MIX CONTROL
Inventory level to produce	Number of seats to protect for the high fare class from the low fare class.
Unit incremental revenue	Fare premium realized by selling a high fare class instead of a low fare class. (price of the high fare class – price of the low fare class)
Unit cost of acquiring inventory	0 <i>Seats take off whatever the protection level.</i>
Unit salvage value of a non sold product	Unit opportunity cost of a non sold seat protected for a high fare class: <i>A seat protected for a high fare class and not sold is not salvaged. Instead, there is a unit opportunity cost of not selling this seat at a low fare class, equal to the unit price of the low fare class.</i>
Unit cost of a shortage in inventory	0 <i>There is no specific cost for the airline if it protects less for high fare class than there is demand for it. Using an opportunity cost of loss if high fare classes would be double counting with the unit incremental revenue.</i>

Fig 2.2.1 Parallel between the Newsboy Problem and Littlewood's Revenue Management Problem.

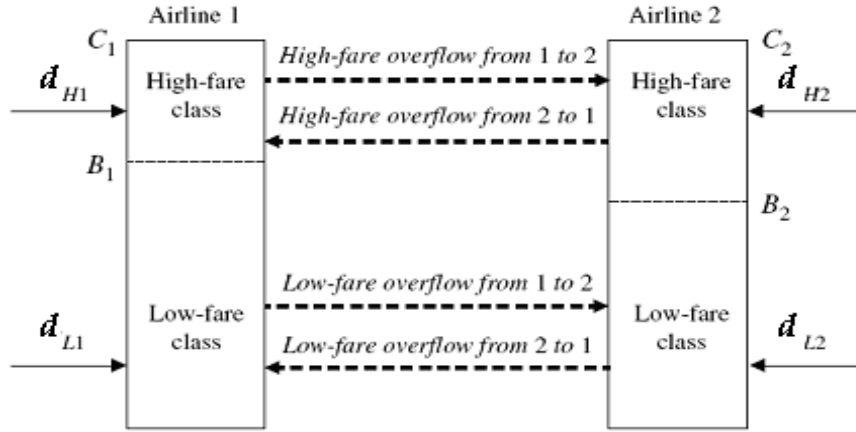


Fig 2.2.2 The model of Netessine and Shumsky [31]

In that model, total demands by airline and by fare class can be written:

$$\begin{aligned}
 (2.1) \quad D_{H1} &= d_{H1} + \text{Max}(0, d_{H2} + \min(D_{L2}, B_2) - C_2) \\
 D_{H2} &= d_{H2} + \text{Max}(0, d_{H1} + \min(D_{L1}, B_1) - C_1) \\
 D_{L1} &= d_{L1} + \text{Max}(0, d_{L2} - B_2) \\
 D_{L2} &= d_{L2} + \text{Max}(0, d_{L1} - B_1)
 \end{aligned}$$

where:

- D_{Li} corresponds to the total demand for airline i, fare class L
- d_{Li} corresponds to the initial allocation of demand to airline i, class L
- B_i corresponds to the booking limit fixed by airline i.
- $\min(D_{Li}, B_i)$ is the remaining capacity for class H, airline i after allocation of L seats.
- C_i is the capacity of airline i
- $\text{Max}(\dots)$ corresponds to the spill of demand coming from the other airline for the fare class.

Each airline's strategy consists of its booking limits for the low fare class.

Under these assumptions, revenue for airline i can be written:

$$(2.2) \quad \Pi_i = E(p_{Li} \min(D_{Li}, B_i) + p_{Hi} \min(D_{Hi}, C_i - \min(D_{Li}, B_i)))$$

where:

- $E(\)$ is an expected value
- p_{Li} is the price of fare class L, airline i

The derivative of airline i's profit with respect to its booking limit is:

$$(2.3) \quad \frac{\partial \Pi_i}{\partial B_i} = p_{Li} \Pr(D_{Li} \geq B_i) - p_{Hi} \Pr(D_{Hi} \geq C_i - B_i, D_{Li} \geq B_i) \\ - p_{Hi} \Pr(d_{Li} \geq B_i, d_{Hj} \geq C_j - \min(D_{Lj}, B_j), D_{Lj} \leq B_j, D_{Hi} \leq C_i - B_i)$$

The first two terms correspond to the equation that leads to Littlewood's rule [27] in a non competitive environment. This rule has been extended to more than two classes by Belobaba[1] with his EMSR model and applied by many airlines. Yet in Netessine and Shumsky's model the demands D_{Li} and D_{Hi} are a function of the competitor's booking limits as expressed in (2.1). The third and more complex term corresponds to a competitive feedback of airline revenue management.

Netessine and Shumsky [31] show that if $(d_{L1}, d_{H1}, d_{L2}, d_{H2})$ are multivariate totally positive of second order in their density function, that is that they are more prone to be all high together or low together rather than mixed high and low, then a pure strategy Nash equilibrium exists, and best-response functions dB_i/dB_j are decreasing. If one player increases his booking limit, then it is optimal for the other player to decrease his booking limit.

For their further results, Netessine and Shumsky restrict themselves to the case where booking limits of both airlines are reached. This is a traditional assumption of airline revenue management methods. This assumption has been made since the beginning in Littlewood's rule [27], and fare class structures are supposed to be designed to meet this condition. In that case, spills for the low-fare class become irrelevant; there is no need to consider them. As long as booking limits are supposed to be reached, the exact value of low fare class spill should not impact the strategies of the two airlines.

Under this condition, the optimality condition where (2.3) is equal to 0 is simply reduced to the basic Littlewood equation but where D_{Hi} is a function of B_j :

$$(2.4) \quad \Pr(D_{Hi} \geq C_i - B_i) = \Pr(d_{Hi} + \max(0, d_{Hj} + B_j - C_j) \geq C_i - B_i) = \frac{P_{Li}}{P_{Hi}}$$

Under this condition, Netessine and Shumsky [31] conclude that whatever the demand functions d_{Li} and d_{Hi} , best-response functions dB_i/dB_j are decreasing. In equation (2.4), the higher B_j , the smaller B_i has to be. This last equation also defines a unique Nash equilibrium. Although in their paper Netessine and Shumsky assert that booking limits under this competitive game are lower than under the cooperative game, we have agreed with the authors that their proof is as of today erroneous, and that the result might be incorrect. Still through extensive numerical simulations, they found that the same conclusions can be drawn in many cases even if the assumption of reached booking limits is not held.

Talluri and van Ryzin [33] propose a graphical interpretation of this game, as provided in Fig 2.2.3. If best response functions are decreasing, an airline's best response arcs cannot cross, and an equilibrium exists.

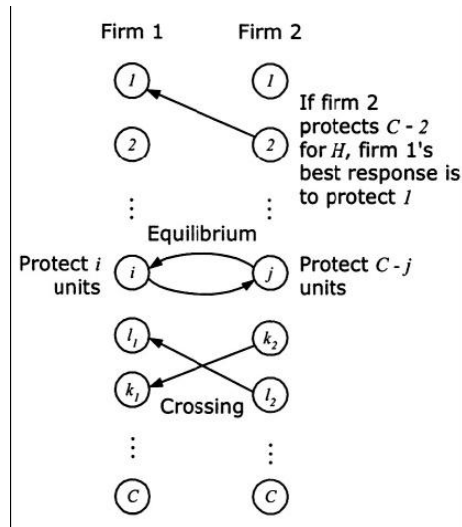


Fig 2.2.3 Bipartite graph of the two airline, independent classes, competitive game.
Source: Talluri and van Ryzin [33]

Li, Zhang A. and Zhang Y. [25] have an approach similar to Netessine and Shumsky [31]. Instead of a random exogenous demand model by airline and by fare class, they introduce a random exogenous demand by fare class only, and split it by airline. The rule for splitting is: an airline's market share for a fare class is proportional to the capacity offered in that fare class by the airline. For the high fare class, the latter would correspond to the protection level, and for the low fare class it would correspond to the booking limit. They refer to this rule as a "proportional rationing rule" found in Tirole[35]. This direct correlation between market shares and capacity is commonly made in the airline industry. Li, Zhang A. and Zhang Y. use a model with identical fares for both airlines, but different costs associated for each fare class and for each airline. This model is equivalent to no costs but different fares for the two airlines. They show that if the ratio p_L/p_H is not 3 times higher from one airline to another, then there exists a unique and stable pure-strategy Nash Equilibrium. They also show that if the two airlines are symmetric in prices, then choosing whether to set collusion optimal booking limits or competition optimal booking limits leads to a Prisoners Dilemma. Strategically, both airlines benefit from implementing revenue management as shown in Belobaba and Wilson [2], but tactically both airlines benefit from cooperating in their seat allocation rather than competing.

In a shorter but similar study, Li, Oum and Anderson [24] had shown that under another set of conditions on incremental expected revenue functions, a symmetrical equilibrium exists, and that this competitive equilibrium has lower booking limits than a cooperative situation. It is interesting to note that Mahajan and van Ryzin [29] also found that under competition, protection for higher classes would be higher than with cooperation, even though their restriction of the newsboy problem does not apply directly to airline revenue management.

2.2.2 Models with dependent demands for classes

All of the above models consider that total demands for classes are independent. This can be considered in markets with fare structures with clear different restrictions for each fare class, also called fully-restricted fare structures. Isler and Imhof [18] introduce a game theoretic model for less restricted fare structure environments (also referred to as unrestricted fare structure environments), where restrictions are the same from one class to another, so that such independence of demands does not hold anymore. The revenue management method they consider consists in the dynamic programming formulation of Gallego and van Ryzin [13]. Isler and Imhof [18] first show that under this scenario, two competing revenue management systems neglecting the fact that they are in competition will generate overall revenues much smaller than if they were fully cooperating. In a second part they show that if the two airlines set their prices simultaneously, if they know the demand behaviour and the inventory counts of each other, these two airlines can reach a pure strategy subgame perfect equilibrium. The total market revenue associated with this equilibrium is higher than the usual competition-blind optimization, but not as high as the perfect cooperation case.

Fig 2.2. 4 sums up Isler and Imhof results. It compares total market revenue generation in three scenarios: full cooperation, full “blind” spiralled-down competition and the game theoretic model for competition introduced in their article.




Fig 2.2.4 Results of Isler and Imhof [18] simulation.

2.2.3 Conclusions

Theoretical research on competition in the field of airline revenue management with fixed exogenous prices has often come to the same conclusions under different sets of scenarios and conditions:

a - There exists a Nash equilibrium if both airlines compete rationally and take into consideration competitive effects.

b - Airlines competing would benefit from taking into consideration competitive effects. They would benefit even more if they cooperated for their allocation of seat inventory to different fare classes.

c - Best-response functions dB_i/dB_j are decreasing. If one player increases his booking limit, then it is optimal for the other player to decrease his booking limit, and conversely.

d – Total market booking limits are smaller under competition than under collusion.

Yet as Gallego, Krishnamoorthy and Phillips [14] pertinently notice, most of these models are single-stage games, and might not be most appropriate to describe a real-world situation. Airlines marketing departments interact on a daily basis and have built a history of tactical moves and responses on which they rely to make their present decisions, including some implicit collusion. A repeated game might be more suited theoretical framework for competitive revenue management.

Furthermore these theoretical models assume that revenue management games are perfect information games. But in reality airlines do not have perfect information on the competitor's availabilities/protection levels. Neither do they even have a perfect knowledge of their own demand, they can only rely on imperfect forecasts.

CHAPTER 3.

IMPACT OF COMPETITION ON CURRENT AIRLINE REVENUE MANAGEMENT SYSTEMS

The fact that the revenue management models widely implemented by airlines do not explicitly take into account competitive effects does not mean that competitive effects do not exist. In this chapter we describe the passive competitive behaviour of current airline revenue management systems, and link our observations with the results of previous literature as summarized in 2.1.3 and 2.2.3. We first use the comparison of monopoly scenarios with competitive scenarios to reveal some behaviours at the total market level of revenue management systems in competition. We then focus on the impact at the level of a single airline of competition between currently used revenue management systems.

3.1 Competition versus Monopoly

The first step towards understanding how a competitive environment affects current airline revenue management is to compare competitive scenarios with their equivalent monopoly scenarios in terms of seat inventory control. Determining how the group of competitors set their booking limits/protection levels as compared to a monopoly helps understand how revenue management is impacted by competitive situations. We first provide a competitive model, a general intuitive idea, and some theoretical results. We then run PODS simulations to see if our ideas, based on restrictive hypotheses, could still be valid in more complex real-world situations.

3.1.1 Preliminary Theoretical Results

Before engaging into simulations of competitive vs. monopolistic environments, we provide theoretical results extending one of the conclusions drawn in the literature review, which concluded that in competitive settings total market booking limits on lower priced fare classes are prone to be smaller than under equivalent monopolies.

Our model is to some extent similar to the model of Parlar [32] as well as Netessine and Shumski [31]. Let us first define our terminology as well as the general intuitive idea behind our results. We place ourselves in the case of stochastic demands. The demand that airlines observe is constrained and consists only of accepted bookings. We first define **non-constrained demand** for a fare class at an airline as the sum of the passengers who book the fare class at the airline plus the passengers who would have booked this fare class if it were available at the airline, but who cannot because the fare class is not available. Let us first suppose that the non-constrained demands for classes are independent. At the same time, with a competitive market, we can view non-constrained demand for an airline as split in two other components (different from the categorization of accepted and rejected bookings), as illustrated in Fig 3.1.1, Fig 3.1.2 and Fig 3.1.3.

On the one hand, **first-choice non-constrained demand** is the initial share of total market non-constrained demand that would book the fare class on the airline as a first-choice preference. The sum of the first choice non-constrained demands for each airline is equal to the total market non-constrained demand. On the other hand, **spill-in** corresponds to passengers who would as a first-choice book on a competitor but who were rejected because the fare product they want to purchase is not available anymore at the competitor. We use the term spill-in for the passengers spilled by competitors towards a specific airline, and **spill-out** designates passengers who are rejected from a specific airline and are added to the non-constrained demands of some competitors. Both spill-in and spill-out are precise terms that refer to the same general phenomenon specific to competitive environments, spill. Among the observed effective bookings for the fare class of an airline, some consist of first-choice bookings, some are bookings by passengers who would have preferred to buy the competitor's fare class but who were spilled-in. Among the passengers that an airline rejects, some considered this airline as a first-choice and will spill-out to the competitors. The others are passengers who were already rejected by one competitor, so they will either spill-out to remaining competitors or not book on any airline. We call the latter **no-go demand**.

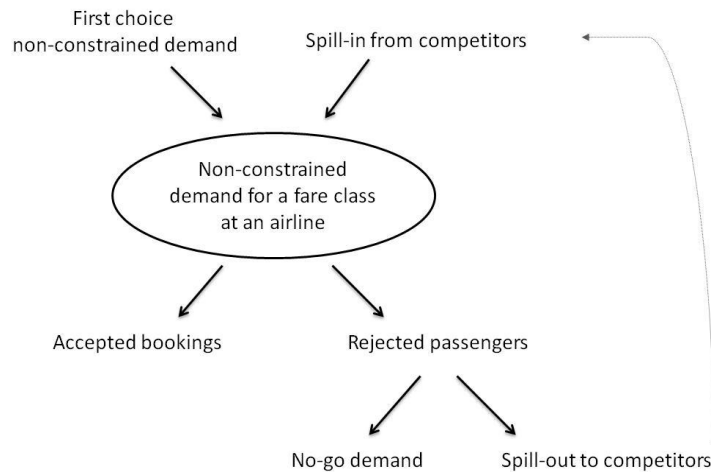


Fig 3.1.1 Decomposition of non-constrained demand for a fare class at an airline

	First-choice non-constrained demand	Spill-in from competitors
Accepted bookings	Captured first-choice demand	Captured spill-in from competitors
Rejected passengers	Spill-out to competitors	Spill-out to competitors or No-go demand

Fig 3.1.2 Decomposition of non-constrained demand for a fare class at an airline. Non-constrained demand for a fare class at an airline is the sum of the four cells.




Fig 3.1.3 Competitive approach, duopoly situation for a set of classes (1..i).

In competition and at the total market level, passengers which are at some point spilled to a competitor are double-counted as part of the non-constrained demand of more than one airline. They increment the non-constrained demand of the airline where they had their first choice preference, as well as the non-constrained demand of the competitor(s) to which they spill-in. Such double-counting of a passenger at the total market level does not occur if there is only a monopoly serving the market. Therefore the sum of the non-constrained demands for each airline of an oligopoly is larger than the non-constrained demand for a monopoly.

Airline revenue management system forecasters integrate an unconstraining module which provides an estimation of non-constrained demand by “unconstraining” the observed constrained demand (accepted bookings), which is the only data available to an airline. This estimation of non-constrained demand is referred to as **unconstrained demand**, and is used to set protection levels. Assuming that the unconstraining modules of forecasters are reasonably accurate, the sum of the unconstrained demands forecasted by airlines in an oligopoly should be larger than the total market non-constrained demand, and larger than the unconstrained demand estimated by a monopoly. As a consequence, assuming that the protection levels set by the airlines are not exponentially increasing in the expected value of forecasted unconstrained demand, the total market protection level set by an oligopoly should be higher than the protection level set by an equivalent monopoly.

Fig. 3.1.6, Fig.3.1.7 and Fig.3.1.8 provide an example, where μ designates the average value of a distribution. Fig 3.1.6 describes a monopoly situation for the determination of a nested protection level. Total market non-constrained demand has an expected value μ of 100. The monopolist observes an average of 78 accepted bookings, with an average protection level of 80 set with an average forecast of unconstrained demand of expected value 99 (the forecaster is not perfect). On average, no-go demand is of 22 passengers. Note that the average of no-go demand is higher than the difference of the average of non-constrained demand and the protection level because an occurrence of no-go cannot be negative. The sum of average no-go and of average accepted bookings is equal to average non-constrained demand. Fig. 3.1.7 displays its equivalent duopoly situation if passengers did not spill between airlines and only booked at their preferred one. Fig 3.1.7 does not represent a realistic situation, but is useful to better understand by comparison the case of Fig 3.1.8. At the total market level, there is not much difference between the case represented in Fig 3.1.6 and the case represented in Fig 3.1.7.

Fig 3.1.8 represents the more realistic case of an equivalent duopoly situation where passengers spill between two identical airlines. Each airline has a first-choice non-constrained demand of expected value 50 corresponding to half of the total market non-constrained demand. Each spills-out an average of 8 passengers to the other. Therefore the non-constrained demand for each airline has an expected value of 58, larger than the expected value of the first choice non-constrained demand for that airline. Each airline accepts on average 44 bookings, with an average booking limit of 45 set with an unconstrained forecast of average value 55. Of the average 8 passengers spilled-out to the other airline, on average 2 book on the competitor, and 6 do not book on any airline (no-go). Each airline is also able to capture more of its first choice non-constrained demand, with on average 3 more seats available for such demand than if there was no spill between competitors as described in Fig 3.1.7. For the duopoly, total market protection level is on average 90 seats (as compared to 80 for the monopoly), total market unconstrained forecasts are on average 110 (as compared to 99 for the monopoly), total market accepted bookings are on average 88 (as compared to 78 for the monopoly), and the total no-go demand is on average 12 (as compared to 22 for the monopoly).

This general intuitive idea that protection levels are higher in competition than in a monopoly can be extended to the case where demands for classes are dependent variables if the total market capacity of the oligopoly is not higher than the seat capacity of the monopoly. When such demands are not independent, there exists **sell-up** between classes. Sell-up designates the decision by a passenger to book a fare class that is higher than his first-choice fare class desired when the latter is not available. If demands for classes were dependent, sell-up from lower classes than the ones of the set of classes considered would be added to the first choice non-constrained demand and to spill-in as shown in Fig. 3.1.4 and 3.1.5.

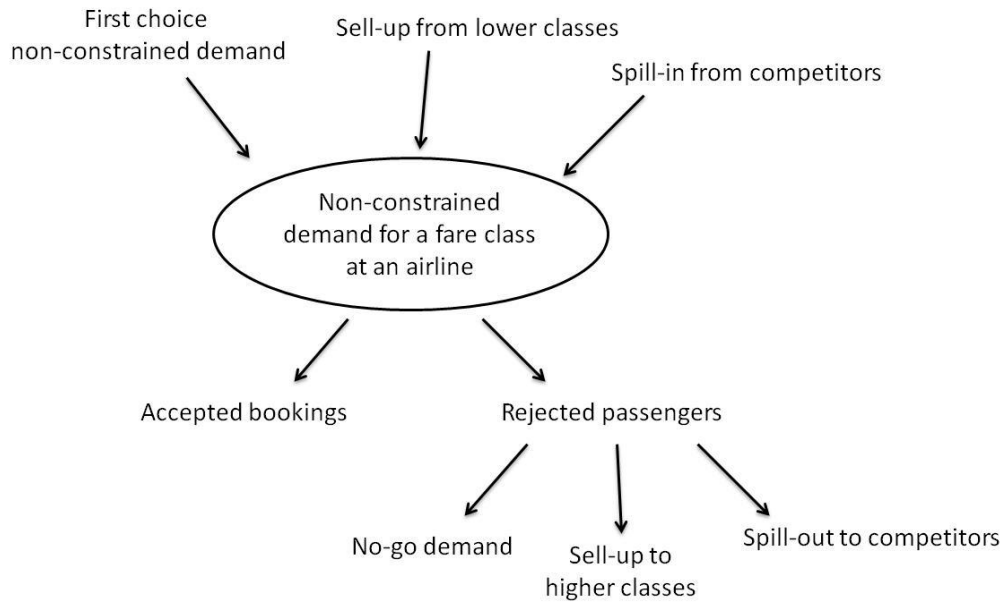


Fig 3.1.4 Decomposition of non-constrained demand with sell-up

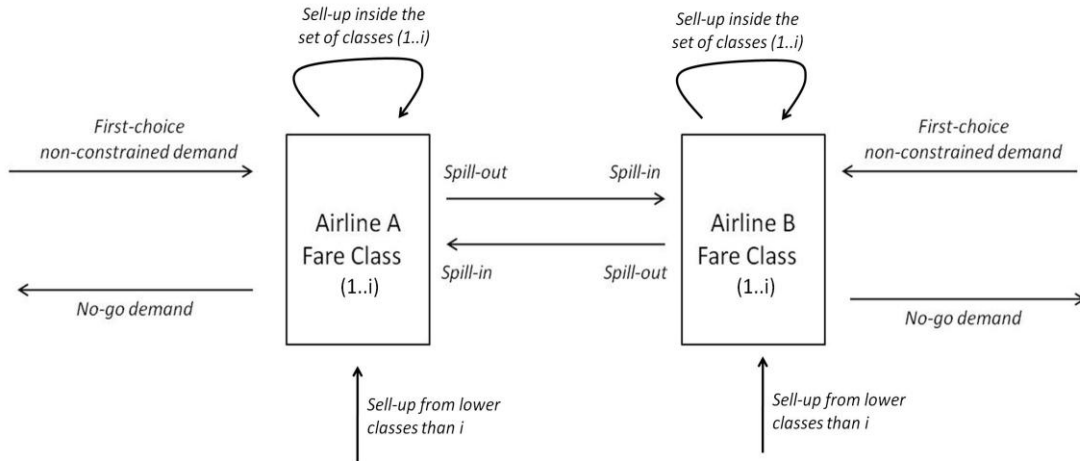


Fig 3.1.5 Competitive approach, duopoly with sell-up for a set of classes (1..i).

If the oligopoly total market capacity is not higher than the capacity of the monopoly, even with sell-up total market protection levels should be higher in competition than in a monopoly, because of larger double-counting. First, spill does not exist in a monopoly situation and is strictly positive in the oligopoly situation. Second, we believe that at the total market level, sell-up should then be more important in the oligopoly than in the monopoly. In monopoly, the only reason for sell-up is that total market booking limit has been reached. In competition, sell-up can happen when the total market booking limit is reached, but also when the total market booking limit is not reached yet, but the passenger has a preference for an airline whose booking limit is reached. Even though higher sell-up reduces spill-in/spill-out, at the total market level the difference with a monopoly is that competition introduces spill-in (that is non-existent in a monopoly) and introduces additional sell-up from lower classes as compared to the monopoly. Hence total market protection levels should be higher in competition than under monopoly, which will also have the feedback effect of increasing even more sell-up at the total market level for the oligopoly.

Fig 3.1.6bis and Fig 3.1.8bis represent respectively an example of a monopoly situation with sell-up and its equivalent duopoly situation (the same capacity for the monopoly and the oligopoly is assumed). The figures represent the fare class mix optimization problem for a set of fare classes, so there are two types of sell-up: sell-up from strictly lower classes, and sell-up within the set of fare classes. At the total market level, both increase in the oligopoly as compared to the monopoly. Sell-up encroaches on spill, which is smaller in the model with sell-up. Yet spill does not exist in the monopoly case anyways. Double-counting of passengers in total market unconstrained demands is consequently higher in the competition case than in the monopoly case because of an increase in overall sell-up and because of the notion of spill.

Last but not least, our general intuitive idea is reinforced if one assumes that competition boosts total market non-constrained demand, even with fixed total market seat capacity. Our model does not take into account no-show behaviour and corresponding overbooking techniques. For relevance we only focused on understanding the mechanisms of competition influencing fare class mix optimization.

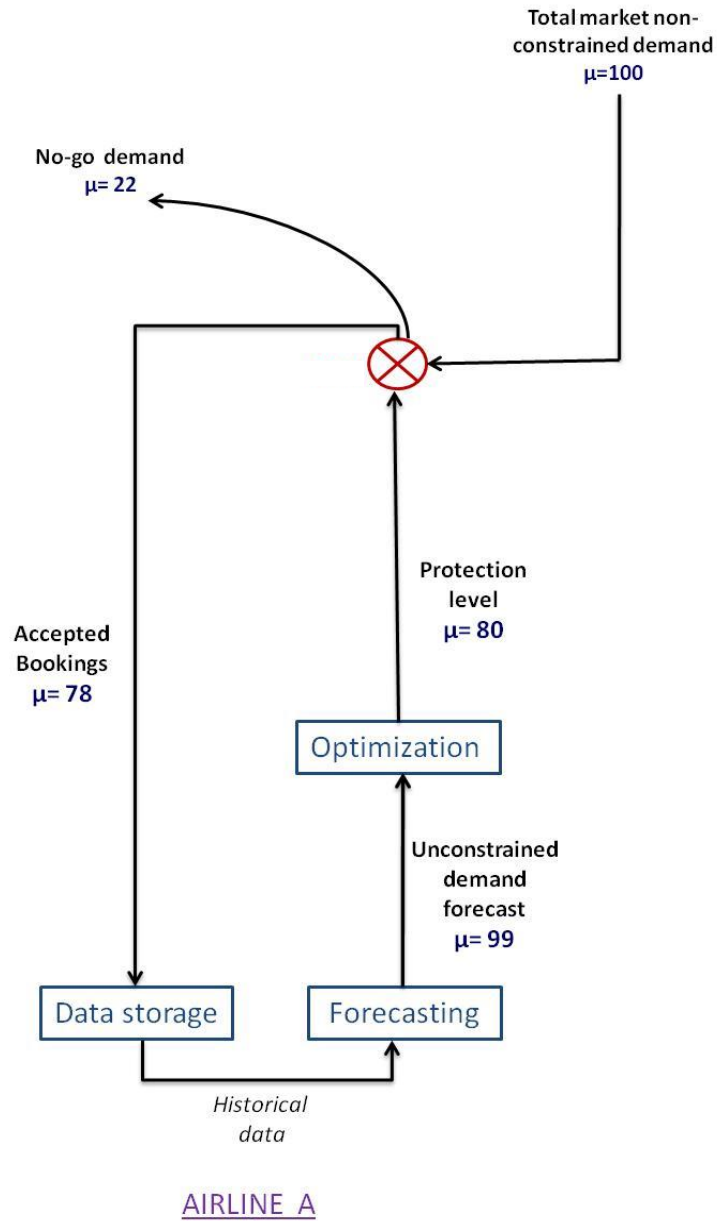
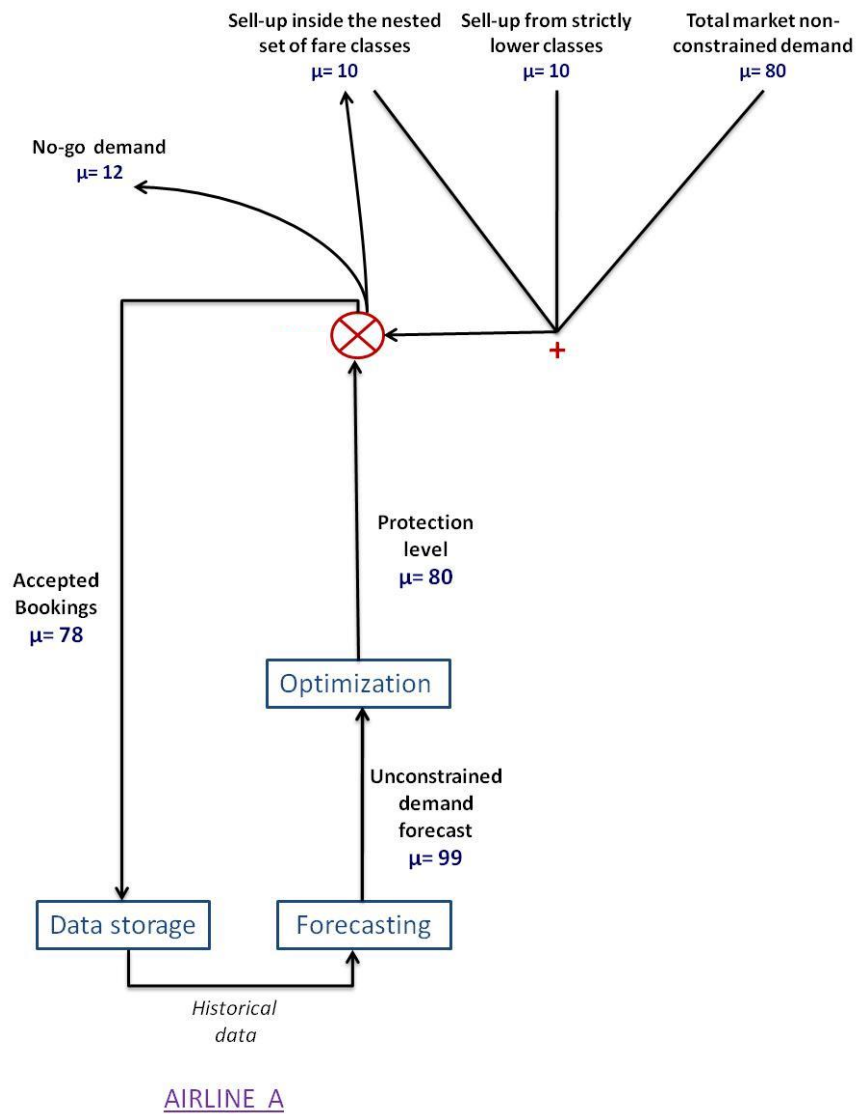



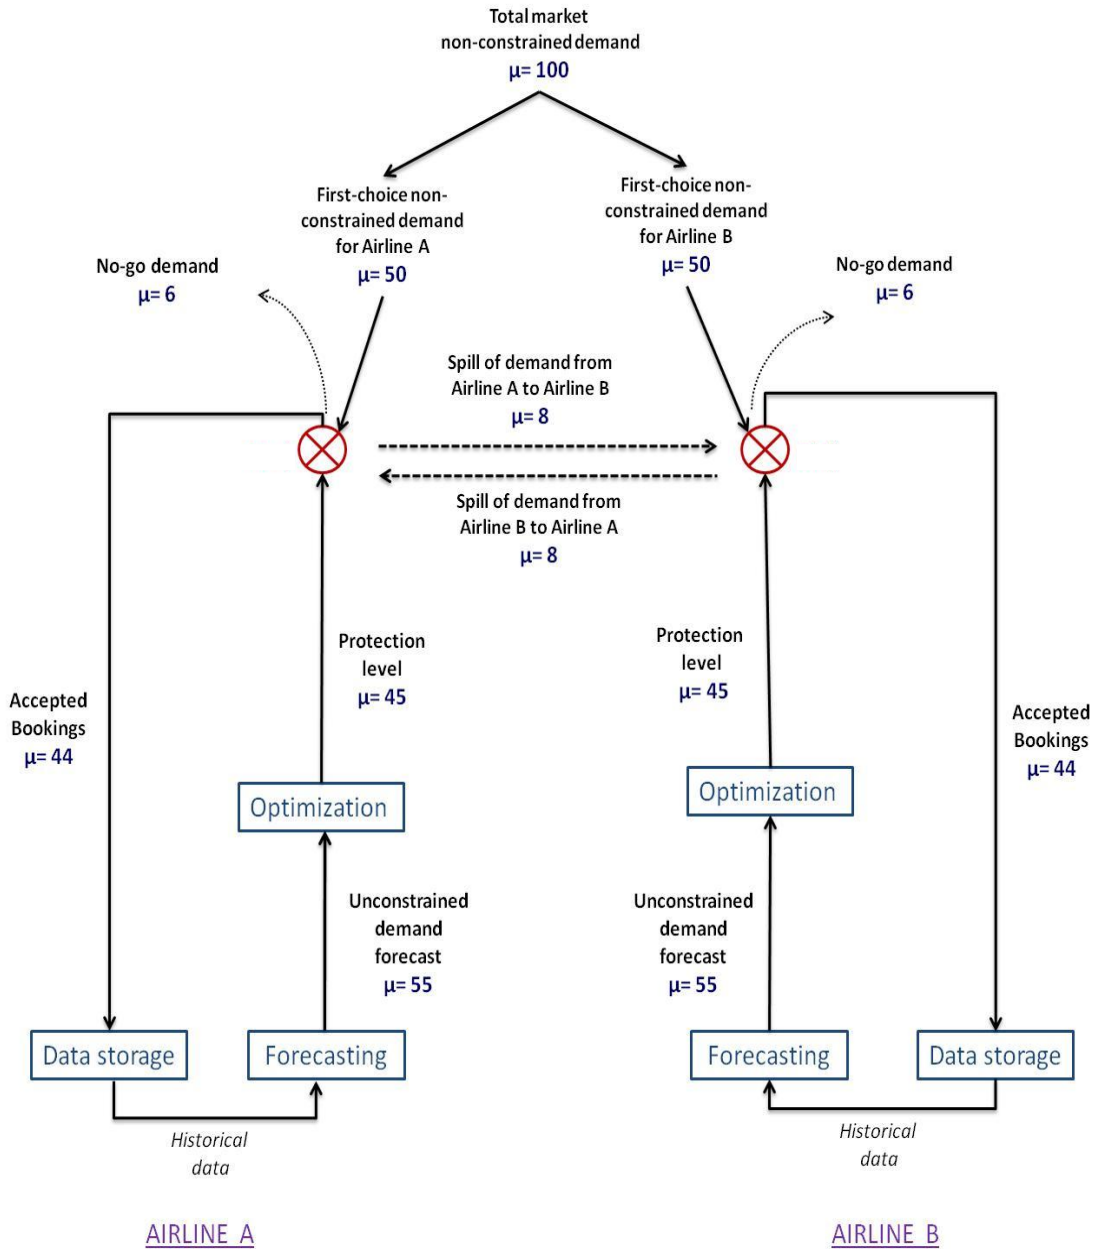
Fig 3.1.6 Determination of a nested protection level
 Example of monopoly situation.



*Fig 3.1.6bis Determination of a nested protection level
 Example of monopoly situation with passenger sell-up
 “+” refers to a sum operator*



*Fig 3.1.7 Determination of a nested protection level
 Example of a duopoly situation where passengers do not spill-in between airlines
 This case does never happen but is useful for a better comprehension.*



*Fig 3.1.8 Determination of a nested protection level
Example of a duopoly situation with passenger spill-in between airlines*

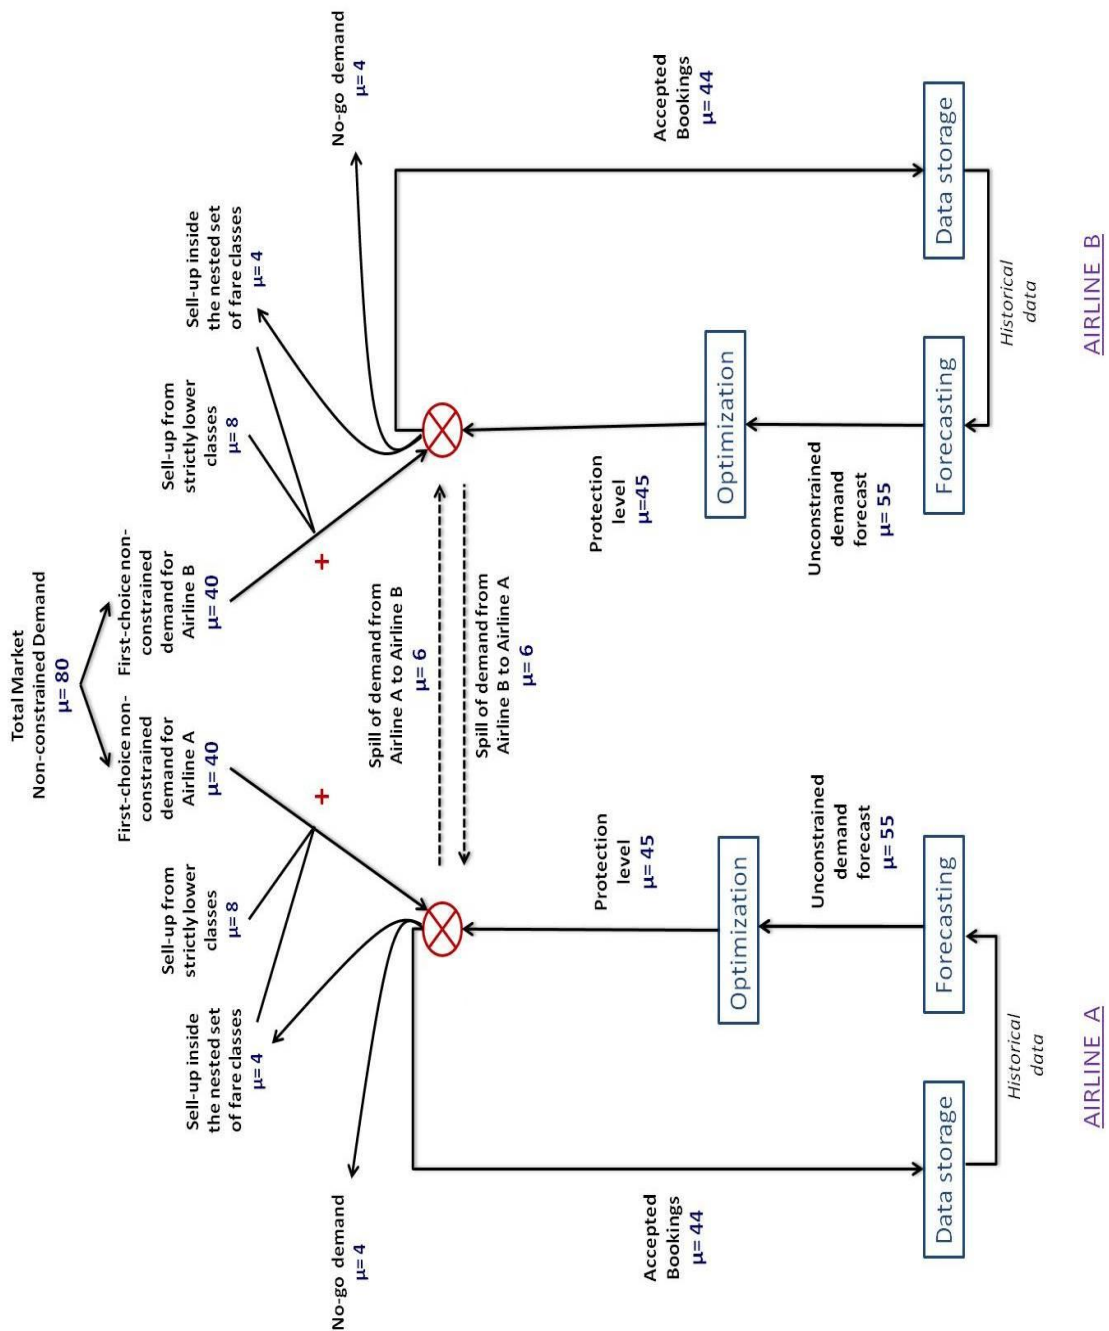


Fig 3.1.8bis Determination of a nested protection level
 Example of a duopoly situation with passenger sell-up within each airline and passenger spill-in between airlines.
 “+” refers to a sum operator

The existing literature has proved the result of this general intuitive idea in the case of two airlines each offering two classes with independent first-choice non-constrained demands. It has also made additional assumptions on these demands. The following Proposition A extends this result to competition between N airlines each offering 2 classes, subject to total market independent non-constrained demands. The restrictions of proposition A are that all airlines have the same fare structure, that the initial share of the total market non-constrained demand by airline for the high fare class is known, and that booking limits are reached (Littlewood's rule).

Proposition A:

Suppose:

- A market where N airlines have an identical 2 class fare structure.
- Total market non-constrained demand for each fare class is a stochastic independent variable.
- The booking limits for the low fare class are reached.
- There is an initial share between airlines of the total market non-constrained demand for the high fare class based on passenger first choice preference.
- Unconstrained forecasts accurately estimate non-constrained demands (non-constrained demand is equal to unconstrained demands).

Then:

- The total market protection level for the high fare class is higher under competition than under monopoly, regardless of the capacities.
- If the oligopoly does not offer a strictly higher total market capacity than the monopoly, total market booking limits are smaller in the oligopoly than in the monopoly.

Proof:

Let:

- d_1 be the stochastic total market non-constrained demand for the high fare class 1.
- p_i be the fare of fare class i .
- $\lambda^j \in [0;1]$ be the proportion of the total market non-constrained demand for the high fare class 1 that has a preference for airline j . By definition $\sum_{j=1..n} \lambda^j = 1$.
- $d_1^j = \lambda^j \cdot d_1$ be the stochastic first-choice non-constrained demand for the high fare class 1 of airline j

The monopoly non-constrained demand for the high fare class is d_1 . The monopoly optimizes its revenues using Littlewood's two class rule (a detailed description of this rule can be found in Talluri and van Ryzin [33] p.35-36) and with unconstrained forecasts accurately estimating demand d_1 .

Let:

- $\varepsilon_1^j \geq 0$ be the stochastic spill-in of demand towards airline j , fare class 1
- $D_1^j = d_1^j + \varepsilon_1^j$ be the stochastic total non-constrained demand for airline j , fare class 1

Let :

- $X : d \rightarrow X(d)$ be the application associating a positive stochastic demand d for the high class with its Littlewood protection level. $X(d)$ is the solution of

$$(3.1) \quad \Pr(d \geq X(d)) = \frac{P_2}{P_1}$$

- X does not depend on the airline.
- X is a positive homogenous application: $\forall a \in \mathfrak{R}^+, X(a.d) = a . X(d)$

This results directly from the fact that $\Pr(d \geq X) = \Pr(a.d \geq a.X) = \frac{P_2}{P_1}$

- $X(d + d^*) \geq X(d)$ for any set of stochastic positive demands d and d^* .

This can be shown knowing that:

1. by definition of X , $\Pr(d + d^* \geq X(d + d^*)) = \Pr(d \geq X(d)) = \frac{P_2}{P_1}$
2. for d, d^* positive stochastic demands, x a real constant, $\Pr(d + d^* \geq x) \geq \Pr(d \geq x)$
3. for a given distribution of the demand d , $x \rightarrow \Pr(d \geq x)$ is decreasing in the real variable x
4. thus to reach 1. from 2. one must decrease x on the left side of equation 2

Therefore we obtain:

$$(3.2) \quad \begin{aligned} \forall (i,j) \quad X(d_i) &= \sum_j \lambda^j . X(d_1^j) \\ &= \sum_j X(\lambda^j d_1) \\ &= \sum_j X(d_1^j) \\ X(d_i) &\leq \sum_j X(d_1^j + \varepsilon_1^j) = \sum_j X(D_1^j) \end{aligned}$$

$X(d_i)$ is the protection level set by the monopoly and $\sum_j X(D_1^j)$ is the sum of the protection levels set by the competing airlines. This inequality shows the first point of Proposition A. The second point of Proposition A comes from the definition of a booking limit as the flight capacity minus a protection level.

End of Proof

With the same general reasoning and provided additional assumptions, Proposition B extends the result of Proposition A to the case of n nested protection levels and EMSRb optimization.

Proposition B

Suppose:

- A market where N airlines have an identical n-class fare structure.
- Total market non-constrained demand for each fare class is a stochastic independent variable.
- Nested booking limits on all lower classes are reached (Littlewood's rule).
- First-choice non-constrained demands is a fixed proportion of total market non-constrained demand for all the classes of an airline.
- Each airline optimizes its revenues with an EMSRb rule where the fare levels considered in the EMSRb equation are not adjusted (a detailed description can be found in Talluri and van Ryzin [33] p. 47-50).
- Unconstrained forecasts accurately estimate non-constrained demands (non-constrained demand is equal to unconstrained demands).

Then:

- The total market nested protection level for classes is higher under competition than under monopoly, regardless of the capacities.
- If the oligopoly and the monopoly offer the same total market capacity, total market nested booking limits are smaller in the oligopoly than in the monopoly.

Proof:

The overall reasoning is the same as for proposition A, with complications due to nesting.

Let:

- d_i be the stochastic total market non-constrained demand for fare class i .
- $d_{1..i} = \sum_{k=1..i} d_k$ be the total stochastic market non-constrained demand for the set of classes $(1..i)$.
- $(d_i)_{i=1..n}$ constitutes a set of independent stochastic variables.
- $(d_{1..i})_{i=1..n}$ thus also constitutes a set of independent stochastic variables.
- p_i be the fare of fare class i
- $\lambda^j \in [0;1]$ be the proportion of the total market non-constrained demand that has a preference for airline j . By definition $\sum_{j=1..n} \lambda^j = 1$
- $d_i^j = \lambda^j \cdot d_i$ be the stochastic first-choice non-constrained demand for class i , airline j .
- $d_{1..i}^j = \lambda^j \cdot d_{1..i}$ be the stochastic first-choice non-constrained demand for the set of classes $(1..i)$ of airline j .

Let:

- $\varepsilon_i^j \geq 0$ be the stochastic spill-in of demand towards airline j , fare class i
- $D_i^j = d_i^j + \varepsilon_i^j$ be the stochastic total non-constrained demand for airline j , class i
- $\varepsilon_{1..i}^j = \sum_{k=1..i} \varepsilon_k^j$ be the stochastic total spill-in of demand towards airline j for the set of fare classes $(1..i)$
- $D_{1..i}^j = \sum_{k=1..i} D_k^j = d_{1..i}^j + \varepsilon_{1..i}^j$ be the total non-constrained demand for airline j for nested classes $(1..i)$

Let :

- $X_{1..i} : (\partial_1, \dots, \partial_i) \rightarrow X_{1..i}(\partial_1, \dots, \partial_i)$ be the application associating a set $(\partial_1, \dots, \partial_i)$ of demands for the classes of a set $(1..i)$ with the nested protection level given by the non-adjusted EMSRb equation:

$$(3.3) \quad Pr \left(\sum_{k=1..i} \partial_k \geq X_{1..i}(\partial_1, \dots, \partial_i) \right) = \frac{P_{i+1}}{\sum_{k=1..i} p_k \cdot \overline{\partial}_k} \cdot \frac{\sum_{k=1..i} \overline{\partial}_k}{\sum_{k=1..i} \overline{\partial}_k}$$

- $X_{1..i}$ does not depend on the airline because they have the same fare structure and the same initial share of the total market non-constrained demand by fare class.

- $X_{1..i}$ is a positive homogenous application:

$\forall a \in \mathfrak{R}^+, X_{1..i}(a \cdot \partial_1, \dots, a \cdot \partial_i) \rightarrow a \cdot X_{1..i}(\partial_1, \dots, \partial_i)$ This results from the fact that:

$$Pr \left(\sum_{k=1..i} \partial_k \geq X_{1..i} \right) = Pr \left(\sum_{k=1..i} a \cdot \partial_k \geq a \cdot X_{1..i} \right) = \frac{P_{i+1}}{\sum_{k=1..i} p_k \cdot \overline{\partial}_k} \cdot \frac{\sum_{k=1..i} \overline{\partial}_k}{\sum_{k=1..i} \overline{\partial}_k}$$

- $X_{1..i}(\partial_1 + e_1, \dots, \partial_n + e_n) \geq X_{1..i}(\partial_1, \dots, \partial_n)$ for any set of initial stochastic non-demands $(\partial_1, \dots, \partial_n)$ and of stochastic spill-in (e_1, \dots, e_n) . This can be shown as:

1. for $(\partial_1, \dots, \partial_i)$ given, $x \rightarrow Pr(\sum_{k=1..i} \partial_k \geq x)$ is decreasing in the real variable x
2. if x is a real constant, $Pr(\sum_{k=1..i} \partial_k + \varepsilon_k \geq x) \geq Pr(\sum_{k=1..i} \partial_k \geq x)$
3. in the equation that defines $X_{1..j}$, when adding spill-in all else equal, and in “non-extreme situations” the proportional change in the weighted average revenue in the right term is smaller than the proportional increase in the left probability term.

Therefore we obtain:

$$\begin{aligned}
 (3.4) \quad \forall i \quad X_{L..i}(d_1, \dots, d_i) &= \sum_j \lambda^j \cdot X_{1..i}((d_1, \dots, d_n)) \\
 &= \sum_j X_{1..i}((\lambda^j d_1, \dots, \lambda^j d_n)) \\
 &= \sum_j X_{1..i}(d_1^j, \dots, d_n^j) \\
 X_{L..i}(d_1, \dots, d_i) &\leq \sum_j X_{1..i}(d_1^j + \varepsilon_1^j, \dots, d_n^j + \varepsilon_n^j) = \sum_j X_{1..i}(D_1^j, \dots, D_n^j)
 \end{aligned}$$

$X_{L..i}(d_1, \dots, d_i)$ is the protection level set by the monopoly for $(L..i)$ and $\sum_j X_{1..i}(D_1^j, \dots, D_n^j)$ is the sum of the protection levels set by the competing airlines for $(L..i)$. This inequality shows Proposition B.

End of Proof

3.1.2 Experimental methodology

We run different simulations whose representation of passenger and airline behaviour is closer to the real world of airline revenue management than the restrictive hypotheses of the presented theoretical results. We compare their results with the conclusions of the literature review and with our general intuitive idea and our own Propositions A and B above. We also try to use the results of the simulations to extract some more insight that theoretical results might have overlooked.

The simulations are run using the PODS simulator with parameters as defined in Chapter 1.5. We study a single market situation to better discriminate the differences in the behaviours of the revenue management systems in competitive versus monopoly environments. The monopoly scenario consists of one airline operating three flights a day, at three different times, in one direction and with three aircraft of a 240 seat capacity each. The equivalent competitive scenarios simulated are 2, 3 or 4 airlines competing. As in the monopoly situation, each airline operates three flights a day, at the same times, in the same direction, and the capacity of the aircraft in the competitive scenario is such as the total capacity offered in the market remains of 240 seats per departure time. In the competitive case, airlines all keep the same fare structure, forecasting and optimization methodologies as the monopoly. Fig.3.1.9 gives a representation of the competitive scenarios versus the monopoly scenario. No airline uses overbooking techniques.

We run the following set of scenarios of competition versus monopoly:

1. A single market with the fully-restricted fare structure provided in Fig 3.1.10. Both airlines use EMSRb optimization and pick-up forecasting with booking curve unconstraining. This case is very similar to the conditions of Proposition B.
2. A single market with the fully-restricted fare structure provided in Fig 3.1.10. Both airlines use Lautenbacher dynamic programming optimization and pick-up forecasting with booking curve unconstraining.
3. A single market with the unrestricted fare structure provided in Fig 3.1.11. Both airlines use EMSRb optimization and Q-forecasting and booking curve unconstraining.







































		08:00	13:00	17:00		
 origin	<p>AL4  60 seats</p> <p>AL3  60 seats</p> <p>AL2  60 seats</p> <p>AL1  60 seats</p>	<p> 60 seats</p> <p> 60 seats</p> <p> 60 seats</p> <p> 60 seats</p>	<p> 60 seats</p> <p> 60 seats</p> <p> 60 seats</p> <p> 60 seats</p>	 destination	Competitive scenario : 4 airlines	
 origin	<p>AL3  80 seats</p> <p>AL2  80 seats</p> <p>AL1  80 seats</p>	<p> 80 seats</p> <p> 80 seats</p> <p> 80 seats</p>	<p> 80 seats</p> <p> 80 seats</p> <p> 80 seats</p>	 destination	Competitive scenario : 3 airlines	
 origin	<p>AL2  120 seats</p> <p>AL1  120 seats</p>	<p> 120 seats</p> <p> 120 seats</p>	<p> 120 seats</p> <p> 120 seats</p>	 destination	Competitive scenario : 2 airlines	
 origin	AL1  240 seats	 240 seats	 240 seats	 destination	Monopoly scenario	

Fig 3.1.9 Monopoly scenario vs. competitive scenarios in a single market

Fare Class	Fare	Advance purchase (days)	Saturday Night Stay	Change Fee	Non Refundable
1	500	0	No	No	No
2	400	3	No	No	Yes
3	315	7	No	Yes	Yes
4	175	10	Yes	Yes	Yes
5	145	14	Yes	Yes	Yes
6	125	21	Yes	Yes	Yes

Fig 3.1.10 Fully-restricted fare structure considered

Class	Fare	Advance purchase (days)	Saturday Night Stay	Change Fee	Non Refundable
1	500	0	No	No	No
2	400	0	No	No	No
3	315	0	No	No	No
4	175	0	No	No	No
5	145	0	No	No	No
6	125	0	No	No	No

Fig 3.1.11 Unrestricted fare structure considered

3.1.3 Simulation results

Scenario 1: Fully-restricted fare structure, EMSRb optimization

The first scenario tested is kept as close as possible to the hypotheses made by the theoretical results, and especially Proposition B. The assumption of independent classes made in the literature and our propositions A and B is best justified with fully-restricted fare structures, although in reality sell-up still exists. EMSRb optimization is used (cf. Talluri and van Ryzin [33] for a detailed description). We keep the forecaster simple, with pick-up forecasting and with booking curve unconstraining (cf. Hopperstad[17] for a detailed description).

We first run simulations where total daily market non-constrained demands are equally distributed among the competing airlines. Fig 3.1.12 shows the sums of the average values of the forecasts of unconstrained demands made by airlines for the 8am flights, 24 days before departure. Fig 3.1.13 shows the associated average total market nested booking limits. Fig 3.1.14 shows the average total bookings by fare class for a day of departures. Fig 3.1.15 shows average total market daily revenues, passengers and load factors.

We find that in the competitive scenario the sum of the average values of the airline unconstrained forecasts is not always higher than the expected value of the forecast made by the equivalent monopoly. It is true for the low classes 4 and 5, but it is not for the high classes. In this competitive scenario, total market nested booking limits are smaller than under the monopoly scenario only for the low classes 6 and (5, 6). Still, under competition the sum of the airline unconstrained forecasts for all classes is higher than the unconstrained forecast for all classes made by the monopoly. The difference in forecasts by class between the monopoly case and the competition case is much larger for the low classes than for the high classes. Under competition total market bookings in the highest fare classes are slightly smaller, total market bookings in the middle classes are higher and total market bookings in the lowest classes are lower. Because of this difference in total market fare class mix, the oligopoly generates higher total market revenues than an equivalent monopoly through higher yields and lower load factors, as shown in Fig 3.1.15. The oligopoly gets fewer total average passengers, but has a higher average fare.




Fig 3.1.12 Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.13 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.14 Average total bookings by fare class for a day of departures




Fig 3.1.15 Average total daily market revenues, passengers and load factors

The PODS simulator allows us to check that this noticed behavior of competitive environments as compared to monopolies is actually due to spill-in and sell-up of demand. The PODS simulator can be run in a “First Choice Only Choice” (FCOC) mode, in which if a passenger does not get his first choice for a fare class, he will not spill-out to another airline or sell-up to another fare class. This mode does not model realistic situations, but helps us estimate the impact of spill-in and sell-up. In such a mode, the non-constrained demand for a competing airline is only its first-choice non-constrained demand. Fig 3.1.16 through Fig 3.1.19 show the results of the “First-Choice-Only-Choice” simulations. We no longer find the same results as above.

In an oligopoly:

- For all classes total market unconstrained forecasts are higher.
- For all classes total market booking limits are higher.
- For all classes overall total market bookings slightly decrease.
- Both total market revenues and load factors decrease.

First, this confirms that spill-in and sell-up of demand are responsible for the smaller booking limits for the low classes observed in the competitive case as compared to the monopoly case. In the competitive case, spill-in and sell-up can be accounted for 40% of the total market forecasts for the set of classes (1,2,3,4,5).

Second, this suggests that the behaviour for the high classes found with spill-in and sell-up which are contrary to the theory exposed both in the literature and 3.1.1 could be caused by another factor than spill-in and sell-up. An in depth study of the simulation data reveals that total market unconstrained forecasts for the high classes are smaller under competition than under monopoly because of an imprecision of booking-curve unconstraining modules of forecasters that is larger than the value of spill-in and sell-up. Even if spill-in and sell-up should theoretically slightly lead to higher total market forecasts for high classes, imperfections of the unconstrainers lead to the opposite.

Last, FCOC simulations also suggest that if there was neither spill-in nor sell-up, oligopolies would not generate as much total market revenues as a monopoly. An explanation is that the accuracy of the forecast of individual members of the oligopoly is not as good as the accuracy of the forecast by the monopoly. The monopoly benefits from more data, its aggregate forecast should be more precise. The more airlines in the oligopoly, the less accurate the individual airlines’ forecasts, further affecting total market revenues.




Fig 3.1.16 FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.17 FCOC - Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.18 FCOC - Average total bookings by fare class for a day of departures




Fig 3.1.19 FCOC - Average total daily market revenues, passengers and load factors

We can check that the same results are found whatever the initial allocation of total market first-choice non-constrained demand among the competing airlines. The former results were found with an equal distribution of demand among the competing airlines. For instance in the duopoly case, on average one half of the passengers preferred one airline, and the other half preferred the other airline.

We compare below the following cases:

- A monopoly situation
- A duopoly situation with initial distribution of non-constrained demand of 1/2 vs. 1/2. 50% of the people prefer airline 1 all else equal.
- A duopoly situation with initial distribution of non-constrained demand of 1/3 vs. 2/3. 33% of the people prefer airline 1 all else equal.
- A duopoly situation with initial distribution of non-constrained demand of 1/4 vs. 3/4. 25% of the people prefer airline 1 all else equal.

with capacities remaining the same as shown in Fig 3.1.9.

Results suggest that the distribution among airlines of first-choice non-constrained demand has naturally an impact on the individual performance of airlines. Yet at the scale of the total market, the differences between a duopoly and a monopoly enumerated above hold whatever the initial distribution of non-constrained demand among airlines. Fig.3.1.20 through Fig.3.1.23 compare the total market statistics with respect to the initial distribution of non-constrained demand among airlines in a duopoly. “1/2 vs. 1/2” refers to passengers equally preferring the two competitors. “1/3 vs. 2/3” refers to 1/3 of the passengers preferring one airline, and 2/3 preferring the other. “1/4 vs. 3/4” refers to 1/4 of the passengers preferring one airline, and 3/4 preferring the other.




Fig 3.1.20 Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.21 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.22 Average total bookings by fare class for a day of departures




Fig 3.1.23 Average total daily market revenues, passengers and load factors

Scenario 2: Fully-restricted fare structure, Lautenbacher dynamic programming

EMSRb is a heuristic widely used in airline revenue management systems. Although not widely used in practice, dynamic programming is another optimization technique widely studied for possible real-world application. We therefore run the same scenario as scenario 1 but with Lautenbacher dynamic programming optimization (a detailed description can be found in Lautenbacher and Stidham [21]) to see whether theoretical results still hold even with dynamic programming optimization.

We first run simulations where total daily first-choice non-constrained demands are equally distributed among the competing airlines. Fig 3.1.24 shows the average sum of the average values of the unconstrained forecasts made by airlines for the 8am flights, 24 days before departure. Fig 3.1.25 shows the associated average total market nested booking limits. Fig 3.1.26 shows the average total bookings by fare class for a day of departures. Fig 3.1.27 shows average total daily market revenues, passengers and load factors.

The results are very similar to when airlines use EMSRb. In the competitive scenario the sum of the average values of the airline unconstrained forecasts is lower than in the monopoly scenario for classes 1,2,3,4. In the competitive scenario, total market nested booking limits are smaller than under the monopoly scenario only for the low classes 6 and (5,6). Under competition total market bookings in the highest fare classes is slightly smaller, total market bookings in the middle classes is higher and total market bookings in the lowest classes is lower than under the monopoly. Yet with dynamic programming, the oligopoly does not always generate higher total market revenues than an equivalent monopoly, as shown in Fig 3.1.27.




Fig 3.1.24 Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.25 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.26 Average total bookings by fare class for a day of departures




Fig 3.1.27 Average total daily market revenues, passengers and load factors

Running the same scenarios in a “First Choice Only Choice” provides the same conclusions with dynamic programming than with EMSRb. Fig 3.1.28 through Fig 3.1.31 show the results of the “First-Choice-Only-Choice” simulations.

In the oligopoly with FCOC mode:

- For all classes total market unconstrained forecasts are higher.
- For all classes total market booking limits are higher.
- For all classes overall total market bookings slightly decrease.
- Both total market revenues and load factors decrease.

First, in the competitive case, spill-in and sell-up account for 35% of the total market forecasts for the set of classes (1,2,3,4,5). This confirms that the spill-in and the sell-up of demand are responsible for the smaller booking limits for the low classes observed in the competitive case as compared to the monopoly case.

Second, we find again that even if spill-in and sell-up should theoretically lead to higher total market unconstrained forecasts for high classes, imperfections of forecasters lead to the opposite.

Last, FCOC simulations also suggest that if there was neither spill-in nor sell-up, oligopolies would not generate as much total market revenues as a monopoly because of less accurate forecasts.

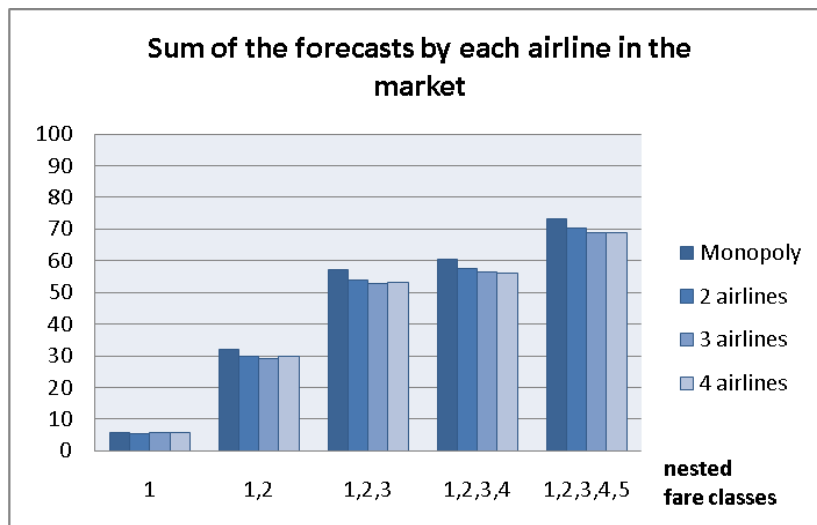


Fig 3.1.28 FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure



Fig 3.1.29 FCOC - Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.30 FCOC - Average total bookings by fare class for a day of departures




Fig 3.1.31 FCOC - Average total daily market revenues, passengers and load factors

Eventually we check that the same results are found whatever the initial allocation of first-choice non-constrained demand among the competing airlines.

We compare again the following cases:

- A monopoly situation
- A duopoly situation with initial distribution of non-constrained demand of 1/2 vs. 1/2
- A duopoly situation with initial distribution of non-constrained demand of 1/3 vs. 2/3
- A duopoly situation with initial distribution of non-constrained demand of 1/4 vs. 3/4

with capacities remaining the same as shown in Fig 3.1.9.

Results suggest again that the distribution among airlines of first-choice non-constrained demand has naturally an impact on the individual performance of airlines. Yet once more no matter how the total market non-constrained demand is initially spread among the competitors, we find the same overall behaviour of the duopoly as compared to the monopoly.

Fig.3.1.32 through Fig.3.1.35 compare the total market statistics with respect to the initial distribution of non-constrained demand among airlines in a duopoly.




Fig 3.1.32 Average sum of airline forecasts, 8am flight, 24 days prior departure



Fig 3.1.33 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.34 Average total bookings by fare class for a day of departures




Fig 3.1.35 Average total daily market revenues, passengers and load factors

Scenario 3: Unrestricted fare structure, EMSRb optimization, Q forecasting

The assumption of independent classes made in the literature and propositions A and B is best justified in fully-restricted structure environment. We run here simulations of an unrestricted fare structure to check whether as mentioned in 3.1.1 results should hold with important sell-up. In this test, both airlines use EMSRb optimization (a detailed description can be found in Talluri and van Ryzin [33] p. 47-50) and Q-forecasting (a precise description can be found in Cleaz Savoyen[7]) and booking curve unconstraining (a detailed description can be found in Hopperstad[17]).

Fig 3.1.36 through 3.1.39 display results of the simulations of the unrestricted fare structure scenario with first-choice non-constrained demands equally distributed over airlines. For all classes the sum of the expected values of airline unconstrained forecasts in competition is larger than the expected value of the unconstrained forecast made by the monopoly. This causes total market booking limits to be smaller under competition than under monopoly whatever the fare class. The lower the fare class, the larger the proportional decrease in booking limits between the monopoly case and the oligopoly case. Because of such booking limits, the oligopoly has fewer bookings in the lowest fare class, but higher bookings in all the upper classes. That way the oligopoly is able to generate more revenues through lower load factors and a higher yield than the monopoly.

The difference with theory for high classes due to the imperfection of the unconstraining module of forecasts is not found in this case. This is because with Q-forecasting, forecasts for the highest classes are computed as a given proportion of a total demand for all classes. This total demand for all classes being in many cases computed with historical bookings for low classes, for which the error of the forecaster is smaller than spill-in and sell-up. In that case, whatever the fare class, the error of the forecaster is hence smaller than spill-in and sell-up, so that forecasts are higher under competition than under monopoly.




Fig 3.1.36 Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.37 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.38 Average total bookings by fare class for a day of departures




Fig 3.1.39 Average total daily market revenues, passengers and load factors

Fig 3.1.40 to Fig 3.1.42 show the results of the equivalent FCOC simulation. We find the opposite results than without FCOC. In competition, total market unconstrained forecasts have a lower expected value than in a monopoly. Booking limits are higher for all classes in competition, and revenues lower. These FCOC simulations show us that the spill-in, completed by sell-up, is responsible for the differences in booking limits between oligopoly and monopoly: If there was neither spill-in nor sell-up, the oligopoly would generate lower total revenues due to less accurate individual airlines' forecasts made with smaller samples than the monopoly.




Fig 3.1.40 FCOC - Average sum of airline forecasts, 8am flight, 24 days prior departure



Fig 3.1.41 FCOC - Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.42 FCOC - Average total daily market revenues, passengers and load factors

Similar results are found when the initial allocation of first-choice non-constrained demand among the competing airlines is not even. No matter the tested allocation of non-constrained demand between airlines, total market unconstrained forecasts are higher and total market booking limits are smaller under competition, moving the fare class mix towards upper classes, and generating more revenues through smaller load factors and higher yields. Fig.3.1.43 through Fig.3.1.46 compare the total market statistics with respect to the distribution of first-choice non-constrained demand among airlines in a duopoly.




Fig 3.1.43 Average sum of airline forecasts, 8am flight, 24 days prior departure




Fig 3.1.44 Average total market booking limits, 8am flight, 24 days prior departure




Fig 3.1.45 Average total bookings by fare class for a day of departures




Fig 3.1.46 Average total daily market revenues, passengers and load factors

3.1.4 Conclusions

We extended the theoretical results of the literature stating that all else equal, an oligopoly sets higher total market protection limits for classes than an equivalent monopoly would do. We first introduced our terminology and our model, and provided a general intuitive reasoning that is valid for N airlines with n classes and with dependent non-constrained demands among classes. We then proved the result with different more restrictive assumptions. We finally performed PODS simulations to validate the general intuitive idea and to examine whether the proved results can be extended to more complex real world situations. We found that in a single market, when airlines have the same fully-restricted or unrestricted fare structure, use the same forecasting and optimization method:

1. Theoretical results apply to more realistic cases where airlines use optimization methods such as EMSRb or dynamic programming, where non-constrained demands for classes are not independent, and whatever the distribution of first-choice non-constrained demand among competitors.
2. If forecasters were perfect, total market unconstrained forecasts would be higher under competition than under a monopoly. Total market protection levels set by an oligopoly would therefore be higher than the ones set by a monopoly for the same total market seat capacity.
3. Yet this theoretical result is not always verified in practice if the forecasters (and especially the unconstraineders) happen to have an imprecision larger than spill-in and additional sell-up. This is more prone to be the case for fare classes with small spill-in and additional sell-up, which are the highest classes.
4. In the simulations run, when oligopolies set higher protection levels than monopolies, it is often associated with an increase in revenues through lower load factors and higher yields. This stems from the fact that the used revenue management methodologies used might be slightly under-aggressive as compared to the optimum. If there was neither spill-in nor additional sell-up, total market revenues generated by the oligopoly would be lower than the revenues generated by an equivalent monopoly, because individual airlines' forecasters would not be as accurate as the monopoly's due to smaller historical samples.

3.2 Competitive Interactions with Current Airline Revenue Management

In 3.1 we distinguished a competitive environment from a monopolistic environment as regards current airline revenue management systems, both with theoretical results and with more complex PODS simulations. We made conclusions on the impact of competition at a total market level. In this part we focus on the impact of competition at the level of a single airline. We describe the passive competitive interactions between current airline revenue management systems. We first examine the forecasting modules that provide unconstrained estimates of future non-constrained demands. We then concentrate on optimization modules which use these forecasts to set booking limits. We try to keep our theoretical reflection as close as possible to real world cases. We do not run specific simulations, but we link our findings to the results of the PODS simulations described in the Literature Review (Part 2.1).

3.2.1 Forecasting

In practice most airline automated forecasters rely on raw historical bookings data only and do not take into account the competitive environment. Revenue management systems use huge amounts of data so forecasting methods have usually been kept simple to insure short run times. The time and financial investment necessary to improve forecasts is often not as justified as improving other facets of airline revenue management. Therefore the majority of current automated forecasting methods simply extrapolate raw past bookings and try to best estimate their logical sequel. They do not try to differentiate spill-in from first-choice demand. They do not balance past bookings with the competitive environment in which they were made. Neither do they balance future forecasts with the current competitive environment.

Airline revenue management in practice mostly relies on human analysts to take into account the competitive environment. Analysts override the forecasts to adapt them to specific market conditions. It is in fact wiser for a human analyst to modify the forecasts rather than the booking limits: an experienced analyst knowing market conditions can make a better estimate of demand than a computer, but rarely can he optimize booking limits better than a computer.

Automatically incorporating knowledge on competition presents two major obstacles. The first one is the nature of the data about the competitors that is available to airlines. Knowing the competitive environment consists of recording the class availabilities and the fare level of the competitor. Obtaining such data is not easy. Exact inventories, past and current bookings are private data. Remaining availabilities and associated fares are public data. Some of the private data could be guessed through reverse engineering of the public data. Yet it would be possible for an airline to confuse such reverse engineering from its competitors. As a consequence, **as of today the airline revenue management competitive game cannot be considered as a complete information game**. In this thesis we use the term “**LOCO**” to designate the lowest available class of the competition at a given time. This information on the competitor’s availability is imprecise but easily accessible by airlines.

The second and more difficult problem with incorporating competitive environments in the automated forecasts is that future competitor moves are unpredictable. The competitor analysts could override their own system to implement specific tactics modifying high fare protection. But competitors might not also act in order to directly optimize revenues. Overriding systems to open more low classes is a marketing strategy used by legacy carriers to maintain market share (cf. Lua [28]). Besides in the area of airline revenue management and pricing, it is possible but not proven that fares and availabilities could be used as a means of communication between airlines. It is conceivable that some fare/availability changes in an origin-destination market could be measures of retaliations or warnings relative to competitive moves in other origin-destination markets. A limited possibility for forecasting is to suppose that the competitor analysts will not override their system and that the competitor will rely exclusively on its automated system. In that case it would be possible to estimate the future behavior of the competitor knowing its past demand behavior.

Forecasts can actually reflect competitive moves if these moves are sustained over the period of time covering the historical data sample used to generate a forecast. In that case only, the extrapolation of the historical data sample will reflect the competitor move. We call this characteristic period of time **the inertia of a forecaster**. If a competitive move is sustained over a period at least as long as the inertia of the forecaster, the forecast will have adapted with a delay equal to its inertia. If a competitive move is not sustained long enough, the extrapolation of historical data will not reflect this move. In this case, only an adjustment of the forecast by a human analyst could make it reflect the short-term competitive move.

3.2.2 Optimization of the Booking Limits

The airline revenue management optimization models published until now have not explicitly taken into consideration competitive feedback effects. Yet because of the inertia of forecasters, any revenue management system that does not purposely take into account competitive interactions for optimization can still have a passive response to a long-term change in the booking limits by its competitors. As shown in Fig 3.1.8 and 3.1.8bis, by changing their booking limits, competitors modify spill. If this modification of booking limits is done over a long enough period of time, the other airline will detect a change in its accepted bookings, which eventually will lead it to change its own booking limits.

The problem of competitive revenue management interactions between two airlines can be expressed as the following game:

- Two competing airlines offering n equivalent fare classes.
- The strategy of airline j consists of a set of n nested protection levels $(X_{1..i}^j)_{i=1..n}$
- The payoff of the combination of strategy $(X_{1..i}^1)_{i=1..n}$ and strategy $(X_{1..i}^2)_{i=1..n}$ is the associated expected revenues π^1 and π^2 .

Because airline revenue management systems can have delayed automated competitive reactions, it is important to distinguish two different cases when studying this revenue management game:

- Games with strategy moves lasting less than the inertia of the forecasters.
(We will call these **short-term moves/short-term games**)
- Games with moves that last at least as long as the inertia of the forecasters.
(We will call these **long-term moves/long-term games**)

Before going into the details of long-term and short-term competitive reactions between airline revenue management systems, we want to point out that the more an airline is preferred, the smaller the ratio between its observed spill-in and its first-choice non-constrained demand. The less an airline is preferred, the higher the ratio between its observed spill-in and its first-choice non-constrained demand. This means that an airline that is initially preferred is less sensitive/vulnerable to the revenue management decisions of its competitors. All else equal, it is reasonable to believe that a legacy carrier is preferred to a low-cost airline. Hence in terms of seat inventory control a low-cost airline is more vulnerable/sensitive to the seat inventory decisions of legacy carriers than the opposite.

Furthermore, as seen in part 3.1, the imprecision of forecasters could in some cases dominate spill-in. In part 3.2, we showed that spill-in could account for as high as 40% of the demand for an airline and a class. All of the following reasoning assumes that the competitive effects on demand are larger than the imprecision of forecasters.

Passive competitive reactions to long-term moves

For long-term revenue management games, the notion of Nash Equilibrium is weak. If one airline makes a unilateral long-term move, this will cause an automatic reaction from its competitor. For airline 2's strategy to remain unchanged with a unilateral long-term move by airline 1, airline 2 must actually intervene and override its revenue management system. Therefore even if a unilateral move is not revenue positive for its instigator, it could still want to make it, expecting the automated reaction from its competitor that could ultimately lead it to better revenues. Airlines could possibly have an incentive to move unilaterally from a Nash Equilibrium. Conversely if its competitor makes a long term move, an airline will have the choice between letting its automated revenue management system choose the new strategy to adopt, or overriding this system and choosing "manually" a new long-term move. Note that the response to a long-term move must be a long-term move, because the effect of a short term move is then negligible.

Proposition D presents a result on long-term moves. Proposition C introduces a result on the spill-in generated by the EMSRb method useful for the proof of proposition D.

Proposition C:

Suppose:

- An airline optimizes a set of classes (1..i) using EMSRb.
- Non-constrained demands for classes follow a normal distribution truncated at 0 (demands are positive).
- Proportional changes in the average revenue weighted by class non-constrained demands for the set (1..i) are negligible.

Let:

- $Spill : d \rightarrow Spill(d)$ be the application that associates a stochastic non-constrained demand d for a set of classes of an airline with the stochastic spill-out it generates using EMSRb.

Then:

- $Spill(d)$ is also a stochastic variable following a normal distribution truncated at 0.
- The distribution of $Spill(d)$ is a function of the standard deviation of d only.
- The expected value and the standard deviation of $Spill(d)$ are increasing in the standard deviation of d .

Proof:

In that case the protection level set by the airline can be written $X = E[d] + \sigma[d].constant$. The distribution of $Spill(d)$ is then the distribution of $d - X = d - E[d] - \sigma[d].constant$ truncated in 0. d follows a normal distribution truncated in 0. Therefore $Spill(d)$ has a distribution equal to $\sigma[d]$ ($SND + constant$) truncated in 0 where SND is the standard normal distribution. Hence the results.

End of Proof

Proposition D

Suppose:

- 2 airlines, choosing their protection level $X_{1..i}^j$ for the set of classes (1..i)
- Non-constrained demands for classes are independent, follow a steady normal distribution truncated in 0
- Nested booking limits are reached
- Forecasters are accurate (unconstrained demand is equal to non-constrained demand)

Then:

- In most of the real-world cases, the long-term competitive response of an airline using EMSRb is negative:
 - When airline 2 makes a long term move where it increases a protection level, this decreases the expected value and the standard deviation of the associated spill-in observed by airline 1. With delay, airline 1 using EMSRb will decrease its equivalent protection level. According to proposition C the expected value and the standard deviation of the spill-in towards airline 2 are decreased. The opposite is true when airline 2 makes a long term move decreasing its protection level.
 - The feedback effect changing the spill-in towards airline 2 reinforces the first-order effect on airline 1.

Proof:

Keeping the notations from 3.1.1, to optimize its revenues with an EMSRb rule, airline 1 would choose the booking limit for a set of classes ($I..i$) so that:

$$(3.1) \quad \Pr(D_{1..i}^1 \geq X_{1..i}^1) = \Pr(d_{1..i}^1 + \varepsilon_{1..i}^1 \geq X_{1..i}^1) = \frac{P_{i+1}^1}{\sum_{k=1..i} p_k^1 (d_k^1 + \varepsilon_k^1)} = \frac{P_{i+1}^1}{\sum_{k=1..i} (d_k^1 + \varepsilon_k^1)}$$

When airline 2 increases its protection level, it affects spills in the following way:

$$\frac{\partial E(\varepsilon_{1,i}^1)}{\partial X_{1,i}^2} \leq 0 \text{ and } \frac{\partial \sigma(\varepsilon_{1,i}^1)}{\partial X_{1,i}^2} \leq 0$$

This implies that in “*non-extreme cases*” the passive EMSRb response $\frac{\partial X_{1,i}^1}{\partial X_{1,i}^2}$ is negative.

“*Extreme cases*” include cases where the ratios between the fares of the fare classes are very close to 1, cases where the proportional change in the average fares weighted by non-constrained demands is not negligible, cases where the first-choice non-constrained demands have very small standard deviation etc.

The feedback change in the spill-in towards airline 2 can be deduced from proposition C, and reinforces the effect of the initial move by airline 2.

End of Proof

Note that as opposed to 3.1.1, the result stating that the competitive response $\frac{\partial X_{1,i}^1}{\partial X_{1,i}^2}$ of an airline using EMSRb is negative does not require that the airlines have the same fares for competing classes nor that first-choice non-constrained demands for each airline be a constant fraction of a given total market non-constrained demand. Also note that this result remains valid with $N > 2$ airlines if we replace the change in the protection level by airline 2 by the overall change in the sum of the $N-1$ competitors’ booking limits.

Suppose Airline A is using EMSRb and that stochastic first-choice non-constrained demands are steady. If as a long-term move, its competitor Airline B increases the protection level for the set of classes (1..i), proposition D suggests that in the long term Airline A will fix its protection level for (1..i) using EMSRb with a lower level of total non-constrained demand. In that case, the expected total revenues of Airline A should decrease, and its spill-out to Airline B is reduced. Fig 3.2.1 illustrates this passive response from Airline A with four successive steps. Step 1 is the first move by Airline B. Step 2 is the consequence of this move on spill-in. Step 3 is the passive response from Airline A. Step 4 is the consequence on spill-in of this move by Airline A. Note that step 4 just reinforces step 2.

Suppose Airline A is using EMSRb and that stochastic first-choice non-constrained demands are steady. If as a long-term move, its competitor Airline B decreases the protection level for the set of classes (1..i), proposition D suggests that in the long term Airline A will fix its protection level for (1..i) using EMSRb with a higher level of total non-constrained demand. In that case, the expected revenues of Airline A should increase, as well as its spill-out to Airline B. Fig 3.2.2 illustrates this passive response from Airline A with four successive steps. Step 1 is the first move by Airline B. Step 2 is the consequence of this move on spill-in. Step 3 is the passive response from Airline A. Step 4 is the consequence on spill of this move by Airline A. Note that step 4 just reinforces step 2.

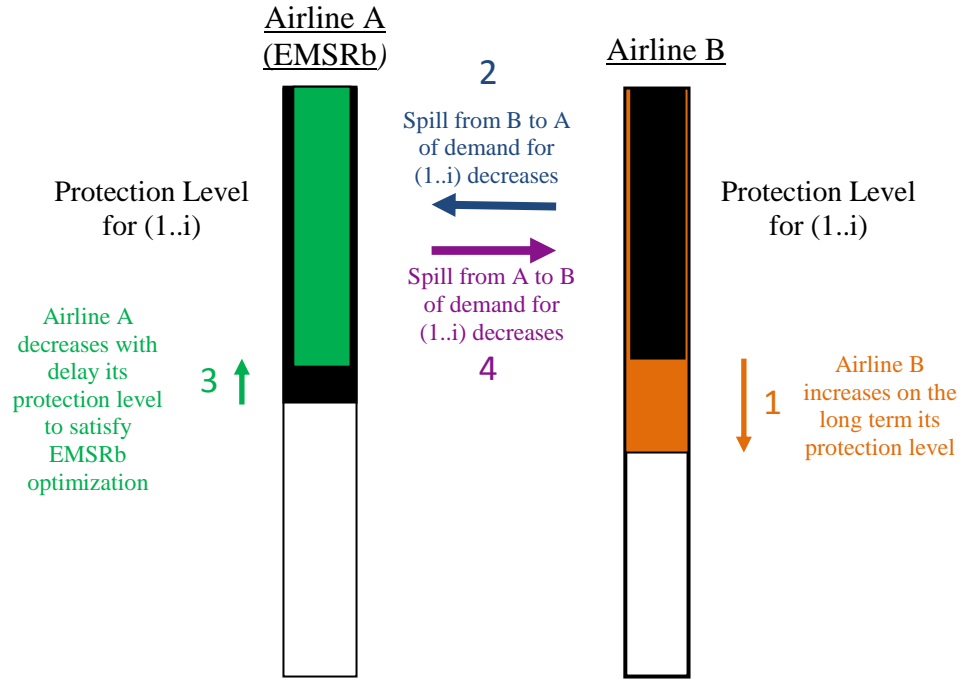


Fig 3.2.1 Passive competitive reaction of Airline A's EMSRb revenue management system when Airline B decreases its booking limit on the long term

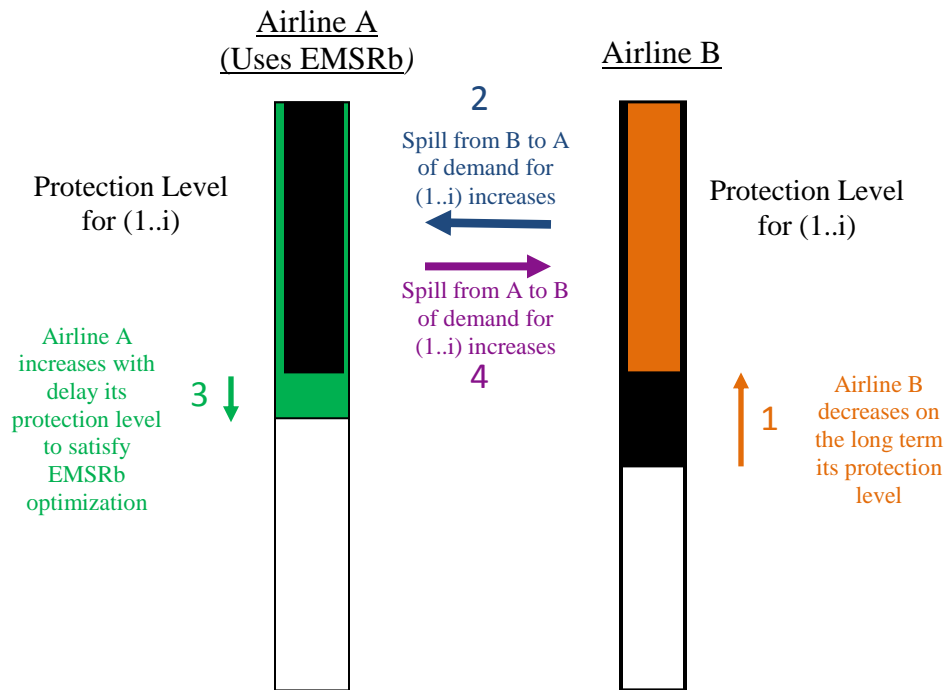


Fig 3.2.2 Passive competitive reaction of Airline A's EMSRb revenue management system when Airline B increases its booking limit on the long term

These theoretical results on passive reactions of revenue management systems to long-term competitive moves explain the results found in part 2.1. Simulations from the literature suggested that if an airline is using EMSRb and its competitor improved its revenue management system, it would decrease the passive airline's revenues. Improving one's revenue management system often means higher protection levels and fewer spill-in of high classes to the competitor. Our result indicates that in that case, the passive airline using EMSRb gets smaller non-constrained demand for high classes. As a consequence it will decrease its protection levels to optimize revenues but will have lower total expected revenues than before the competitor made its move. All the past simulations of competitive environments as described in 2.1 indicate that this remains true whatever the revenue management method used by the passive airline. If its competitor improves its revenue management method by lowering its booking limits, the passive airline will get lower non-constrained demands for high classes, and consequently will increase its booking limit to reoptimize its revenues, but will end up in a worse-off situation after all. Similarly, when an airline catches up with the better revenue management of its competitor, the competitor sees its revenues decrease as compared to when he had the advantage of the revenue management method. Nevertheless each airline's revenues are then higher than when no airline had implemented this better revenue management method. Fig 3.2.3 illustrates our point by showing the changes in booking limits and revenues of airlines due to airlines implementing better revenue management.

	AirlineB: no move	AirlineB: better RM		
	Airline A revenues	Airline B revenues	Airline A revenues	Airline B revenues
AirlineA: no move	0	0	-	++
AirlineA: better RM	++	-	+	+

↘ (- +)

→ (+ -)

Fig 3.2.3 Revenue implications of the passive interactions of revenue management systems when an airline implements a better form of revenue management. Arrows represent the change in booking limits set by the airlines at the equilibrium.

These results also confirm the results by d'Huart and Belobaba [10] that when a carrier open-matches its competitor, it leads to better revenues for the latter. By open-matching its competitor, an airline consistently decreases its protection levels for higher classes, leading to a long-term increase in its competitor's non-constrained demand for higher classes, in its protection levels, and in its revenues.

Moreover we can deduce from these results on passive competitive reactions of EMSRb systems that when two airlines use EMSRb to set their booking limit for a given set of classes (1..i), they should on the long term reach a stable equilibrium. We can use a graphical reasoning as a demonstration, and an example is provided in Fig 3.2.4. Former results indicate that for given initial demands, the EMSRb protection level for (1..i) and the EMSRb revenues by one airline are decreasing functions of the protection level for (1..i) set

by the competitor. Besides the EMSRb protection level of an airline is strictly positive when the competitor is protecting its entire capacity because there is an initial demand specific for this first airline. This implies that if we draw the two EMSRb reaction curves in a same graph where axes are the protection levels set by the airlines, the two curves intersect. Because the two curves intersect, there is at least one stable equilibrium of protection levels for (1..i). In Fig 3.2.4 there is one stable equilibrium. If the EMSRb reaction curves with respect to the competitor's protection level have no inflexion point, then there is necessarily a unique equilibrium. In this thesis we do not show whether reaction curves have an inflexion point. But our intuition is that in most real world cases, there is no inflexion point, and there is a unique stable equilibrium.




Fig 3.2.4 Graphical representation of the case of two competing airlines using EMSRb

Passive competitive reactions to short-term moves

For short-term revenue management games, the notion of Nash equilibrium holds. If an airline makes a unilateral short-term move, by definition the competitor's automated revenue management system will not react to this short-term move. Fig 3.2.5 and Fig. 3.2.6 illustrate the competitive impact of a short-term competitive move.

Suppose that first-choice non-constrained stochastic distributions of demands are steady. If as a short-term move, Airline B increases a protection level for the set of classes (1..i), the revenue management system of Airline A will not detect this move by Airline B. Airline B will decrease the expected value and the standard deviation of the spill-in and therefore of the non-constrained demand observed by Airline A for (1..i). In return Airline A will spill-out less to Airline B than if Airline B had made no move. As Airline A is not changing its protection levels, if it is using EMSRb optimization, it is likely to be over-protecting. This should result in lower expected revenues for Airline A for the departure dates where Airline B is making the short term move.

Suppose that first-choice non-constrained stochastic distributions of demands are steady. If as a short-term move, Airline B decreases the protection level for the set of classes (1..i), the revenue management system of Airline A will not detect this move by Airline B. Airline B will increase the expected value and the standard deviation of the spill-in and therefore of the non-constrained demand observed by Airline A for (1..i). In return Airline A will spill-out more to Airline B than if Airline B had made no move. As Airline A is not changing its protection level, if it is using EMSRb optimization, it is likely to be under-protecting. Still its mix of demand is improved. It should result in higher expected revenues for Airline A for the departure dates where Airline B is making the short term move.

Note that whatever the scenario, the average daily revenues of the passive airline are lower after a short term move than after an equivalent long term move followed by a reaction by the EMSRb optimization system. Also note that all of the above results require the assumption that booking limits are reached whatever the move made by the competitor.

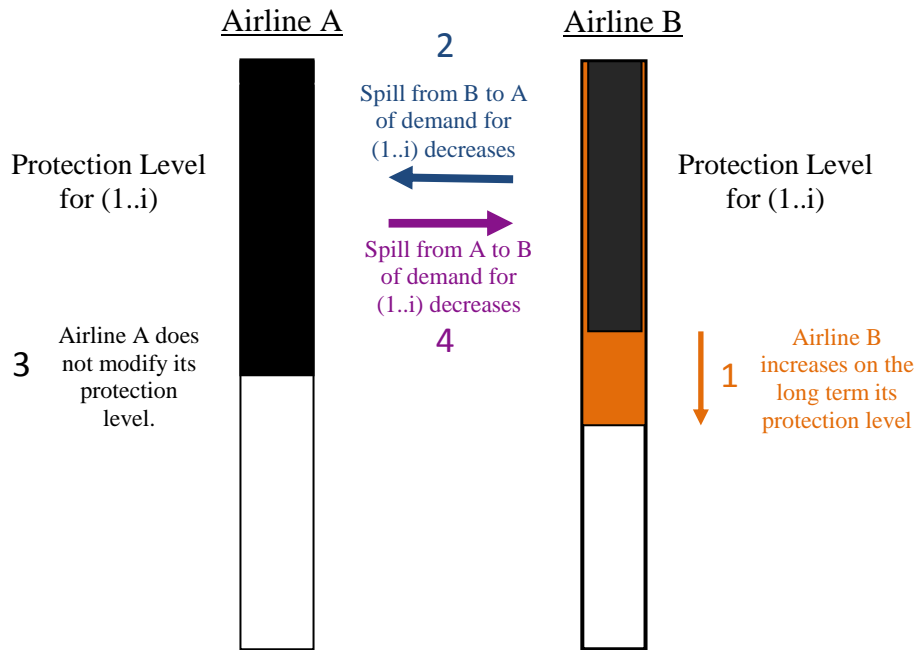


Fig 3.2.5 Competitive interaction, Airline B decreases a booking limit on the short term




Fig 3.2.6 Competitive interactions, Airline B increases a booking limit on the short term

Extension of the results to cases with dependent demands among classes

We believe that a dependence of the non-constrained demands between classes can attenuate the amplitude but not the nature of the preceding results. Suppose there exists sell-up between classes. Our justification is the following:

- If Airline B increases its protection level for classes $(1..i)$, it is going to increase its sell-up from classes $(i+1..n)$ to $(1..i)$. But this increase in sell-up is necessarily smaller or equal to the change in the protection level. Hence the spill-in from Airline B to Airline A for classes $(1..i)$ should still decrease in expected value and standard deviation because of the increase of the protection level of Airline B. If Airline B's move is sustained on the long term, it will lead Airline A to decrease its protection level. This will decrease Airline A's sell-up from classes $(i+1..n)$ to $(1..i)$, which will not lead Airline A to increase its protection level.

- The same reasoning can be applied if Airline B decreases its protection level.

3.2.3 Conclusions

We define the airline revenue management competitive game as:

- Two competing airlines offering n equivalent fare classes.
- The strategy of airline j consists of a set of n nested protection levels $(X_{1..i}^j)_{i=1..n}$
- The payoffs of the combination of strategy $(X_{1..i}^1)_{i=1..n}$ and strategy $(X_{1..i}^2)_{i=1..n}$ are the associated expected revenues π^1 and π^2 .

A revenue management competitive move is a change in a nested protection level for given steady distributions of demands.

Automated forecasters in airline revenue management systems have usually been kept simple and extrapolate past data without balancing it with past competitive situations or current competitive situations. They adapt to competitive moves but with a certain delay that we call the inertia of a forecaster. The following conclusions require that the impact of competitor moves on demand be larger than a possible imprecision of forecasters.

Competitor moves that last for a duration shorter than the forecaster inertia are not taken into account by the forecaster or by the optimizer. The passive airline does not modify its booking limits. On average, it experiences higher revenues and spills-out more for the associated high classes if the competitor decreases a protection level. It will experience lower expected revenues and will spill-out less for the associated high classes if the competitor increases a protection level.

Competitor moves that last at least as long as the inertia of the forecaster are taken into account. Suppose a passive airline uses an automated revenue management system. If the competitor increases a booking limit, this should cause the passive airline to get more first-non-constrained demand for the associated high classes, to decrease its equivalent booking limits, and to get more revenues. If the competitor decreases a booking limit, this will cause the passive airline to get less non-constrained demand for the associated high classes, to increase its equivalent booking limit, and to get less revenues. Therefore the protection level for a set of classes (1..i) and the revenues set by the passive airline with its optimization method are decreasing functions of the competitor's protection level for the set of classes (1..i). Furthermore, if both airlines have implemented an automated optimization method, they should reach a stable equilibrium situation.

These results on the long-term passive competitive reactions of airline revenue management systems are demonstrated in reasonable scenarios under the assumptions of independent non-constrained demands by classes and reached booking limits for airlines using an EMSRb method. They should still hold but be attenuated even with dependent demands for classes, when there exists sell-up between classes. They match observations made on competitive interactions of airline revenue management systems in former complex PODS simulations described in the literature review. The rest of the former PODS simulations described in the literature review seem to allow the extension of the results to other automated optimization methods than EMSRb.

Last but not least airlines preferred to their competitor all else equal are less sensitive to their competitive environments. This suggests that in the field of revenue management, low-cost carriers are all else equal more sensitive to moves by legacy carriers than the opposite.

CHAPTER 4.

PROPOSED COMPETITIVE APPROACHES

In this chapter, we examine possible airline revenue management methodologies to take into account the competitive environment. In the first part we present an approach for heuristic adjustment of forecasts which incorporates information on the current seat inventory situation of competitors. In the second part we offer an overview of the impact of overrides of revenue management systems as competitive moves.

4.1 LOCO-based Forecast Multiplication

LOCO is an abbreviation designating the “lowest open class of the competitor”. LOCO-based Forecast Multiplication is a heuristic adjustment of the unconstrained forecasts of demand by one airline based on the simultaneous fare class availability offered by its competitors. We first present the theoretical principles of this method. We then use PODS simulations to estimate how much it can improve airline revenue management performance.

4.1.1 Concept

As its name suggests, LOCO-based Forecast Multiplication involves multiplication of the forecasts of unconstrained demand distributions. It takes into account the competitive environment, and more specifically the simultaneous seat inventory availability of the competitor. It compares the current competitive situation with the past competitive situations so as to adjust the current unconstrained demand distributions forecasted by traditional airline forecasters.

In part 3.2.1 we indicated two challenges facing the development of automated competitive approaches for revenue management: The unpredictability of competitor moves and the limited availability of competitive inventory data. LOCO-based Forecast Multiplication addresses these challenges by supposing that the competitor is letting its automated revenue management system run without overriding it, and by using the available competitive indicator that is LOCO (Lowest Open Class of the COmpetitor). LOCO is not an as precise indicator of availability as the actual competitors’ booking limits and bookings-in-hand, but it is more easily available. For a given LOCO, there are many possible combinations of the competitor’s current booking limits and bookings-in-hand.

The general idea behind LOCO-based Forecast Multiplication is the following: Suppose that the stochastic distribution of non-constrained demands and the ratios between the average fares of classes are steady. Suppose the competitor is nesting its protection levels. If without overriding its automated revenue management system, the competitor is more open (i.e. the LOCO is a lower fare class than usual) for a given time frame before departure, then this competitor is likely to compete more than usual for passengers for the given flight. It will spill-out less than historically in the future time frames before the same departure. Our forecasts of unconstrained demand should be decreased. Conversely, if the competitor is more closed than usual (i.e. the LOCO is a higher class than usual) for a given time frame

before departure, the competitor is likely to compete less than usual for passengers for the given flight. It will spill-out more than historically in the future time frames before the departure. Our forecasts of unconstrained demand should be increased.

Here is a more detailed reasoning. Suppose the stochastic distribution of non-constrained demands and the fare ratios between the average fares of classes are steady. If the competitor is not overriding its automated revenue management system, and if forecasters are relatively accurate, its protection levels should themselves remain mostly steady. If the competitor is more open for a fare class at the current time frame before departure than it has historically been for the same time frame, then there is higher likelihood that its bookings-in-hand for this class and the lower ones are currently smaller than on the comparable historical average. For any set of fare classes containing this class and the lower ones, the probability is small that the competitor will experience a non-constrained demand sufficiently high to make up for its currently low bookings-in-hand. Besides, the probability is small that the current low bookings-in-hand for this class and the lower ones are due to a low level of non-constrained demand for one class only. Therefore the probability is greater that the competitor's bookings-in-hands for any bottom-up set of classes will remain lower than usual in the future time frames for that same departure. The competitor will have higher chances to be more available than usual for all bottom-up sets of classes in the future time frames for the given departure. This should decrease its spill-out, and decrease the non-constrained demands for all classes of the other airlines. Our forecasts of unconstrained demand should be decreased. The opposite is true if the competitor is more closed than usual for a set of fare classes for a given time frame before departure.

For best performance of LOCO-based Forecast Multiplication, the definition of LOCO has to be restricted to the lowest open competitor class on the same path and for the directly competing flights in terms of schedule. LOCO has to be recorded for each time frame, and the current LOCO has to be compared to the average historical LOCO for the same time frames of preceding departure dates.

Formulations of LOCO-based Forecast Multiplication should be specific to the forecaster used by the airline to avoid double-counting. We have seen that a competitor more open than usual will decrease our level of non-constrained demand. There are two imaginable reasons why a competitor could be more open than usual with no override of its revenue management system: Either its own current non-constrained demand is lower than usual, or the total market level of non-constrained demand is lower than usual. Some forecasters detect the latter case, some do not. If they do, then there would be some double counting by decreasing our forecasts if the competitor is more open than usual. In this thesis we present a LOCO-based forecaster multiplier suitable for **additive pick-up forecasting**. This forecasting method considers that the non-constrained demands to come for each future time frame for a given departure are independent. To forecast future demands-to-come for a given departure, it merely adds up the historical distributions of demand for each of the remaining time-frames before departure. According to Talluri and van Ryzin [33] (p.470-472) additive pick-up forecasting is a simple heuristic used by many airlines and reported to perform well.

We have explained why the unconstrained forecasts for all classes and future time frames of a departing flight should change when a competitor is more open/closed than usual for the current time frame before departure. It is correct to say that all the classes are not equally affected by the competitive situation. Yet to keep the heuristic simple, we multiply unconstrained forecasts for all the classes by the same value.

Different formulations for LOCO-based Forecast Multiplication for pick-up forecasters have been developed as part of the PODS Consortium research at MIT:

- The forecast multiplier could be a function of the fraction of the historical observations of LOCO that were a higher/lower class than the current LOCO for the current time frame before departure:

$$(4.1) \quad FM_{LOCO} = 1 + \alpha(f_{lower} - f_{higher})$$

- Where:
- FM_{LOCO} is the value of the LOCO-based forecast multiplier
 - f_{lower} is the fraction of historical observations where the LOCO was a strictly lower class than the current LOCO.
 - f_{higher} is the fraction of historical observations where the LOCO was strictly higher class than the current LOCO.
 - α is a scaling constant

In that case the multiplier is equal to 1 when the current LOCO is equal to the median value of historically observed LOCOs. It is higher than 1 if the current LOCO is higher than the median, and lower than 1 if the current LOCO is lower than the median value of LOCOs. α is a scaling constant that controls the amplitude of the forecast multiplication. Note that the value of FM_{LOCO} belongs in that case to the interval $[1 - \alpha; 1 + \alpha]$. An implementation has to make sure the values of FM_{LOCO} do not become excessive, and for instance restrict them within the interval $[0.1; 10]$.

Example:

Given the frequencies that each class has historically been the LOCO for the current time frame:

LOCO	Frequency in the past observations used for forecasting
Class 1	0.1
Class 2	0.2
Class 3	0.3
Class 4	0.4

We can compute the different values of FM1 depending on the current LOCO observed (here we set $\alpha=1$):

Current LOCO	f_{lo}	f_{hi}	FM_{LOCO}
Class 1	0.9	0	1.9
Class 2	0.7	0.1	1.6
Class 3	0.4	0.3	1.1
Class 4	0	0.6	0.4

- The forecast multiplier could be a comparison of the current value of LOCO with the average value of LOCO over past observations of the current time frame before departure:

$$(4.2) \quad FM_{LOCO} = 1 + \alpha(\overline{LOCO} - LOCO)$$

- Where:
- FM_{LOCO} is the value of the LOCO-based forecast multiplier
 - \overline{LOCO} is the average index of LOCO over past observations
 - $LOCO$ is the current index of LOCO
 - α is a scaling constant

Classes are given an index, so that the lower the class in the nesting hierarchy, the higher the index. In that case the multiplier is equal to 1 when the current LOCO is equal to the average value of the historically observed LOCOs. α is a scaling constant that controls the amplitude of the forecast multiplication. Note that for the same value of α , the amplitude of the multiplier given by (4.2) is larger than the amplitude of the multiplier given by (4.1). An implementation would have to make sure the values of the forecast multiplier cannot leave a set interval, such as [0.1; 10].

Example:

LOCO	index
Class 1	1
Class 2	2
Class 3	3
Class 4	4

We can compute the different values of FM1 depending on the current LOCO observed (here we set $\alpha=0.5$, and $\overline{LOCO} = 2.9$):

Current LOCO	\overline{LOCO}	FM_{LOCO}
Class 1	2.9	1.95
Class 2	2.9	1.45
Class 3	2.9	0.95
Class 4	2.9	0.45

- The forecast multiplier could be a function of the relative fares of the current LOCO and the historical LOCO over the past similar time frames before departure.

$$(4.3) \quad FM_{LOCO} = 1 + \alpha \left(\frac{\text{fare}_{LOCO} - \overline{\text{fare}_{LOCO}}}{\overline{\text{fare}_{LOCO}}} \right)$$

- Where:
- FM_{LOCO} is the value of the LOCO-based forecast multiplier
 - $\overline{\text{fare}_{LOCO}}$ is the average value of the fare of the LOCO over past observations
 - fare_{LOCO} is the current value of the fare of the LOCO
 - α is a scaling constant

This model is similar to the preceding model, but here the competitor's classes are weighted by their average fares. Note that again, an implementation would have to make sure the values of the forecast multiplier cannot leave a set interval, such as [0.1; 10].

Example:

LOCO	fare
Class 1	1000
Class 2	800
Class 3	400
Class 4	300

We can compute the different values of FM1 depending on the current LOCO observed (here we set $\alpha=1$, and $\overline{\text{fare}_{LOCO}} = 500$):

Current LOCO	$\overline{\text{fare}_{LOCO}}$	FM_{LOCO}
Class 1	500	2.0
Class 2	500	1.6
Class 3	500	0.8
Class 4	500	0.6

For LOCO-based Forecast Multiplication to be unbiased, the average value of the LOCO-based forecast multiplier over the set of the possible LOCOs must be equal to 1. Mathematically this can be written as:

$$(4.4) \quad \sum_{i=1..n} f_i . FM_{LOCO}(i) = 1$$

- Where: - There are n classes.
- $FM_{LOCO}(i)$ is the value of the LOCO-based forecast multiplier if class i is the current LOCO.
 - f_i is the probability that class i be the current LOCO, also equal in a steady state to the past frequency at which class i has historically been the LOCO for the current time frame.

If it did not verify this property, the method would be biased and would either be over-multiplying or under-multiplying forecasts.

The multipliers suggested in (4.1), (4.2), (4.3) verify this property.

Proof for (4.1):

$$\begin{aligned} \sum_{i=1..n} f_i . FM_{LOCO}(i) &= \sum_{i=1..n} f_i . (1 + \alpha \sum_{j=i+1..n} f_j - \alpha \sum_{k=1..i-1} f_k) \\ &= 1 + \alpha \sum_{i=1..n} (\sum_{j=i+1..n} f_i f_j - \sum_{k=1..i-1} f_i f_k) \\ &= 1 + \alpha (\sum_{i=1..n} \sum_{j=i+1..n} f_i f_j - \sum_{i=1..n} \sum_{k=1..i-1} f_i f_k) \\ &= 1 + \alpha (\sum_{i=1..n} \sum_{j=i+1..n} f_i f_j - \sum_{k=1..n} \sum_{i=k+1..n} f_i f_k) \\ &= 1 \end{aligned}$$

Proof for (4.2):

Let x_i be the index given to class i

$$\begin{aligned} \sum_{i=1..n} f_i . FM_{LOCO}(i) &= \sum_{i=1..n} f_i . (1 + \alpha \sum_{j=1..n} f_j x_j - x_i) \\ &= 1 + \alpha (\sum_{i=1..n} \sum_{j=1..n} f_i f_j x_j - \sum_{i=1..n} f_i x_i) \\ &= 1 + \alpha (\sum_{i=1..n} \sum_{j=1..n} f_i f_j x_j - (\sum_{j=1..n} f_j) (\sum_{i=1..n} f_i x_i)) \\ &= 1 + \alpha (\sum_{i=1..n} \sum_{j=1..n} f_i f_j x_j - \sum_{i=1..n} \sum_{j=1..n} f_j f_i x_i) \\ &= 1 \end{aligned}$$

Proof for (4.3):

Identical to the proof for (4.2), where x_i is replaced by the ratio between the fare of the current LOCO and the historical average fare of the LOCO.

End of proofs

There are different ways to control the magnitude of LOCO-based Forecast Multiplication. As additive pick-up forecasting differentiates the forecasts by time frame, one could decide to apply forecast multiplication only to the most adjacent time frames that are to come in the booking period. Theoretically the closer the time frame to the current time frame, the more justified the assumptions of Forecast Multiplication, and the more robust Forecast Multiplication should be. A second way to control the magnitude of Forecast Multiplication lies in the choice of the scaling constant α .

The practice of LOCO-based Forecast Multiplication nevertheless entails two limits:

- It supposes that the competitor is not overriding its automated revenue management system. Such overrides include the practice of LOCO-based Forecast Multiplication itself. The theory does not hold if both competing airlines use LOCO-based Forecast Multiplication. In this thesis we will simulate such cases of two airlines using LOCO-based Forecast Multiplication to see what the implications could be for airlines.
- The concept of Forecast Multiplication is to better adapt to a specific external competitive environment. It is best justified if it simply fits better protection levels with non-constrained demand without impacting the other airline. Yet in reality by automatically multiplying forecasts and modifying its protection levels, as explained in 3.2, an airline will have an impact on its competitor, and experience a feedback effect from its own forecast multiplication. Suppose the competitor is more open than usual at a given time frame for a departing flight. LOCO-based Forecast Multiplication will decrease forecasts, protections will be lower, thus spilling-out less for low classes, increasing the chances that the competitor will be more open than usual in the future. The opposite is true if the competitor is more closed than usual. Forecast multiplication reinforces the cause that sets it off. This does not question the validity of the method. Yet it requires that simulations be done to check that this competitive feedback of LOCO-based Forecast Multiplication within the entire booking period preceding the departure of a flight does not lead to an overall decrease in expected revenues.

However, across successive flight departures, the impact of LOCO-based Forecast Multiplication does not progressively trigger the method more strongly. If the distributions of non-constrained demands are steady, the more successive flights have been “more open than usual” over their booking period, the smaller the chances that the following flights will themselves be “more open than usual”. The more successive flights have been “more closed than usual” over their booking period, the smaller the chances that the following flights will themselves be “more closed than usual”.

4.1.2 Experimental Methodology

We tested the three presented formulations of LOCO-based Forecast Multiplication in various simulation environments. In this thesis we only show the simulation results for the first formulation of equation (4.1), and we use the LOCO of the flight directly competing in terms of schedule, rather than the LOCO of the other competitor flights departing on the same day. It proved to be the most consistently performing formulation. We will call this formulation **FM1**, and we recall here its definition:

$$(4.1) \quad FMI = 1 + \alpha(f_{lower} - f_{higher})$$

- Where:
- FM_i is the value of the LOCO-based forecast multiplier
 - f_{lower} is the fraction of historical observations where the LOCO was a strictly lower class than the current LOCO.
 - f_{higher} is the fraction of historical observations where the LOCO was a strictly higher class than the current LOCO.
 - α is a scaling constant
 - values of FMI are kept within [0.1; 10].

We use the PODS simulator to model complex real world airline revenue management interactions to evaluate the impact of LOCO-based Forecast Multiplication. For all simulations, we compare a case where forecast multiplication is used with a **base case** where no forecast multiplication is used by any competitor. The most pertinent data are the proportional changes of airline statistics from the base case due to the use of forecast multiplication. We more particularly report the impact on both competitors, the performance of LOCO as a competitive indicator, the robustness of the method, the spiral situations evocated in 4.1.1, the amplitude of the value of the forecast multipliers, and the contradictory situation of both competitors using LOCO-based Forecast Multiplication.

The simulations are run with parameters as defined in Chapter 1.5. We study a single market duopoly situation to keep the experiment simple and to better discriminate the impact of LOCO-based Forecast Multiplication. We use the same market parameters as presented in Chapter 3. Airline 1 and Airline 2 each offer three flights a day at the same times, in the same direction in the same origin-destination market, with aircrafts of a capacity 120 seats each. We run the following scenarios:

1. A single market representing the competition between a legacy carrier and a low-cost carrier. Airlines have the unrestricted fare structure provided in Fig 4.1.2 (which is the same as the unrestricted fare structure used in Chapter 3, Fig 3.1.3). Airline 1, the legacy carrier, uses EMSRb optimization, a combination of Q-forecasting and pick-up forecasting with booking curve unconstraining, and implements LOCO-based Forecast Multiplication against Airline 2, a low-cost competitor.
2. A single market representing the competition between two legacy carriers. Airlines have the fully restricted fare structure provided in Fig 4.1.3 (which is the same as the fully restricted fare structure used in Chapter 3, Fig 3.1.2). Both airlines use EMSRb optimization, pick-up forecasting with booking curve unconstraining. Only Airline 1 is implementing LOCO-based Forecast Multiplication.
3. The same case as 2, but with both legacy carriers implementing LOCO-based Forecast Multiplication.

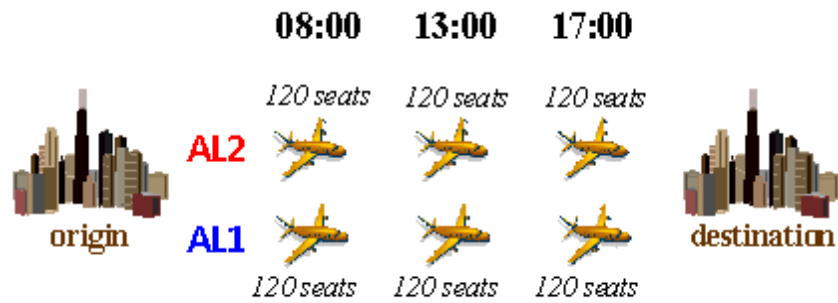


Fig 4.1.1 Single market duopoly scenario studied

Class	Fare	Advance purchase (days)	Saturday Night Stay	Change Fee	Non Refundable
1	500	0	No	No	No
2	400	0	No	No	No
3	315	0	No	No	No
4	175	0	No	No	No
5	145	0	No	No	No
6	125	0	No	No	No

Fig 4.1.2 Unrestricted fare structure considered

Fare Class	Fare	Advance purchase (days)	Saturday Night Stay	Change Fee	Non Refundable
1	500	0	No	No	No
2	400	3	No	No	Yes
3	315	7	No	Yes	Yes
4	175	10	Yes	Yes	Yes
5	145	14	Yes	Yes	Yes
6	125	21	Yes	Yes	Yes

Fig 4.1.3 Fully-restricted fare structure considered

4.1.3 Simulation results

Scenario 1: Unrestricted fare structure, legacy carrier using LOCO-based Forecast multiplication against a low-cost competitor

In this simulation we wish to represent a single market with competition between a legacy carrier and a low-cost carrier. Airline 1, the legacy carrier, uses EMSRb optimization, a combination of Q-forecasting with Frat5c, and pick-up forecasting with booking curve unconstraining. Frat5c is an estimate of the fare ratio curve at which 50% of the passengers are expected to sell-up for each time-frame. A precise description of Q-forecasting and Frat5c are provided in Cleaz Savoyen[7]. The legacy carrier Airline 1 implements LOCO-based Forecast Multiplication against Airline 2, a low-cost competitor.

To model the revenue management behaviour of a low-cost airline in PODS, an optimization method called “Accordion Threshold” is used. It corresponds to an adaptive dynamic method that optimizes booking limits in order to achieve a high target load factor (hence the image of an “accordion” method). The method sets thresholds which are the booked load factors at which each class is closed. “Accordion Threshold” multiplies fixed initial thresholds by an accordion parameter ap . ap is computed as a function of current bookings, past load factors, and the past values of ap itself. “Accordion Threshold” has five inputs:

- A load factor goal $lfgoal$
- Initial threshold levels
- A minimum accordion parameter allowed $apmin$
- A maximum accordion parameter allowed $apmax$
- A convergence smoothing constant α

ap is computed according to the following formula:

$$(4.5) \quad ap = \max[apmin, \min[apmax, \overline{ap}(1 + \alpha(\frac{lf}{\overline{lf}} - 1))]]$$

Where:

- \overline{ap} = average historical accordion parameter
- \overline{lf} = average historical load factor, leg l
- lf is an adjusted target load factor taking into account current bookings:

$$(4.6) \quad lf = \frac{C.lfgoal - B}{C - B} \quad \text{where } C \text{ is the leg capacity and } B \text{ the current bookings}$$

Note that this optimization method does not require any forecast of non-constrained demand.

In this simulation airline 2 sets this load factor goal $lfgoal$ to 90%, and initial thresholds are given in the following table. $apmin$ is set to 0.5, $apmax$ to 1.5 and the smoothing constant α is set to 0.75.

Fare classes	Initial nested load factor thresholds
1,2,3,4,5,6	1.00
2,3,4,5,6	0.90
3,4,5,6	0.80
4,5,6	0.65
5,6	0.50
6	0.35

Fig 4.1.4 Initial thresholds set by Airline 2

We provide the results when Airline 1 uses FM1 with a scaling constant equal to 1 and FM1 is applied to all the future time frames.

As shown in Fig 4.1.5 FM1 significantly increases the average revenues of Airline 1 (+8.3%) while very slightly decreasing the average revenues of Airline 2 (-0.2%). This simulation result is comforting, because LOCO-based Forecast Multiplication is originally designed to improve forecasts, to better adapt the protection levels to external non-constrained demand, but without any specific intent to impact the competitor. The fact that Airline 2 is very little impacted by the implementation of FM1 by Airline 1 could mean that the competitive feedback effect within one booking period for a single departure evoked in 4.1.1 is very small. The increase in revenues for Airline 1 is achieved through both an increase in average load factor and in average yield. FM1 generates higher average bookings for the highest and the lowest fare classes to the detriment of bookings for the “middle classes” 4 and 5. The reason for this is explained in part 4.1.1: When the competitor is more closed than usual, with FM1 Airline 1 increases its forecasts and protects more for the highest classes. When the competitor is more open than usual, with FM1 Airline 1 decreases its forecasts and increases the booking limits for the lowest classes. This overall pushes the bookings to “extreme” classes. Average daily bookings per fare class are shown in Fig 4.1.6.




Fig 4.1.5 Average proportional change in revenues, load factor, and yield for each airline when Airline 1 uses FM




Fig 4.1.6 Comparison of Airline 1 average bookings by fare class when it is using FM1 and when it is not

It is possible to check that across successive flight departures, the impact of LOCO-based Forecast Multiplication does not progressively trigger the method more strongly. Fig 4.1.7 represents statistics for 200 successive days of departure with the use of FM1, and no specific spiral situation is noticeable. We provide in Fig 4.1.8 the distribution of the proportional increase in revenues that FM1 would generate for Airline 1 over the same 200 successive departure dates. Even though the standard deviation of the distribution is larger than its mean value, FM1 increases revenues for 74% of the departure dates. This result is statistically significant, because if Forecast Multiplication was purely random and either increasing or decreasing revenues like a coin flip, the probability to have at least 74% of increases in revenues over 200 departure dates would be smaller than one out of a million.

(This probability is $\sum_{n=148..200} C_{200}^n \frac{1}{2^{200}} = 3.6 \times 10^{-12}$)




Fig 4.1.7 Statistics for Airline 1 using FM1 over 200 successive days of departures.




Fig 4.1.8 Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 over 200 successive departure dates

We also provide a distribution of the impact on Airline 2 of the use of FM1 by Airline 1 in Fig 4.1.9 to check how variable the impact of FM1 is on the competitor. The impact is mostly grouped around the mean value of -1% over the 200 successive departures studied. Therefore the feedback competitive effect within a single departure date can be considered to be small.




Fig 4.1.9 Distribution over 200 successive departure dates of the percentage change in revenues for Airline 2 when Airline 1 uses FM1

As shown in Fig. 4.1.10 when FM1 decreases Airline 1 revenues, it corresponds to departure dates with a low level of total market demand (There is a direct correlation between base case revenues and total market demand). FM1 increases revenues in all scenarios of total market demand. It increases more than one half of the departure dates where total market demand was low to medium, and it increases the revenues of most of the high demand scenarios. When total market demand is high, Q-forecasting is not aggressive enough and does not set protection levels for the high classes high enough. FM1 improves Q-forecasting by detecting that the competitor is then more closed than usual and by protecting more for high classes. Therefore when comparing in Fig 4.1.11 the distribution of Airline 1 revenues when it uses FM1 and when it does not, the effect of using FM1 is to push a significant part of the medium to high revenues to larger values. For low values, the distribution of Airline 1 revenues is nearly unchanged by the use of FM1.




Fig 4.1.10 Distribution over 200 successive departure dates of Airline 1 base case (not using FM1) revenues when FM1 increases revenues and when it decreases revenues




Fig 4.1.11 Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1

Fig 4.1.12 provides a representative sample of the evolution of the value of FM1 over a booking period for different flights. The first feature to notice is that the values of FM1 for a single departure tend to remain either superior or inferior to 1. This could indicate that the feedback competitive effect within one departure date evoked in 4.1.1 could take place. Yet it could just be that total market demand is very variable. Besides the fact that the impact of FM1 on Airline 2 is on average much smaller than the impact of FM1 on airline 1 leads us to think that the spiral situation within a single booking period for a departure date is small. The second feature to notice is that LOCO is an imperfect discontinuous indicator of the competitor's availability, occasionally causing "jumps" of the value of FM1 from one time frame to another. As an indicator, LOCO does not allow to compare situations where the current LOCO is about to be closed with situations where the LOCO is still allocated plenty of available seats. If the competitor is lacking demand and just closed a class, LOCO-based Forecast Multiplication can consider that the competitor caught up with its lack of demand, when it is in fact falling short of demand.





Fig 4.1.12 Examples of values of FM1 over sample flights at different departure dates.


We provide in Fig 4.1.13 the evolution of the distribution over 2000 departure dates of the values of FM1 for successive time frames of a booking period. The time frames are defined in Fig 1.2.5 (Chapter 1). The distributions are widespread except in the first time frames (early days of bookings for a flight departure) and in the last time frames (last days before departure). These peaks of distribution of the value of FM1 in the first and last time frames are due to the fact that the set of possible past LOCOs is then reduced, making of LOCO an imprecise indicator of the availability of the competitor. In the first time frames, whatever the speed at which the competitor is observing bookings, there are very high chances that its LOCO will be its lowest class. Similarly, in the last time frames, there are strong chances that the low-cost competitor will be sold out, making of the LOCO a poor indicator of how much the demand of Airline 1 should be impacted if its competitor Airline 2 has sold out its entire flight capacity.

Last but not least, the nature of all the results presented above remains unchanged when the scaling constant α is kept reasonable (between 0 and 2), or when Forecast Multiplication is applied to a limited number of time frames to come. A different scaling constant or number of time frames to which LOCO-based Forecast Multiplication is applied only changes the amplitude of the mechanisms that we described here.


Time Frame 3




Time Frame 5



Time Frame 7



Time Frame 9



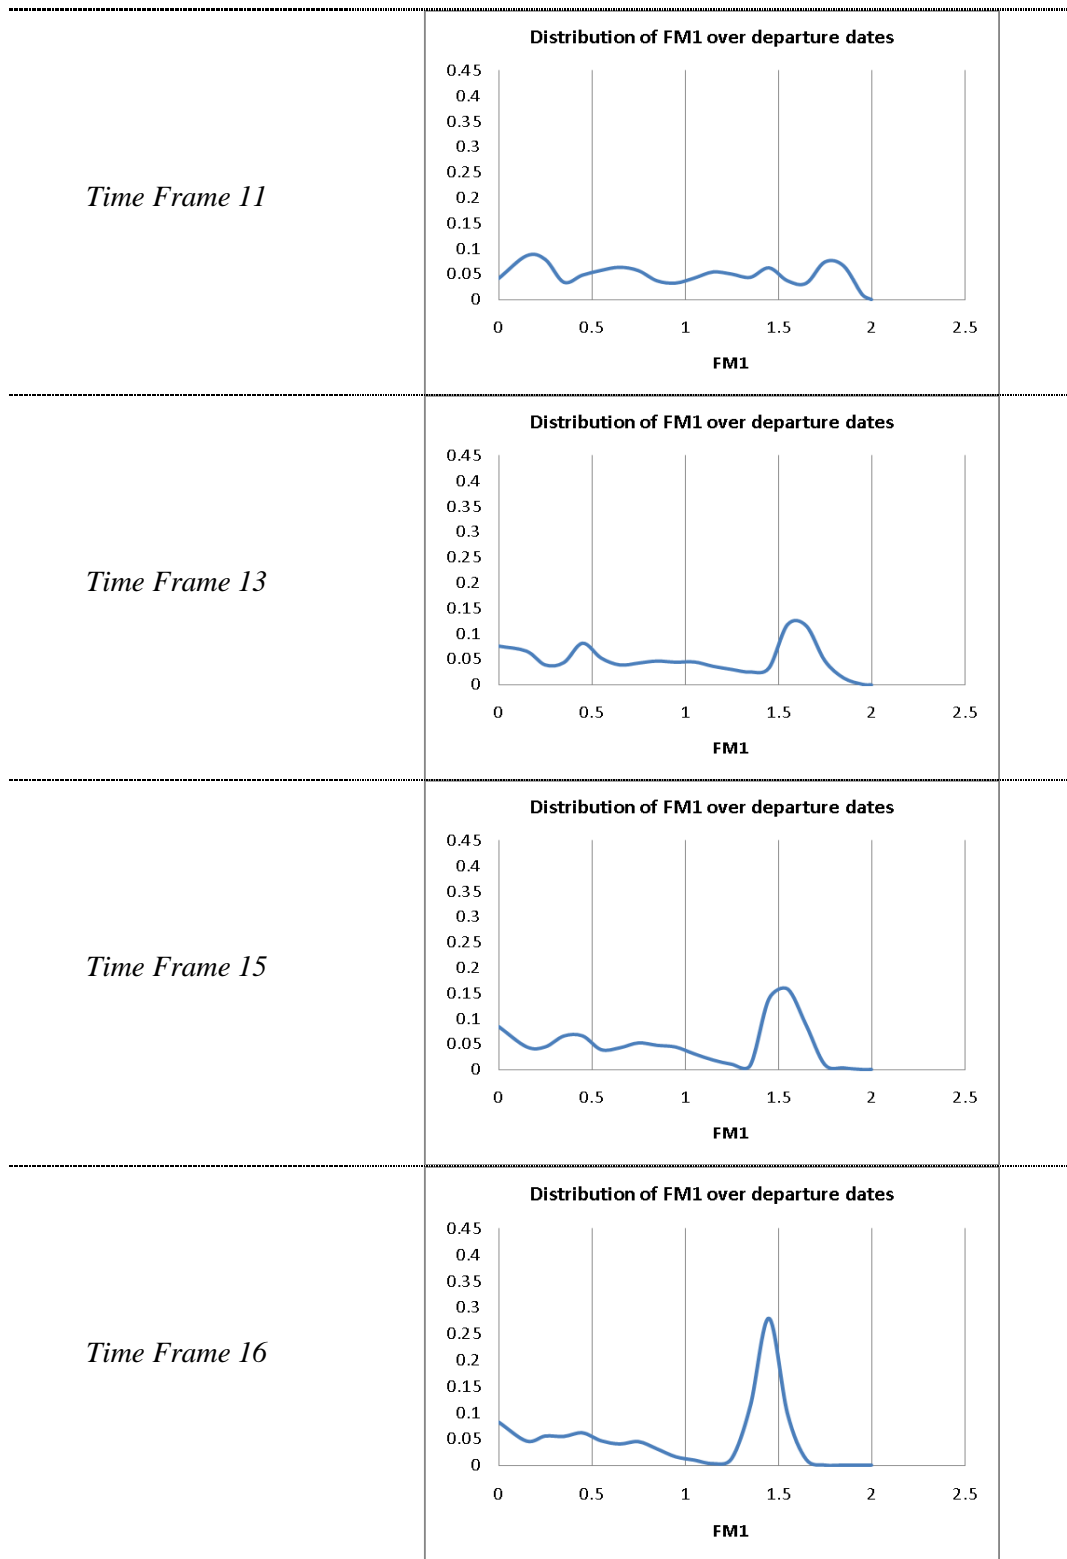


Fig 4.1.13 Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates

Scenario 2: Fully-restricted fare structure, two identical legacy carriers using EMSRb, one of them using LOCO-based Forecast Multiplication.

We saw the positive impact of FM1 if used by a legacy carrier in a single market with an unrestricted environment where it is facing the competition of a low-cost competitor. We now want to test the impact of FM1 in a single market if it is applied by a legacy carrier facing the competition of an identical legacy carrier, with a fully-restricted fare structure environment. In this simulation, both competitors Airline 1 and Airline 2 are identical in schedule and capacity, in passenger preference, and they use EMSRb optimization, additive pick-up forecasting with booking curve unconstraining. Only Airline 1 applies FM1. We provide the results when Airline 1 uses FM1 with a scaling constant equal to 1 and FM1 is applied to all the future time frames.

As shown in Fig 4.1.14 FM1 increases the average revenues of Airline 1 by +1.9% while also increasing the average revenues of Airline 2 by the smaller value of 0.6%. Although the increase in revenues for Airline 1 is smaller than in the unrestricted case, it is still important and is similarly achieved through both an increase in average load factor and in average yield. FM1 pushes bookings to extreme classes through the same mechanism: When the competitor is more closed than usual, FM1 makes Airline 1 close more, and when the competitor is more open than usual, FM1 makes Airline 1 open more. Fig 4.1.15 displays the difference in average bookings per class and per departure date for Airline 1 when it does not use FM1 and when it does.




Fig 4.1.14 Average proportional change in revenues, load factor, and yield for each airline when Airline 1 uses FM




Fig 4.1.15 Comparison of Airline 1 average bookings by fare class for its 3 daily flights when it is using FM1 and when it is not

As shown in Fig 4.1.16, the simulation confirms that across successive flight departures, the impact of LOCO-based Forecast Multiplication does not progressively trigger the method more strongly.




Fig 4.1.16 Statistics for Airline 1 using FM1 over 200 successive days of departures

We provide in Fig 4.1.17 the distribution of the proportional increase in revenues that FM1 would generate for Airline 1 over the same 200 successive departure dates. The standard deviation of the distribution is the same as in scenario 1, but the mean increase in revenues is smaller. FM1 is not as effective in a fully-restricted environment against a competitor using EMSRb than in an unrestricted environment against a low-cost competitor. Revenues are increased in 56% of the cases. However this result remains statistically significant, because if forecast multiplication was purely random and either increasing or decreasing revenues like a coin flip, the probability to have at least 56% of increases in revenues over 200 departure dates would be 5.2%. We are thus 95% confident of the statistical significance of the result.

(This probability is $\sum_{n=112..200} C_{200}^n \frac{1}{2^{200}} = 0.052$)




Fig 4.1.17 Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 over 200 successive departure dates

We also check how variable the impact of FM1 is on the competitor. The distribution of the impact on Airline 2 revenues is provided in Fig 4.1.18. The impact is mostly grouped around the mean value of +0.1% over the 200 successive departures studied. The feedback competitive effect within a single departure date can be considered to be relatively small.




Fig 4.1.18 Distribution over 200 successive departure dates of the percentage change in revenues for Airline 2 when Airline 1 uses FM1

Fig. 4.1.19 shows that FM1 would lead to an increase in Airline 1 revenues of slightly more than one half of the departures, when it would decrease Airline 1 revenues in slightly less than one half of the departures, whatever the total market demand for the departure. The departures for which FM1 is much more probable to increase Airline 1 revenues than to decrease them are scenarios of low demand and revenues. It is possible that when total market demand is low, both airlines are too aggressive, especially with pick-up forecasting. Fig 4.1.20 compares the distribution of Airline 1 revenues when it uses FM1 and when it does not. The use of FM1 by Airline 1 spreads this distribution.




Fig 4.1.19 Distribution over 200 successive departure dates of Airline 1 base case (not using FM1) revenues when FM1 increases revenues and when it decreases revenues





Fig 4.1.20 Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1


We also find that one of the main reasons why in a restricted environment with a competitor using EMSRb LOCO-based Forecast Multiplication does not perform as well is that the discontinuous imperfections of LOCO as an indicator of the competitor's availability are increased if the competitor is using a fully-restricted structure. We provide in Fig 4.1.21 the evolution of the distribution over 2000 departure dates of the values of FM1 for successive time frames of a booking period. For the same reasons why the distributions present a peak in the first and last time frames, the set of possible past LOCOs is reduced in time frames 9, 11, 12, 13, 15 because then the advance purchase rules set by the competitor successively force its lower classes to close. Whatever the speed at which the competitor is observing booking, when an advance purchase rule just forced a class to close, there are very high chances that its current LOCO be its lowest possible open class considering advance purchase requirements, so the distribution of the values of FM1 will present a peak at a value lower than 1. As a consequence, advance purchase rules reduce the precision of LOCO as an indicator of the competitor's relative availability. In Fig 4.1.21, the distribution of the values of FM1 is widespread from time frame 5 to 8 only. From time frame 1 to 4, there is a peak of distribution at a value lower than 1 which corresponds to cases where the competitor has still not closed its lowest class. Starting time frame 9, advance purchase restrictions of the competitor's classes repeatedly create peaks of distribution at a value lower than 1. This is when the lowest possible open class considering the competitor's advance purchase requirement is open. In the last two time frames appears a peak of distribution at a value superior to 1 which corresponds to a competitor with its entire capacity booked. The peak at a value lower than one then corresponds to its class 1 open.

Finally, the nature of all the results presented above remains unchanged when the scaling constant α is kept reasonable (between 0 and 2), or when forecast multiplication is applied to a limited number of time frames to come. A different scaling constant or the number of time frames to which LOCO-based Forecast Multiplication is applied only changes the amplitude of the mechanisms and results that we described here.


Time Frame 3




Time Frame 5



Time Frame 7



*Time Frame 9:
Class 6 closed by advance
purchase requirement*



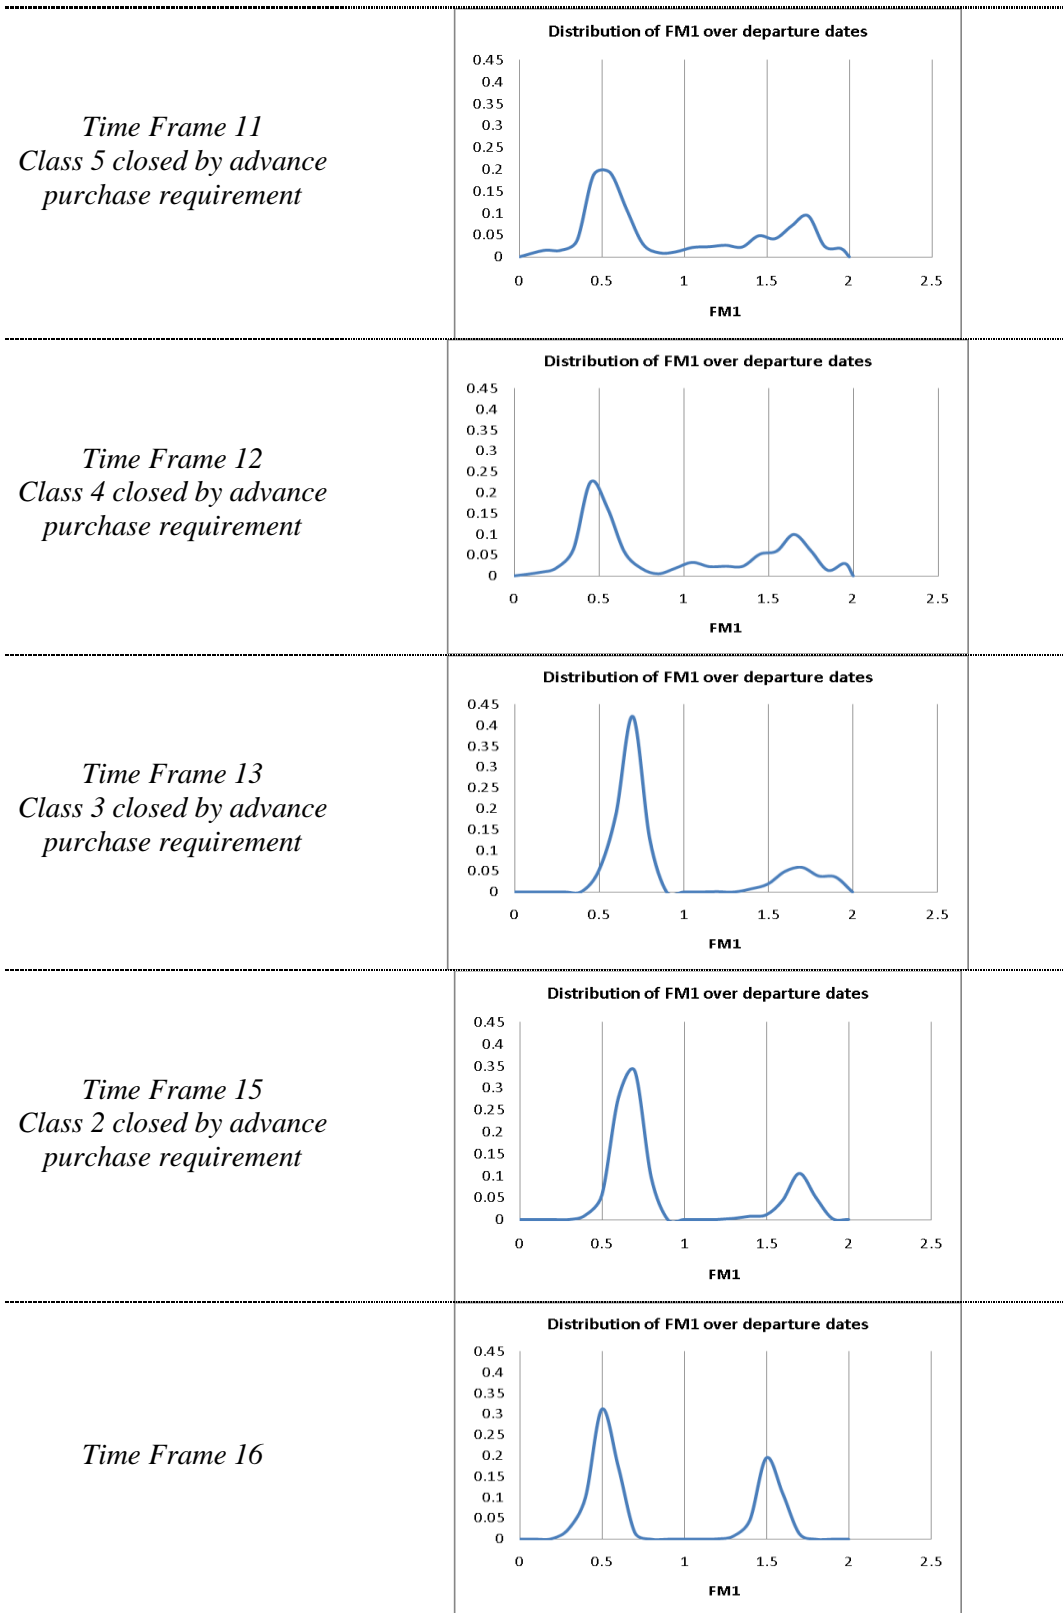


Fig 4.1.21 Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates

Scenario 3: Fully-restricted fare structure, two identical legacy carriers using EMSRb, both of them using LOCO-based Forecast Multiplication.

LOCO-based Forecast Multiplication assumes that the competitor is not overriding its automated revenue management system. Overrides include the practice of LOCO-based Forecast Multiplication itself. Our objective with this scenario is to test the possible situation of both competitors simultaneously using such Forecast Multiplication. To do so we take Scenario 2 and have Airline 2 also use FM1 with a scaling constant equal to 1 and FM1 applied to all the future time frames.

We find as shown in Fig 4.1.22 that its use by both airlines still increases their individual average revenues and yields. Total market revenues and total market yield are on average increased by 1.2%. By using FM1, both competitors push average bookings to “extreme classes” as displayed in Fig 4.1.23 and Fig 4.1.24. Overall Airline 2 is slightly advantaged by the use of FM1, but this constitutes an effect of the stochasticity modeled in the PODS simulator.

We provide in Fig 4.1.25 the distributions of the proportional increase in revenues that FM1 would generate for Airline 1 and Airline 2 over 200 successive departure dates. The mean increase in revenues for each airline is relatively small (+0.5% for Airline 2, +0.3% for Airline 1) especially compared with the standard deviations. Revenues are increased respectively in 52% and 53% of the cases for Airline 1 and Airline 2, so that it is difficult to assert that the increase in revenues due to FM1 over these 200 departure dates is statistically significant. When both airlines use FM1, its positive impact on an individual airline is more limited and questionable. However over the overall 2000 departure dates simulated, the increase in revenues for each airline is still larger and statistically significant. As regards the distribution of the proportional increase in total market revenues given in Fig 4.1.25, it increases revenues in 55% of the cases, so it is slightly more statistically significant than the increase in individual airlines’ revenues, but it is still not as statistically significant as the increase in revenues of Airline 1 when it is the only one to implement FM1.

When looking at Fig. 4.1.27 and Fig 4.1.28, it appears that FM1 implemented at both airlines decreases the revenues of Airline 1 in nearly half of the cases whatever the total market demand, when it decreases Airline 2 revenues in more than one half of the cases of low total market demand and increases revenues in more than one half of the scenarios of medium total market demand. As regards total market revenues, they are increased by the simultaneous use of FM1 by both airlines in more than one half of the scenarios of medium total market demand, and decreased in more than one half of the scenarios of low or high total market demand, as shown in Fig 4.1.29.

Surprisingly, if used by both airlines, FM1 does not impact much the overall distribution of the revenues of individual airlines and hence of the total market revenues. But when looking at the distributions of the values of FM1 given in Fig 4.1.33, when both airlines use FM1, one can observe a spiral competitive situation of FM1, with fare class 6 being open longer at the beginning of a booking period than they are when only one airline uses FM1. Indeed if both airlines use LOCO-based Forecast Multiplication, the historical distribution of LOCO should be pushed to extreme classes. When an airline is more open (respectively closed) than usual, it makes its competitor open more (resp. closed), increasing the chances that the initial airline also opens (resp. closes) more by using LOCO-based Forecast Multiplication.




Fig 4.1.22 Average proportional change in revenues, load factor, and yield when both airlines use FM




Fig 4.1.23 Comparison of Airline 1 average bookings by fare class when both airlines are using FM1 and when they are not




Fig 4.1.24 Comparison of Airline 2 average bookings by fare class when both airlines are using FM1 and when they are not




Fig 4.1.25 Distribution of the percentage increase in revenues that FM1 would generate for Airline 1 and Airline 2 over 200 successive departure dates




Fig 4.1.26 Distribution of the percentage increase in total market revenues that FM1 would generate over 200 successive departure dates




Fig 4.1.27 Distribution over 200 successive departure dates of Airline 1 base case (not using FM1) revenues when FM1 increases or decreases Airline 1 revenues




Fig 4.1.28 Distribution over 200 successive departure dates of Airline 2 base case (not using FM1) revenues when FM1 increases or decreases Airline 2 revenues




Fig 4.1.29 Distribution over 200 successive departure dates of base case (not using FM1) total market revenues when FM1 increases or decreases total market revenues




Fig 4.1.30 Distribution over 200 successive departure dates of Airline 1 revenues without and with the use of FM1 by both airlines

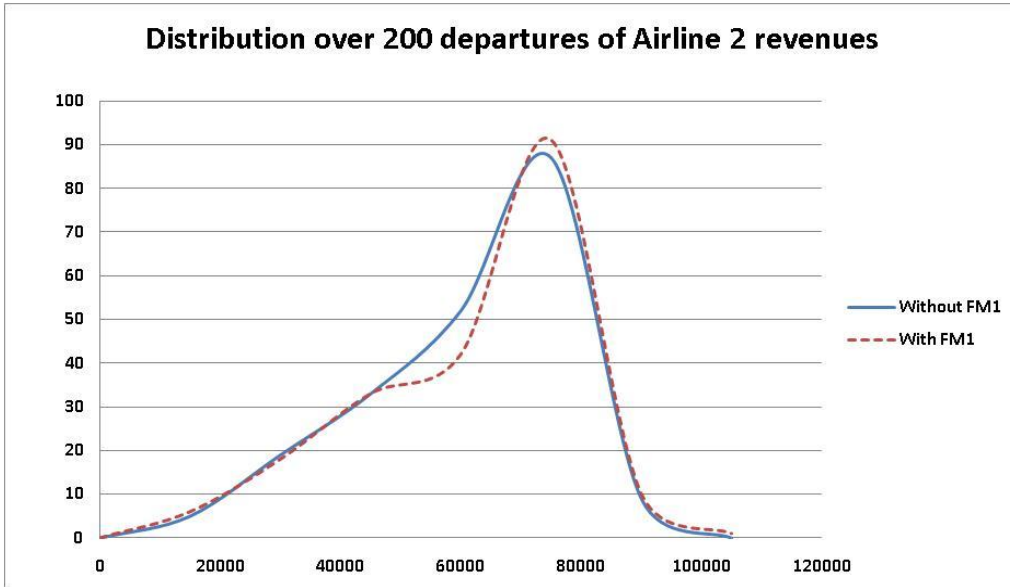


Fig 4.1.31 Distribution over 200 successive departure dates of Airline 2 revenues without and with the use of FM1 by both airlines


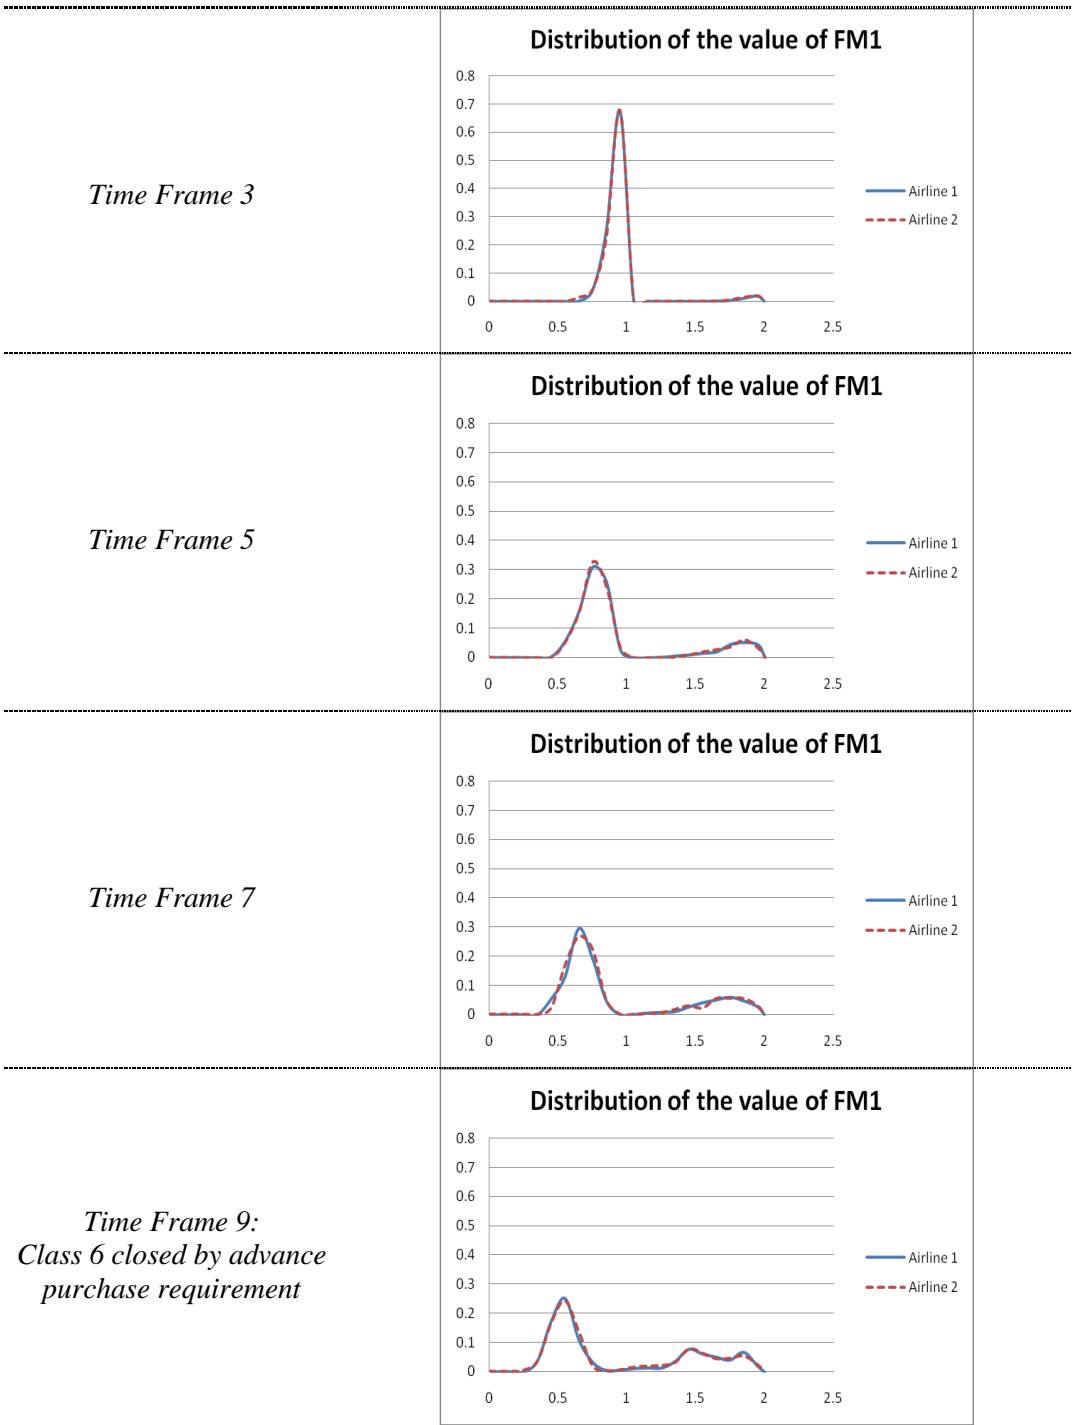


Fig 4.1.32 Distribution over 200 successive departure dates of total market revenues without and with the use of FM1 by both airlines



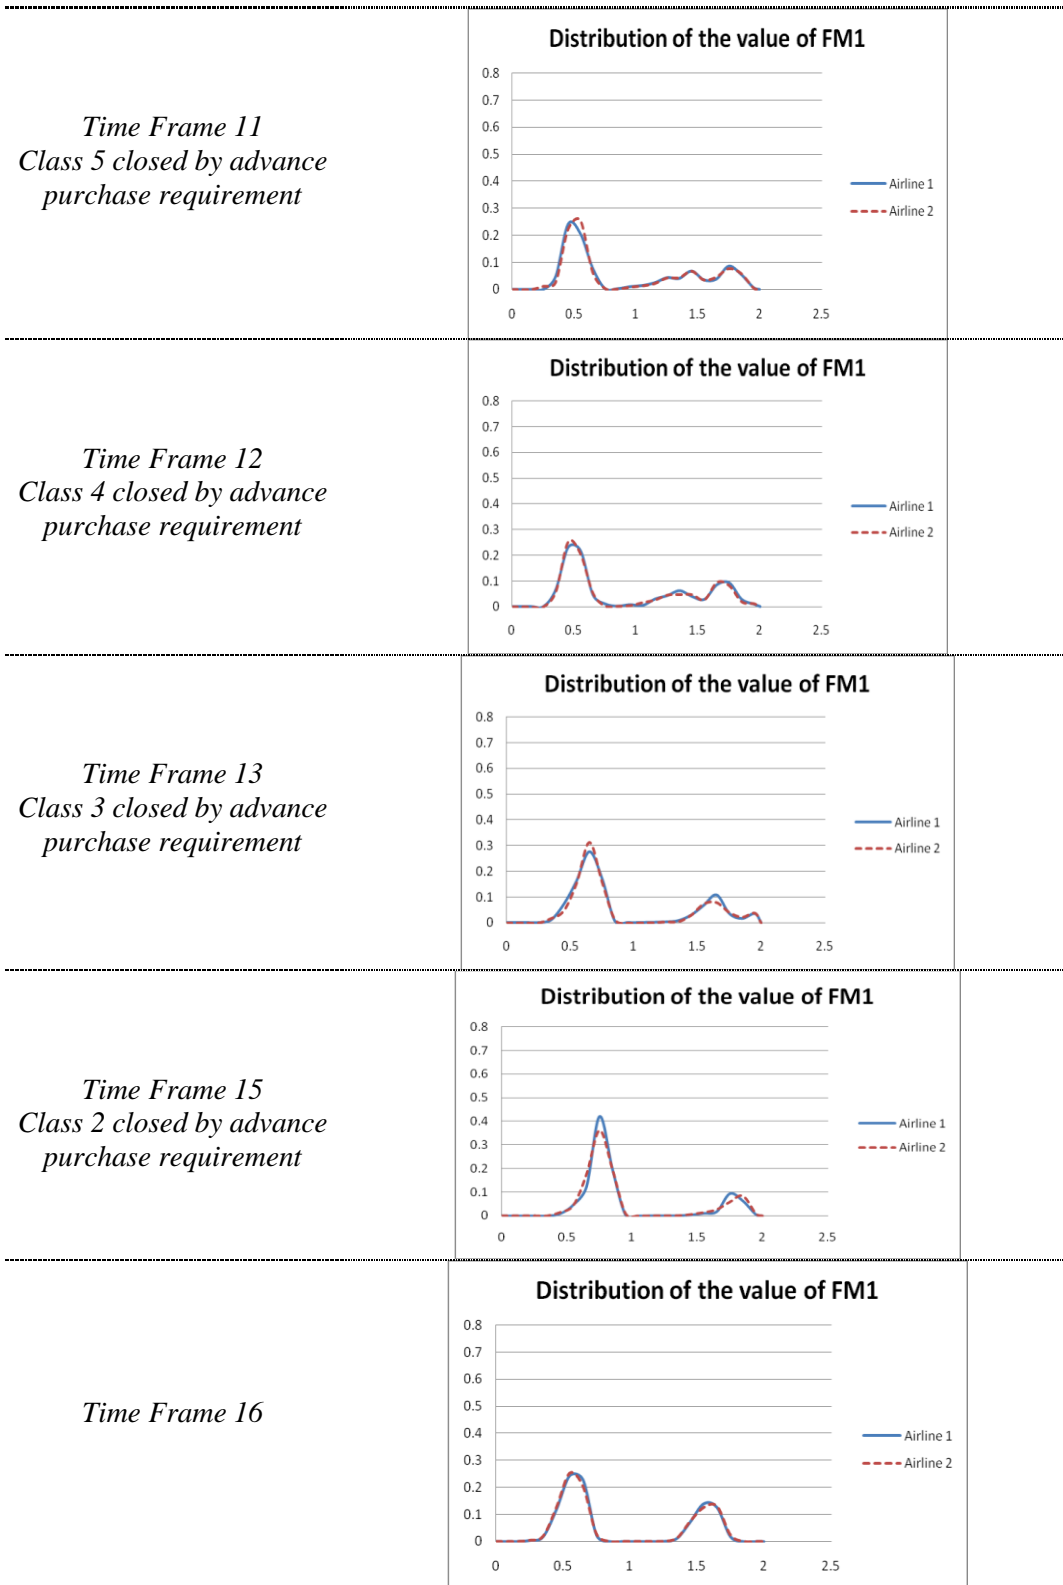


Fig 4.1.33 Distribution of the values of FM1 by time frame of the booking period over 2000 departure dates

4.1.4 Conclusions

We introduced LOCO-based Forecast Multiplication, a heuristic multiplication of additive pick-up forecasts developed by the members of the PODS Consortium. It is based on the comparison of the competitor's current fare class availability with its corresponding historical availability. Its main idea is that a competitor more available (respectively less available) than usual for a flight at a given point of the booking period is prone to remain so until the departure, resulting in more (respectively less) competition, so that forecasts of unconstrained demands should be decreased (respectively increased). The indicator of availability used is LOCO, the lowest open competitor class. The assumptions made are that the non-constrained demand distributions are steady over time, that the competitor does not override its forecast or its optimization method. Overrides include LOCO-based Forecast Multiplication itself. We presented PODS simulations results of the impact of FM1, one of the suggested formulations for LOCO-based Forecast Multiplication, which is a function of the fraction of the historical observations of LOCO that were higher/lower fare classes than the current LOCO.

When applied by one airline, we found that FM1 can increase revenues significantly for this airline, with a smaller impact on its competitor. LOCO-based Forecast Multiplication pushes bookings to "extreme classes", by making the airline open more when its competitor is more open than usual, and close more when its competitor is more closed than usual. LOCO-based Forecast Multiplication also improves airline revenue management systems for cases of unusually high or unusually low total market demand for a specific isolated departure. This includes inadequate high protection levels in fully-restricted environments when total market demand is unusually low for a departure and the airline uses EMSRb with pick-up forecasting. Another example is when in an unrestricted environment total market demand is unusually high and the airline uses EMSRb, Q-forecasting with Frat5c, which is then not aggressive enough. Therefore FM1 tends to spread out the distribution of an airline's revenue over departure dates.

LOCO was found to be an easily accessible but not a perfect indicator of the competitor's availability. It is imprecise because it is discontinuous, and advance purchase restrictions of the competitor reduce the performance of LOCO by decreasing the number of possible LOCO of the competitor at a given time before departure. LOCO-based Forecast Multiplication thus performs better in unrestricted environments than in environments where the competitor uses advance purchase requirements.

When applied by both competitors using EMSRb in a fully-restricted environment, we found that FM1 can still increase both competitors' revenues. Even though the method is not theoretically justified anymore, it is important to study what could become a real world situation. In such a case, the positive impact of FM1 on the airlines is not robust, with revenues decreased for nearly half of the departure dates. Both airlines push their bookings to extreme classes even more than when only one airline implements FM1. However we found that on average, individual airlines' and total market revenues are not decreased.

4.2 Deliberate Competitive Overrides

In part 3.2 we described how a passive airline is impacted and adapts because of changes in its competitive revenue management environment. In this part, we take the point of view of an airline that would deliberately make a seat inventory competitive move. We present its available options and we describe their respective impacts. We use the results of part 3.2 and we keep the same assumption that the competitive effects on demand are larger than the imprecision of forecasters. We only consider for simplification the 2 airline duopoly revenue management game. Yet any reasoning made on the competitor's behaviour in the duopoly case can be extended to the equivalent resulting move of the sum of the competitors in a $n > 2$ airlines competitive scenario. We first examine possible active overrides of the automated revenue management system to respond to a move by the competitor. We then focus on deliberate competing moves that an airline could initiate from the EMSRb protection levels. We do not run specific simulations, but we link our findings to the results of the PODS simulations described in the Literature Review (Part 2.1).

4.2.1 Deliberate active reactions to competitor moves

In this part we present the responses an airline can take when it is subjected to a specific change in its competitive environment, that is unusual competitor availabilities or protection levels. The main challenge of deciding what action to take to respond to an unusual competitive environment lies in determining the reason behind it. Whether the current situation is due to an override by the competitor of its revenue management system or whether it is simply due to an unusual occurrence of market demand without override of its automated revenue management system leads to opposite recommendations. An analyst judgement seems necessary to identify whether the competitor could be overriding its system or whether it is simply experiencing an unusual occurrence of market demand but still using automated availability calculations. We directly use the conclusions of part 3.2.2, so our results hold with dependent demands between classes. In parallel to reading the following two paragraphs, the reader can also refer for better understanding to Fig 3.2.1, 3.2.2, 3.2.5 and 3.2.6.

Assume an airline duopoly nesting protection levels for higher fare classes, and assume that Littlewood's assumption of reached booking limits is met. Assume demands between fare classes are dependent, with sell-up but no buy-down of demand. If Airline 2 is unusually open for its low classes, it could be a "special offer" override decreasing protection levels for higher classes with steady demands to improve its load factors at the expense of its yields and expected revenues. As we showed in part 3.2.2, such a move should result in an increase in Airline 2 average spill-out for high-classes and Airline 1 average spill-in for high-classes. This "special offer override" by Airline 2 should improve the fare class mix of the non-constrained demand for Airline 1, which will necessarily increase the expected revenues of Airline 1. When the move of Airline 2 is sustained long enough, it is automatically taken into account by Airline 1's revenue management system that will protect more for high classes and optimize revenues. If the competitive move has not been sustained long enough, the best for Airline 1 is to override its system and protect more for high classes.

However if the unusual availability of Airline 2 is rather merely due to an unusually low occurrence of market demand, the reasoning of part 3.2.2 does not apply anymore. As explained in part 4.1 it is then better for Airline 1 to decrease protection levels because its non-constrained demand-to-come for high classes is prone to be smaller than usual.

Fig 4.2.1 displays a decision diagram for the response to a competitor unusually available. If the competitor is unusually open, it is necessary to know whether the competitor is overriding its system with steady demands, or whether it is just experiencing lower bookings than usual. In the first case, our protection levels for higher classes should be increased to better adapt to the predictable higher average spill-in for higher classes than usual. If the competitor move is sustained over many departure dates, the revenue management system automatically adapts in that manner. In the case where the competitor is simply experiencing lower bookings than usual, spill-in for higher classes should on the contrary be smaller than usual. Protection levels for higher classes should be decreased, such as with LOCO-based Forecast Multiplication.

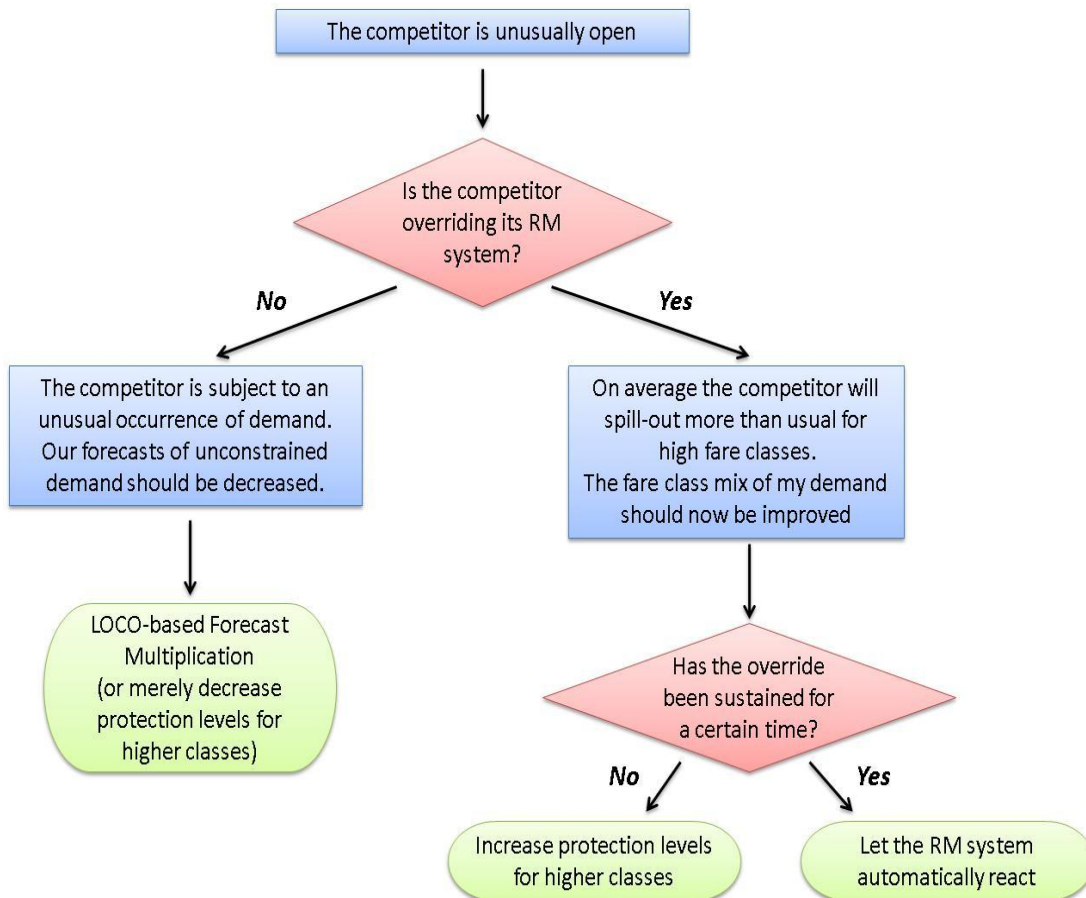


Fig 4.2.1 Decision diagram. RM response to a competitor unusually open.

Let us now make a similar reasoning if Airline 2 is unusually closed. We still assume an airline duopoly, setting nested protection levels for higher classes, where booking limits are reached (Littlewood's rule), with dependent demands between fare classes with sell-up but no buy-down of demand. If Airline 2 is unusually closed due to a deliberate aggressive override with steady demands where it increases protection levels for high classes, we showed in part 3.2.2 that it will lead to a decrease in the average spill-in experienced by Airline 1 for higher classes. By deteriorating the fare-class mix of the non-constrained demand for Airline 1, its expected revenues necessarily decrease. Airline 1 could adapt and decrease protection levels as its automated revenue management system would do anyway if the move by Airline 2 was sustained over many successive departures. This decrease in protection levels by Airline 1 is a lesser evil optimizing short term revenues knowing a given worse competitive environment. The other option for Airline 1 is to increase protection levels for higher classes, not to optimize its expected revenues in the short term, but to counter the move by Airline 2 by in turn hurting Airline 2's spill-in and hence Airline 2's expected revenues. Such a reaction by Airline 1 could avoid a downward spiral situation for Airline 1 where it experiences a worse fare-class mix of demand, decreases protection levels for higher classes, and experiences smaller average revenues at the expense of the Airline 2.

However if Airline 2 is unusually closed due to an unusually high occurrence of market demand, the reasoning of part 3.2.2 does not apply anymore. As explained in part 4.1 the best response for Airline 1 is then to increase protection levels for higher classes because its non-constrained demand-to-come for high classes is prone to be larger than usual.

Fig 4.2.2 displays a decision diagram for the response to a competitor more closed than usual. Again, it is necessary to know whether the competitor is overriding its system by setting unusually high protection levels for higher classes with steady demands, or whether the competitor is just experiencing a larger number of bookings than usual. In the second case our spill-in for higher classes should be larger than usual. Protection levels for higher classes should be increased, such as with LOCO-based Forecast Multiplication. In the first case, if it is that the competitor is aggressively overriding its revenue management system and setting unusually high protection levels, it is possible to optimize our short-term revenues by decreasing our own protection levels, because spill-in will probably be smaller than usual. Yet expected revenues would still remain lower than before the move by the competitor, so that we could be inclined to hurt back the competitor by in turn spilling-out less to him, and try to force the competitor to stop setting larger protection levels for high fare classes.

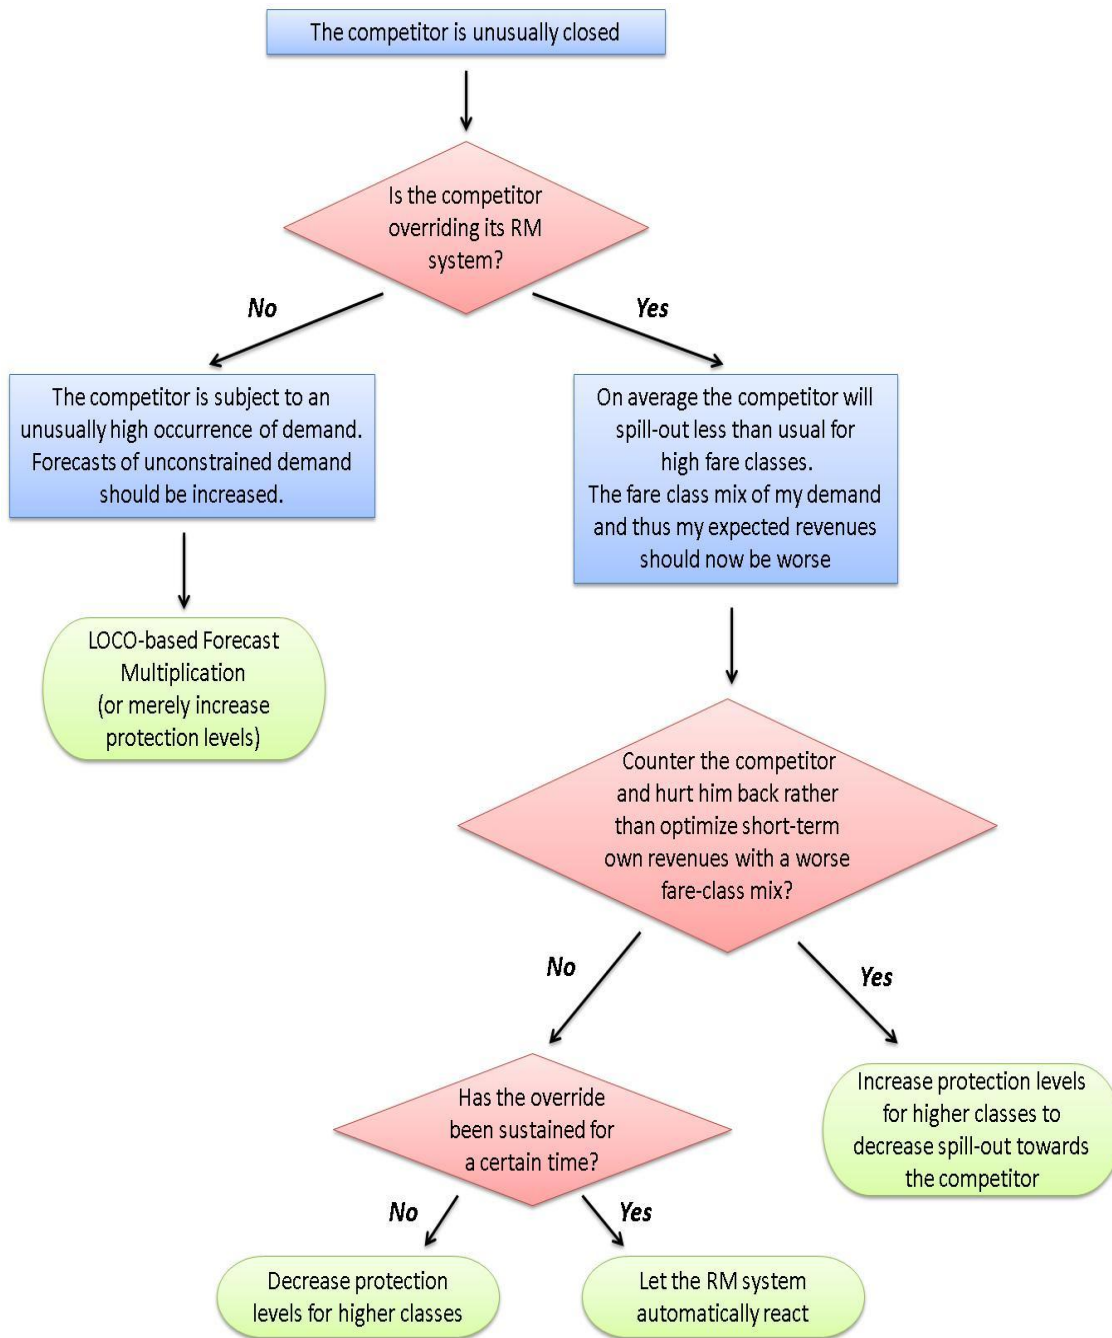


Fig 4.2.2 Decision diagram. RM response to a competitor unusually closed.

4.2.2 Deliberate initial competitive moves

In this part we consider the possible revenue management overrides with steady market demands that an airline can initiate from an EMSRb near-optimum. We use the results of part 3.2.2 and 4.2.1. **We assume that demands between classes are independent** (sell-up does not exist). When making such an initial override, an airline should think of second order effects: the impact on the competitor, its reaction, and the associated feedback effect on itself.

In 3.2.2 we saw that whatever the reaction of its competitor, if an airline overrides its revenue management system by protecting more (respectively protecting less) for high classes it will necessarily decrease (respectively increase) its competitor's expected revenues.

The feedback effect of a move by an airline on its own spill-in of demand depends on the reaction of its competitor, and for long-term moves on the competitor's optimization algorithm. Assume an airline is using EMSRb and has for only objective the optimization of its own revenues. Assume that EMSRb is a near optimum given external demands. Then the only way for an override of its EMSRb optimization to lead to an increase in its expected revenues using competitive considerations would be for this airline to increase its protection level for higher classes and have a feedback effect increasing its spill-in for these high classes. Fig 4.2.3 displays the consequences of overriding from the EMSRb protection level.

If the feedback effect of a move by an airline is to decrease its own spill-in of non-constrained demand for higher classes, then the fare-class mix of this airline's non-constrained demand is worsened, and expected revenues necessarily decrease from the initial EMSRb near-optimum. If the feedback effect of a move by an airline is to increase its own spill-in of non-constrained demand for higher classes, it can only increase expected revenues if the initial move was to increase protection levels for higher classes. If the initial move was to decrease protection levels for higher classes, then the feedback increase in spill-in of demand is due to closures of the associated higher classes. In that case the airline making the initial move cannot benefit from this higher spill-in of non-constrained demand, because it consists in additional non-constrained demand for classes that the airline has already closed. This explains why LOCO-open-matching (cf. Lua[28] or d'Huart and Belobaba[10]), by decreasing protection levels from the EMSRb near-optimum, always decreases expected revenues at the benefit of the competitor.

Hence for an override of an airline revenue management system to increase expected revenues, it is necessary that it increases a protection level for higher classes, and that the feedback effect on the airline overriding its system is an increase of the spill-in of non-constrained demand for the associated high classes.

First note that if protection levels are increased as a short-term move and the competitor does not react, the feedback effect will be a decrease in spill-in towards us. If protection levels are increased as a long-term move and the competitor just lets its automated system react, the feedback on spill-in depends on the optimization method he uses. However we showed in part 3.2.2 Proposition C that if the competitor also uses EMSRb, then spill-in towards us is decreased in most real world situations. This explains why LOCO-closure-matching (cf. Lua[28] or d’Huart and Belobaba[10]) should decrease revenues if both competitors are using EMSRb as a basis. Moreover if the competitor is aware of the move increasing protection levels by the first airline, it knows it will inevitably lead to a decrease in its expected revenues. Among its possible reactions, the competitor has little incentive to choose among its possibilities the one that will increase spill-in to the first airline (which is to under-protect for its high classes). It will rather choose a position that would in turn hurt the first airline rather than engage in a spiral situation where the first airline’s revenues are increased at the expense of its own revenues.

We therefore believe that no override from EMSRb based on competitive feedback considerations can realistically have strong chances to increase expected revenues.

Using results from 3.2.2 and 4.1 we can illustrate this result and assert that if both airlines were using EMSRb, it corresponds to a stable equilibrium because if one airline increases a protection level, there is no incentive for its competitor to make a decision increasing its spill-in towards the first airline. Therefore considering that EMSRb is a near optimum for given demands, both airlines using EMSRb is both a competitive near-optimum and a symmetric cooperation near-optimum for airlines that do not share revenues.

Override from EMSRb rule:	Decrease of nested protection level (top-down)		Increase of nested protection level (top-down)	
Competitive feedback effect	Spill-in decreases	Spill-in increases	Spill-in decreases	Spill-in increases
Impact on expected revenues	↓	↓	↓	↑ (seldom in practice)

Fig 4.2.3 Impacts of an override from the EMSRb rule on an airline’s revenues

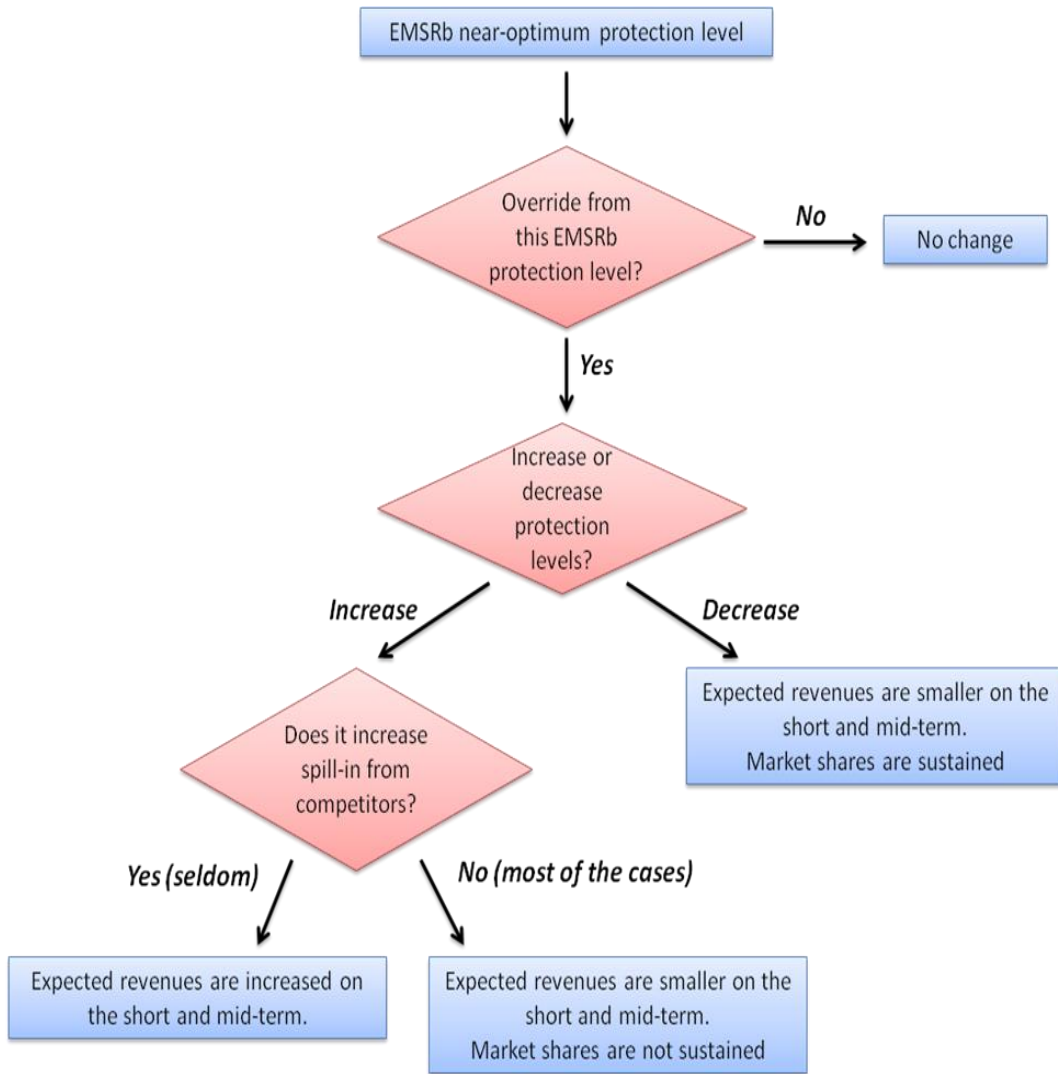


Fig 4.2.4 Decision diagram. Override from an EMSRb near-optimum protection level.

4.2.3 Conclusions

We based our study on competitive considerations, evaluating the possible successive stage strategies for airlines, and for our recommendations we assumed an airline's main objective when overriding its revenue management system as an initial move is merely to optimize its own revenues and not to hurt its competitors'.

We explained why in these conditions and assuming independent demands between classes, no initial override (short-term or long-term) by an airline from an EMSRb near-optimum can be considered to have high chances of increasing expected revenues. An initial override from EMSRb decreasing protection levels, such as LOCO-open-matching (cf. Lua[28] or d'Huart and Belobaba[10]), cannot increase expected revenues, but it can help maintain market shares. The only way an override by an airline could increase its own revenues would be that it set larger protection levels for higher fare classes than dictated by the EMSRb rule, and that in turn the competitor's reaction generated an increase in the overriding airline's incoming spill-in of demand for these higher fare classes. This can happen in rare cases, and should account for the few cases in which LOCO-closure matching (cf. Lua[28] or d'Huart and Belobaba[10]) can be revenue positive with a competitor using an accordion threshold method such as described in 4.1, but not if it uses EMSRb too. However as soon as the competitor is aware of competitive revenue management interactions, it does not have much incentive to react so as to decrease protection levels and increase its spill-out.

We can deduce that both airlines setting EMSRb protection levels can be considered as a stable equilibrium, as well as a near-optimum for both the seat inventory competitive game and the seat inventory cooperative game with no sharing of revenues between the airlines.

We also showed that in deciding how to react to an unusual competitive environment, it is necessary to determine whether the reason behind it is an override by the competitor or whether it is the result of the use of an automated system with an unusual occurrence of passenger demand. In the first case, the airline has the choice between using EMSRb protection levels to optimize its revenues knowing the given competitive environment, or taking action to hurt back the competitor if the initial move by the competitor was inevitably deteriorating the first airline's revenues at the expense of the competitor. If the unusual competitive environment is merely due to an unusual occurrence of demand, then a method such as LOCO-based Forecast Multiplication (Part 4.1) is most effective.

CHAPTER 5.

CONCLUSION

Since the 1980s, the airline industry has known two major changes: Deregulation, which led to an increase in competition, and the development of revenue management systems. Paradoxically the revenue management models used have not taken into account much competitive considerations. This thesis tackles the specific issue of reciprocal interactions between the revenue management practices of competing airlines. They have been mostly overlooked until now, because theoretical models of forecasters and optimizers for airline revenue management have considered most parameters such as price or demand as exogenous, isolated from the competitive environment. It is more realistic to consider that these parameters are partly endogenous, and determined by the simultaneous seat inventory decisions of each competing airline.

The objectives of this thesis are to evaluate the impact of competitive interactions given the current practice of airline revenue management, and more specifically fare class mix control, as well as to suggest possible adjustments to the current practice that would reasonably take into account the competitive environment. In the field of revenue management, the competitive environment is defined on the one hand by the long-term strategies of competitors, that are their rules to set protection levels set for higher classes. On the other hand, the revenue management competitive environment is defined in the short-term by the observed availabilities of competitors for each flight. The experimental tool used throughout the thesis is the Passenger Origin Destination Simulator (PODS), a software simulator of airline revenue management.

5.2 Research Findings and Contribution

We first provided a review of the state of the research on the relationship between competition and revenue management. On the one hand, there are very few theoretical studies that apply directly to airline revenue management. Netessine and Shumsky [31], Li et al. [24] as well as Li et al.[25] have come to similar conclusions under different sets of scenarios and restrictive conditions. There should exist a Nash equilibrium of protection levels if both airlines compete rationally, with perfect information, and take into consideration competitive interactions. Airlines competing would benefit from taking into consideration competitive effects in the field of revenue management. They would benefit even more if they cooperated for their allocation of seat inventory to different fare classes. For a given stochastic demand, the protection level of an airline optimizing its revenues with knowledge of competitive effects should be a decreasing function of the equivalent protection level set by its competitor. Last, total market protection levels for higher classes should be larger under competition than under collusion.

We also provided a review of past simulations made using the PODS simulator. The main purpose of these simulations was not to focus on the competition between revenue management systems, but to test whether some suggested improvements would work in competitive environments. We grouped their results and reached conclusions on competitive interactions. An airline improving its revenue management system most of the time leads to an increase of its revenues at the expense of its competitors. Yet one cannot speak of a zero-sum game as the increase in the airline's revenues is higher than the decrease in revenue of its competitors. It spills less high-fare class and more low-fare class passengers to competitors. Not only does this improve the leading airline's fare class mix and revenues, it also shifts the demand experienced by the competitor towards low fare classes. No matter its optimization method, the latter ends with a worse fare class mix and lesser revenues. Conversely, when a tested revenue management technique happens to decrease the innovating airline's revenues, it usually increases the passive airline's revenues by having improved its mix of demand through spill. When both airlines improve their revenue management system, both benefit. There is no long-term advantage of being the first or the last to move towards the improved revenue management system.

After this review of existing results, and as a first step towards understanding the mechanisms of competitive interaction in the field of airline revenue management, we compared an oligopoly situation with its equivalent monopoly situation. We presented a model and a general intuitive idea justifying that because of the notion of passenger spill between airlines, with the current state of revenue management, an oligopoly sets higher total market protection levels than a monopoly for an identical total market non-constrained demand for the same total market seat capacity. Spill corresponds to passengers who would as a first-choice book on a specific airline but who are rejected because the fare product they want to purchase is no longer available at the airline. In an oligopoly, spilled passengers are at least double-counted at the total market level, because they increment the forecasts of unconstrained demands at more than one airline. Spill of passengers, and thus such double-counting do not occur in a monopoly. This leads the oligopoly to set higher total market protection levels than the monopoly.

The idea presented is valid for N airlines with n classes, nested protection levels, and independent demands. It holds if demands between classes are dependent as long as the oligopoly does not have a total market capacity strictly higher than the monopoly. We proved the result for most real world cases when N airlines with the same n -classes fare structure use the EMSRb-rule to optimize their seat inventory, with independent demands for classes, and provided the initial share of total market non-constrained demand between airlines based on passenger preference is the same for all classes. The PODS simulations that we ran reveal that in practice the result holds with different revenue management optimization methods (EMSRb, Dynamic Programming for fully-restricted fare structure environment, EMSRb Q-forecasting for unrestricted environments) as long as the forecasters, and in particular the unconstrained, are accurate enough to reflect changes in demand due to competitive moves. In some of our simulations, the protection level for the highest classes set by an oligopoly is not higher than the protection level set by the monopoly, because the amplitude of passenger spill is not large enough to be accurately reflected in the forecasts.

After we examined how competitive interactions affect the current practice of revenue management at the level of a total market, we focused on the impact of seat inventory competition at the level of a single airline. We define the revenue management game as airlines setting nested protection levels for higher classes, with steady demand distributions, with pay-off functions equal to their expected revenues. We define a competitive move as a change in a nested protection level with demand distributions remaining steady.

We first highlight that airlines preferred to their competitor all else equal are less sensitive to their competitive environments. This suggests that in the field of seat inventory control, low-cost carriers are all else equal more sensitive to moves by legacy carriers than the opposite. We then point out that automated forecasters in airline revenue management systems have usually been kept simple and extrapolate past data without accounting for past competitive situations or current competitive situations. However, if sufficiently accurate, they adapt to competitive moves with a certain delay that we call the inertia of a forecaster. This remains true for airline revenue management systems that do not incorporate forecasters, but still involve comparing the current situation with past situations, such as what is implemented at some low-cost carriers. Because an airline's revenue management system automatically adapts to a competitor move which lasts for a certain amount of departure dates, the notion of Nash equilibrium is then not relevant.

Competitor moves that last at least as long as the inertia of the forecaster are automatically taken into account. We showed that in that case, the protection levels and the revenues of a "passive" airline letting its revenue management system run without overriding it are decreasing functions of the competitor's equivalent protection levels. This allows us to conclude that two airlines which have implemented an automated optimization method and do not override them should reach a stable equilibrium situation. Competitor moves that last for a duration shorter than the forecaster inertia are not taken into account by the forecaster or by the optimizer. The passive airline does not adapt. On average, it experiences higher spill-in for higher classes and higher revenues if the competitor increases lower classes booking limits and it experiences lower spill-in for higher classes and lower revenues if the competitor decreases its booking limits.

Some results were only demonstrated in reasonable scenarios, under the assumptions of independent demands for classes, of reached booking limits, and for airlines using an EMSRb method. They should still hold but be attenuated with dependent demands for classes. They match observations made on competitive interactions of airline revenue management systems in former complex PODS simulations described in the literature review. The latter seem to allow the extension of the results to other automated optimization methods than EMSRb.

Using these results, and evaluating the possible successive-stage strategies for airlines, we made recommendations assuming an airline's main objective is to optimize its own revenues and that demands for fare classes are independent. We believe that in such conditions, and assuming EMSRb is a near-optimum, no initial override (short-term or long-term) by an airline from EMSRb can have high chances of increasing expected revenues. An initial override from EMSRb decreasing protection levels should not increase expected revenues, but it can help maintain high load factors. To increase revenues, an override from EMSRb must increase protection levels and have the competitor's reaction generate an increase in incoming spill-in. This can happen in rare cases, but not if the competitor uses EMSRb too. These conclusions account for the poor revenue impacts of the practice of LOCO matching

(cf. Lua[28] or d’Huart and Belobaba[10]). As soon as the competitor is aware of competitive revenue management interactions, it does not have much incentive to react in a manner that will increase its spill-out. We deduced that for the two airline revenue management game, using EMSRb protection levels can be considered as a stable equilibrium, as well as a near-optimum for both the competitive game and the cooperative game with no share of revenues between the airlines.

We also showed that to decide how to react to an unusual competitive environment, it is necessary to determine whether it is due to an override by the competitor or whether it is the result of the use of an automated system with an unusual occurrence of passenger demand. In the first case, the airline has the choice between using EMSRb protection levels to optimize its revenues knowing the given competitive environment, or taking action non-optimal on the short-term but to counter a spiral-down situation with own expected revenues decreasing at the benefit of the competitor.

For cases when the unusual competitive environment is merely due to an unusual occurrence of demand, we introduced a method developed by the PODS Consortium called LOCO-based Forecast Multiplication. The general idea is that an airline more open in lower fare classes than usual for a given time frame before departure is likely to spill-out less than historically in the future time frames before the same departure. Its competitors should decrease their forecasts of unconstrained demand. Conversely, if the competitor is more closed than usual for a given time frame before departure, this competitor is likely to spill-out more than historically in the future time frames before the same departure. Forecasts of unconstrained demand should be increased.

LOCO-based Forecast Multiplication is a heuristic multiplication of additive pick-up forecasts. It is based on the comparison of the competitor’s current fare class availability with its historical corresponding availability. The indicator of availability used is LOCO, the “lowest open competitor class”. Different heuristic formulations are suggested. The assumptions made are that the non-constrained demand distributions are steady over time, that the competitor does not override its forecast or its optimization method. Overrides include LOCO-based Forecast Multiplication itself. We validated the method through various PODS simulations of both restricted and non-restricted fare structure environments.

When applied by one airline, we found that LOCO-based Forecast Multiplication can increase revenues significantly for this airline, with a smaller impact on its competitor. Over many departures, it pushes bookings to “extreme classes”, by making the airline open more when its competitor is more open than usual, and close more when its competitor is more closed than usual. LOCO-based Forecast Multiplication also improves airline revenue management systems when they have limitations in cases of unusually high or unusually low total market demand. This includes the over-protection of seats for higher fare classes of EMSRb, classic pick-up forecasting for fully-restricted environments when total market demand is unusually low, and the under-protection of seats for higher fare classes of EMSRb, Q-forecasting for unrestricted environments when total market demand is unusually high. Therefore LOCO-based Forecast Multiplication tends to spread out the distribution of an airline’s revenues over departure dates.

LOCO was found to be an easily accessible but not a perfect indicator of the competitor's availability. It is imprecise because it is discontinuous, and advance purchase restrictions of the competitor reduce the performance of LOCO by decreasing the number of possible LOCOs of the competitor at a given time before departure. LOCO-based Forecast Multiplication thus performs better in unrestricted environments than in environments where the competitor uses advance purchase requirements. When applied by both competitors using EMSRb in a fully-restricted environment, we found that LOCO-based Forecast Multiplication can still increase both competitors' revenues. Even though the method is then not theoretically justified in such a scenario, it is important to study what could become a real world situation. In such a case, the positive impact of LOCO-based Forecast Multiplication on the airlines is not robust, with revenues decreased for nearly half of the departure dates. Both airlines push their bookings to extreme classes even more than when only one airline implements LOCO-based Forecast Multiplication. However we found that on average, individual airlines' and total market revenues are not decreased.

All of our results are limited by the practical limitations of airline revenue management including:

- The accuracy of airline forecasters, and more particularly of demand unconstraining modules. When airline forecasters are not accurate enough to report changes in non-constrained demands such as caused by changes in spill-in, an oligopoly does not necessarily protect more for high classes, and our results on reactions to long-term changes by the competitor of its rule setting protection levels could not be verified. If an airline protects more (respectively less) for high classes, the reaction of the competitor could be after some time to increase (respectively decrease) protection levels for high classes too because of an inaccuracy of the unconstrained forecast of demand. Through simulations we noticed that such inaccuracy of forecasters, if any, concerns more high-fare classes, because the magnitude of spill (and sell-up) is smaller for such classes than for low-fare classes. Potential inaccuracy of forecasters only concerns our results on the behavior of current revenue management systems. It does not affect our recommendations for reactions to competitive moves and the theory behind LOCO-based forecast multiplication.
- The fact that in the real world, revenue management moves might not always be strictly justified by revenue management theory. If a competitor's seat inventory protection levels are not strictly justified by revenue management optimization theory, our results on how to react and set own protection levels hold. Yet again this fact affects our results on the behavior of current revenue management systems. It does not affect our recommendations for reactions to competitive moves.
- The fact that it is necessary to know whether competitor moves are due to unusual occurrences of demand or due to an aggressive override by the competitor. This is even more complicated than in the real world, revenue management moves might not always be strictly justified by revenue management theory (cf. preceding point).

5.3 Future Research Directions

We suggest here further investigations to extend our work.

We explained why as of today the revenue management competitive game cannot be considered as a perfect information game. It would be interesting to evaluate the benefits/impacts on the industry if the information about inventory levels was openly shared with competitors, and used to improve forecasts. How could such perfect information be used? Would such knowledge of the booking limits and bookings-in-hand of competitors benefit the industry as a whole? Would it benefit passengers?

As we showed that EMSRb constitutes a near optimum equilibrium for the competitive seat inventory control game as well as the cooperative game with no share of revenues, we believe that improvement to take competition into account in the field of revenue management will rely on improving forecasting methods. We suggested a heuristic adjustment of current additive pick-up forecasters. Yet the theory justifying the proposed heuristic does not hold if more than one competitor uses it. One could consider developing a forecasting methodology that would still hold if all competitors were to use it. One could think of evaluating total market non-constrained demand, first-choice non-constrained demand, and evaluate spills and total airline non-constrained demand. Furthermore, another indicator of competitor availability than LOCO could be considered. Accurate forecasting in competitive environments would be facilitated if information was openly shared (not in a collusion manner) with competitors.

All our work needs to be extended to airline networks. Our study focused on single market scenarios, but there could be some complications at the network level.

Last, research could be performed to link our results with the study of revenue management for airline alliances. Our work could be used as a starting point to study situations of origin-destination markets where alliance members operate separate flights but do not share capacity or revenues. Besides, as spill-of-demand also occurs when airlines share capacity on jointly operated flights, our model could be used to further the study of capacity/ revenue-sharing mechanisms for airline alliances.

REFERENCES / BIBLIOGRAPHY

- [1] Belobaba P. (1987). *Air travel demand and airline seat inventory management*. Doctoral Thesis, MIT.
- [2] Belobaba, P. and Wilson J. L. (1997). *Impacts of yield management in competitive airline markets*. *Journal of Air Transportation Management* 3(1) 3–9.
- [3] Belobaba P. , Hopperstad C. and Tam W (2008). *The PODS Primer (Version 1)*. MIT PODS Consortium.
- [4] Belobaba P. , Odoni A. and Barnhart C., *The Global Airline Industry (2009)*. John Wiley and Sons. *Fundamentals of Pricing and Revenue Management*, 73-112.
- [5] Boeing Commercial Airplane Group (1993). *Decision Window, Path Preference Methodology, Time mode*.
- [6] Boyd E. (2007). *The Future of Pricing*. Palgrave Macmillan.
- [7] Cléaz-Savoyen, R. (2005). *Airline Revenue Management Methods for Less Restricted Fare Structures*, SM Thesis, MIT.
- [8] Dar M. (2006). *Modelling the Performance of Revenue Management Systems in Different Competitive Environments*. SM Thesis, MIT.
- [9] Darot J. (2001). *Revenue Management for Airline Alliances: Passenger Origin Destination Simulation Analysis*. SM Thesis, MIT.
- [10] d'Huart O. and Belobaba P. (2009). *Impacts of Matching Competitor RM Seat Availability*. AGIFORS Revenue Management and Cargo Annual Conference Presentation, Amsterdam, Netherlands.
- [11] Ferea J. (1996). *Revenue Management Impacts of Multi-Hub Codeshare Alliances*. SM Thesis, MIT.
- [12] Fiig T., Isler K., Hopperstad C. and Belobaba P. (2010). *Optimization of mixed fare structures: Theory and applications*. *Journal of Pricing and Revenue Management*, 9, 152–170.
- [13] Gallego G. and van Ryzin (1994). *Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizon*. *Management Science*, 40, 8, 999-1020.
- [14] Gallego G., Krishnamoorthy S. and Phillips R. (2006). *Dynamic revenue management games with forward and spot markets*. *Journal of Pricing and Revenue Management*, Vol 51,10-31.
- [15] Gorin T. (2005). *Assessing low-fare entry in airline markets: Impacts of revenue management and network flows*. Doctoral Thesis, MIT.

- [16] Guo J.C. (2008). *Estimation of Sell-up Potential in Airline Revenue Management Systems*. SM Thesis, MIT.
- [17] Hopperstad Consulting (2005). *Passenger O/D Simulator PODS v.10 Technical Specifications*.
- [18] Isler K. and Imhof H. (2008). *A game theoretic model for airline revenue management and competitive pricing*. *Journal of Revenue and Pricing Management*, Vol 7, 4, 384-396.
- [19] Karjalainen, R. (1992). *The newsboy game*. Working paper, Wharton School, University of Pennsylvania.
- [20] Kayser M. R. (2008). *Revenue Management Methods for Multiple Fare Structure Environments*. SM Thesis, MIT.
- [21] Lautenbacher, C. J., S. Stidham. (1999). *The underlying Markov decision process in the single-leg airline yield management problem*. *Transportation Science*. 33(2) 136–146.
- [22] Lee A. (1998). *Investigation of Competitive Impacts of Origin-Destination Controls Using PODS*. SM Thesis, MIT.
- [23] Lee S. (2000). *Modelling passenger disutilities in airline revenue management simulation*. SM Thesis. MIT.
- [24] Li M.Z.F. , Oum T.H. and Anderson C.K. (2007). *An Airline Seat Allocation Game*. *Journal of Pricing and Revenue Management*. Vol. 6, 4, 321-330.
- [25] Li M.Z.F. , Zhang A. and Zhang Y. (2008) *Airline Seat Allocation Competition*. *International Transactions in Operations Research* 15(2008) 439-459.
- [26] Lippman, S. and McCardle, K. (1997). *The Competitive Newsboy*. *Operations Research* Vol. 45, No. 1, January-February 1997, 54-65.
- [27] Littlewood K. (1972). *Forecasting and control of passenger bookings*. Proceedings of the Twelfth Annual AGIFORS Symposium, Nathanya, Israel.
- [28] Lua W. F. (2007). *Matching of Lowest Fare Seat Availability in Airline Revenue Management Systems*. MSc Thesis. MIT.
- [29] Mahajan, S., G. van Ryzin. (2001). *Inventory Competition under Dynamic Consumer Choice*. *Operations Research*, 49(5), 646-657.
- [30] McGill J., van Ryzin G. (1999). *Revenue Management: Research Overview and Prospects*. *Transportation Science* 33(2) 233-256.
- [31] Netessine S. and Shumsky R. A. (2005). *Revenue Management Games: Horizontal and Vertical Competition*. *Management Science*, Vol. 51, No. 5, 813–831.

- [32] Parlar M. (1988). *Game theoretic analysis of the substitutable product inventory problem with random demands*. Naval research logistics, vol. 35, n°3, 397-409.
- [33] Talluri K. and van Ryzin G. (2004). *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.
- [34] Tam C.F.W. (2008), *Airline Revenue Management Based on Dynamic Programming Incorporating Passenger Sell-Up Behavior*. SM Thesis, MIT.
- [35] Tirole J. (1988). *The theory of industrial organization*. MIT Press.
- [36] Vanhaverbeke T.(2006). *Revenue Management Based on Dynamic Programming in Unrestricted and Simplified Fare Structures*. SM Thesis.
- [37] Wilson J. (1995). *The Value of Revenue Management Innovation in a Competitive Airline Industry*. SM Thesis, MIT.
- [38] Zickus J. (1998). *Forecasting for Airline Network Revenue Management: Revenue and Competitive Impact*. SM Thesis, MIT.

