

**Estimation of Sell-up Potential in Airline Revenue
Management Systems**

by

Jingqiang Charles Guo

B.S., Electrical and Computer Engineering
Cornell University (2005)


Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Masters of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008



© Massachusetts Institute of Technology 2008. All rights reserved.

Author
Sloan School of Management
May 9, 2008

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

Certified by
Cynthia Barnhart
Professor, Civil and Environmental Engineering, and Engineering Systems
and Co-Director, Operations Research Center
Thesis Reader

Accepted by
Dimitris Bertsimas
Professor, Sloan School of Management
and Co-Director, Operations Research Center

ARCHIVES

Estimation of Sell-up Potential in Airline Revenue Management Systems

by

Jingqiang Charles Guo

Submitted to the Sloan School of Management
on May 9, 2008, in partial fulfillment of the
requirements for the degree of
Masters of Science in Operations Research

Abstract

The growth of Low Fare Carriers (LFCs) has encouraged many airlines to remove fare restrictions (such as advance purchase requirements and Saturday-night stays) on many of their fare class products, leading to the simplification of fare structures in competitive markets. In the most extreme case, these markets have fare structures that are unrestricted; the fare class products differ only by price since they all lack restrictions.

In these unrestricted markets, passengers buy the lowest possible fare product since there are no longer any restrictions that prevent them from doing so. A forecasting method known as “Q-forecasting” takes into account the sell-up potential of passengers in forecasting the demand in each of the fare products in such markets. Sell-up occurs when passengers upon being denied their original fare class choice, decide to pay more for the next available fare class so long as the price remains below their maximum willingness to pay. Quantifying this sell-up potential either using estimated or input values is thus crucial in helping airlines increase revenues when competing in unrestricted fare markets.

A simulation model known as the Passenger Origin-Destination Simulator (PODS) contains the following 3 sell-up estimation methods: (i) Direct Observation (DO), (ii) Forecast Prediction (FP), and (iii) Inverse Cumulative (IC). **The goal of this thesis is thus to investigate and compare the revenue performance of the 3 sell-up estimation methods.** These methods are tested in a 2-airline (consisting of AL1 and AL2) unrestricted network under different RM fare class optimization scenarios. Both estimated and input sell-up values are tested on AL1 whereas only input sell-up values are tested on AL2.

The findings of the simulations indicate that using FP typically results in the highest revenues for AL1 among all 3 sell-up estimation methods. When compared against simple RM fare class threshold methods that do not consider sell-up, using FP results in up to a 3% revenue gain for AL1. Under some fare class optimization scenarios, using FP instead of input sell-up values even results in a revenue increase of close to 1%. These findings suggest that FP is robust enough under a range of fare class optimizers to be used by airlines as a sell-up estimator in unrestricted fare environments so as to raise revenues.

Thesis Supervisor: Peter P. Belobaba

Title: Principal Research Scientist, Department of Aeronautics and Astronautics

Thesis Reader: Cynthia Barnhart

Title: Professor, Civil and Environmental Engineering, and Engineering Systems
and Co-Director, Operations Research Center

Acknowledgments

Firstly, I must thank Dr. Peter Belobaba for being my research advisor at MIT. For the past 2 years, he has given me the opportunity to conduct airline revenue management research with the PODS consortium. I am grateful for the guidance and mentorship he has shown me during this time.

I must also thank Craig Hopperstad, for helping me understand the intricacies of the PODS simulator, as well as the various airlines in the PODS consortium, for their feedback on research directions as well as their financial support.

I also want to thank my academic advisor, Prof. Cindy Barnhart, for always being willing to offer help and advice.

To all my friends at MIT, especially James, Emmanuel and Chris, thanks for all the good times that we've shared. Life as a graduate student would have been much less fun and interesting without you.

Last but not least, I want to thank my parents for their unconditional love and support all these years. *Mummy* and *Papa*, this thesis is dedicated to you.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction	13
1.1	Revenue Management in the Airline Industry	14
1.2	New Challenges - Low-Fare Airlines and the Web	15
1.3	Improvements to Forecasting and Estimation Techniques	16
1.4	Thesis Objective	17
1.5	Thesis Organization	17
2	Literature Review	19
2.1	Airline Revenue Management	19
2.1.1	Forecasting	20
2.1.2	Inventory Seat Control	21
2.2	Low-fare Airlines and Less-Restricted Fare Environments	23
2.3	Revenue Management Tools for the New Environment	23
2.3.1	Q-forecasting	24
2.3.2	Hybrid-forecasting	24
2.4	Chapter Summary	25
3	Methods for Estimating Sell-up Potential	27
3.1	Price-Oriented Behavior, Sell-up and Historical Bookings	27
3.2	Estimation Methods	29
3.2.1	Direct Observation	29
3.2.2	Inverse Cumulative	30
3.2.3	Forecast Prediction	31
3.3	Sell-up Estimation in PODS	33
3.3.1	Fare-Ratio Sell-up Model	33

3.3.2	Booking Period in PODS	34
3.3.3	Estimation Using The Fare-Ratio Model	35
3.4	Chapter Summary	38
4	PODS Simulation Environment and Simulation Results	39
4.1	PODS General Architecture	39
4.1.1	Passenger Choice Model	40
4.1.2	PODS Revenue Management System	42
4.1.3	PODS Seat Allocation Optimizer Methods	42
4.2	Simulation Environment	43
4.3	Simulation Setup	45
4.4	Simulation Results	47
4.4.1	Test Case 0 - AT90 (AL1) vs. AT90 (AL2)	48
4.4.2	Test Case 1 - EMSRb (AL1) vs. AT90 (AL2)	48
4.4.3	Test Case 2 - EMSRb (AL1) vs. EMSRb (AL2)	51
4.4.4	Test Case 3 - DAVN (AL1) vs. AT90 (AL2)	54
4.4.5	Test Case 4 - DAVN (AL1) vs. EMSRb (AL2)	57
4.4.6	Test Case 5 - DAVN (AL1) vs. DAVN (AL2)	60
4.5	Chapter Summary	62
5	Conclusion	65
5.1	Summary of Findings	66
5.2	Future Research Directions	68

List of Figures

1-1	Differential Pricing - Fewer Unsold Seats and More Revenues	15
2-1	3rd Generation Airline Revenue Management System	20
2-2	Spiral Down Effect on Revenue	24
2-3	Overview of Q-Forecasting	25
3-1	Step 1: Estimation of $c_t^{50\%}$ at interval t	36
3-2	Step 2: Estimation of $c_t^{50\%} = \alpha + \beta.t$ across all 16 intervals to obtain α and β	37
3-3	Example - Plot of $f_{j,n,t}^{50\%}$ over time	37
4-1	PODS Structure	40
4-2	Network D6 (unrestricted) Network	44
4-3	Plot of FRAT5C Input Values	46
4-4	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 1, DM = 0.8)	49
4-5	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 1, DM = 1.0)	49
4-6	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 1, DM = 0.8)	50
4-7	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 1, DM = 1.0)	50
4-8	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 2, DM = 0.8)	52
4-9	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 2, DM = 0.8)	52
4-10	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 2, DM = 1.0)	53
4-11	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 2, DM = 1.0)	53
4-12	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 3, DM = 0.8)	55
4-13	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 3, DM = 1.0)	55
4-14	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 3, DM = 0.8)	56
4-15	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 3, DM = 1.0)	56
4-16	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 4, DM = 0.8)	58

4-17	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 4, DM = 1.0)	. . .	58
4-18	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 4, DM = 0.8)		59
4-19	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 4, DM = 1.0)		59
4-20	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 5, DM = 0.8)	. . .	60
4-21	AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 5, DM = 1.0)	. . .	61
4-22	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 5, DM = 0.8)		61
4-23	AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 5, DM = 1.0)		62

List of Tables

3.1	Flight h at booking interval t with $k = 10$ samples and $n = 4$ fare classes . .	29
3.2	Direct Observation Example with $k = 10$ samples	30
3.3	Flight h at booking interval t with $k = 4$ samples and $n = 4$ fare classes . .	30
3.4	Inverse Cumulative Example with $k = 10$ samples	31
3.5	Forecast Prediction Example with $k = 10$ samples (Iteration 1)	32
3.6	Forecast Prediction Example with $k = 11$ samples (Iteration 2)	32
3.7	Time Frames in PODS	34
3.8	Direct Observation Example with $f_{j,n}$ added for $n = 4$ at interval t	35
4.1	Initial Load Factor Thresholds for AT90	43
4.2	Fare Structure and Restrictions in Network D6 (unrestricted)	44
4.3	FRAT5C Input Values	46
4.4	Simulation Test Cases (DM values of both 0.8 and 1.0)	47
4.5	Revenue values of AL1 in Test Case 0	48
5.1	Revenues Changes for AL1 in switching from AT90 to EMSRb (Q-forecasting)	66
5.2	Revenues Changes for AL1 in switching from AT90 to DAVN (Q-forecasting)	66
5.3	Summary of the best-performing sell-up estimation method	67

THIS PAGE INTENTIONALLY LEFT BLANK

List of Abbreviations

DAVN	Displacement Adjusted Virtual Nesting, page 19
DO	Direct Observation, page 24
EMSR	Expected Marginal Seat Revenue, page 19
FP	Forecast Prediction, page 26
FRAT5	Fare ratio that causes 50% of the demand to sell-up, page 28
FRAT5C	a particular set of input FRAT5 values, page 38
IC	Inverse Cumulative, page 25
LFCs	Low Fare Carriers, page 13
O-D	Origin-Destination, page 13
PODS	Passenger Origin-Destination Simulator, page 13
RM	Revenue Management, page 2
WTP	Willingness to Pay, page 13

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

The rise of Low Fare Carriers (LFCs) in numerous markets has led to the simplification of fare structures in these markets. There now exist fare products that are completely undifferentiated in terms of (i) having no restrictions or (ii) having the same set of restrictions; these products differ only by price. As a consequence, the assumption that the demand of each of these products is independent no longer holds since (for a particular set of undifferentiated fare products) passengers will always buy the lowest possible fare possible.

For an airline to maximize its revenues, it makes sense for it to make each passenger choose from the set of undifferentiated fare products the fare level that is as close to his or her maximum Willingness to Pay (WTP) as possible. Hence, there is a need to model and quantify this sell-up potential of passengers i.e. the probability that a passenger will be willing to buy a higher-priced fare class if the current one is unavailable.

This thesis focuses on Sell-Up Estimation, a technique that is used by a network R.M. system in order to forecast the (corrected for sell-up) demand in the individual fare class products. **The goal of the thesis is thus to comprehensively investigate and compare (in terms of revenue) the different estimation methods available in the Passenger Origin-Destination Simulator (PODS).** PODS is a simulation model that we will present later in Chapter 4 to simulate both airline competition and passenger demand in defined markets.

1.1 Revenue Management in the Airline Industry

Revenue Management has its roots in the rudimentary forecasting and overbooking control systems that a few airlines implemented before the 1970s. (see McGill [27]) Such systems were reasonably successful, consequently paving the way for the greater adoption of mathematical models and approaches in airline reservation systems.

In the 1970s, some airlines started departing from the practice of only offering a single (full) fare product for a given O-D market. For instance, BOAC (the predecessor to today's British Airways) charged discounted fares to passengers who booked at least 3 weeks before the departure date. Such airlines can be considered to be the first to practice differential pricing, since they offered more than one price (fare-product) for travel in the same Origin-Destination (O-D) market. These airlines can also be considered to be the first to practice inventory control, since they now also controlled how many seats they should make available for each of the two fare-products (by setting booking limits on the discounted fare product). During this period, airlines could only practice differential pricing to a limited degree since the airline industry was still regulated i.e. a government agency in each country determined the prices of the air tickets that the airlines sold.

Despite this regulated environment consisting of only 2 fare products (full and discounted), the benefits of practicing differential pricing in terms of incremental revenue is readily observed as shown in Figure 1-1. With just the full-priced fare product being offered, an aircraft would likely depart with unsold seats. With differential pricing, more of the otherwise empty seats can be sold, thereby raising revenues. Indeed, differential pricing has become an important element of revenue management. (see Belobaba [6])

With the end of airline deregulation in 1979 in the USA, airlines now had the freedom to practice differential pricing since they had full control of deciding (i) the number of different fare-products to offer in a given O-D market, and (ii) the prices of these fare-products. In theory, this meant that airlines could maximize their revenues since by offering more fare-products priced at different prices, there is a greater chance of offering a fare product that is very close to a given passenger's maximum willingness to pay.

In order for passengers to reveal their true maximum willingness to pay, differential pricing needs to present passengers with a trade-off between price levels and inconveniences. Such inconveniences include advance purchase requirements, minimum stay periods and

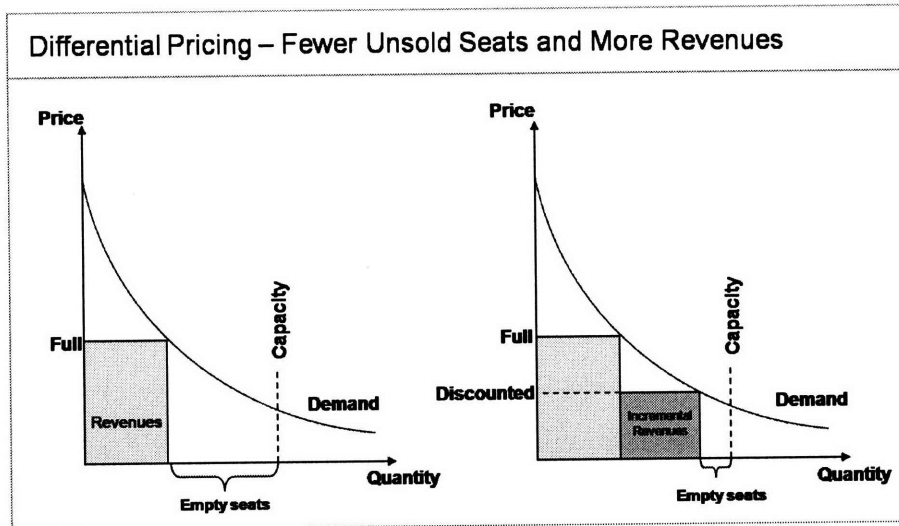


Figure 1-1: Differential Pricing - Fewer Unsold Seats and More Revenues

the lack of refundability. The general assumption is that business-oriented passengers are willing to pay more than leisure-oriented passengers in return for tickets that have fewer restrictions. Indeed, such restrictions achieved substantial success in segmenting passengers into business and leisure groups.

During this period, airlines also began to realize the importance of a reasonably accurate forecasting system. Since business-oriented passengers tend to book later (i.e. closer to the departure date) than leisure-oriented passengers, good estimates of the number of business-oriented passengers are needed so that sufficient seats are set aside in the higher-priced fare classes. Good forecasting is thus crucial in optimizing the allocation of seats among the various fare classes.

1.2 New Challenges - Low-Fare Airlines and the Web

Traditional airlines (often referred to as legacy carriers) have in recent years faced two major challenges. The growth of low-fare carriers in numerous markets has reduced their ability to practice differential pricing, just as the rise of internet travel sites has undermined their pricing power.

Low fare carriers are known for selling tickets in a given O-D market that are often priced lower than comparable ones from the legacies. They are also noted for implementing simplified fare structures in markets that they operate in so as to offer a less confusing

and thus more appealing product to passengers. Such simplified fare structures are often implemented by removing many of the restrictions (such as Saturday night stays) that have been successful in segmenting passenger demand. Traditional carriers have little choice but to simplify their own current fare structures in order to compete with the low-fare entrants, and thus have markedly diminished differential pricing ability.

Similarly, the rise of online travel sites such as Expedia and Orbitz has enabled passengers to easily compare fare product offerings from different airlines for the same O-D market. This has dramatically increased the level of transparency in terms of market knowledge of available airfares. Consequently, airlines have reduced pricing power since they are now compelled to charge equivalent fares in competing for passengers.

1.3 Improvements to Forecasting and Estimation Techniques

Airlines have looked towards the further development of revenue management techniques, especially in forecasting and estimation, in order to respond to the challenges as mentioned in Section 1.2.

Forecasting (for a given flight) refers to using historical unconstrained booking demand in order to estimate the number of future bookings-to-come by fare product. The importance of having accurate forecasting data, as already mentioned at the end of Section 1.1, is made even pertinent in markets that have simplified fare structures. As mentioned at the beginning of Chapter 1, since there now exist fare products that are undifferentiated in terms of restrictions, passengers will always buy the lowest-priced fare product. As a consequence, the assumption that the demand of each of these products is independent no longer holds.

From the airline's perspective, it thus makes sense to make each passenger choose from the set of undifferentiated fare products the fare level that is as close to his or her maximum willingness to pay as possible. Hence, there is a need to model and quantify this sell-up potential of passengers i.e. the probability that a passenger will be willing to buy a higher-priced fare class if the current one is unavailable. **Accurate forecasting thus requires having accurate values of these sell-up probabilities.**

Several estimation methods have thus been developed for use in PODS and will be explained in greater detail in the relevant sections.

1.4 Thesis Objective

A modern revenue management system requires a reasonably accurate forecasting tool, which in turn requires accurately estimating the sell-up probabilities.

Existing research on the revenue improvements arising from using different forecasting techniques exists. However, the impact of using different sell-up estimation methods on airlines' revenues has yet to be comprehensively investigated.

As such, the objective of this thesis is to provide an investigation into the effectiveness of the 3 sell-up estimation methods available in PODS in raising airline revenues. This investigation is performed in both a 2-carrier unrestricted network.

1.5 Thesis Organization

This thesis consists of three main parts: (i) the literature review, (ii) the PODS approach to simulation, and (iii) the PODS simulation results.

In Chapter 2, there will be a survey of existing research on revenue management with a focus on unrestricted and mixed fare structures. Topics covered include forecasting, various RM models, and the growth of simplified fare structures. The aim of this chapter is to provide the rationale for estimating sell-up through the WTP of passengers.

In Chapter 3, there will be an detailed explanation on sell-up estimation. The focus of this explanation will be on how each of the 3 estimation methods computes the sell-up probability values using the historical booking database of a typical RM system, as well as the fare-ratio sell-up model used by PODS to compute these probability values.

Chapters 4 first describes the PODS environment and the simulation setup in the unrestricted fare environment, before presenting the simulation results in detail.

Finally, Chapter 5 summarizes the results of the simulations and provides two directions for future research.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Literature Review

As revenue management has grown in sophistication over the past 3 decades, the quantity of relevant literature has similarly increased. Reviewing several of the most important published works related to revenue management in this chapter will help to (i) explain the use of revenue management in the airline industry, and (ii) explain key concepts that are needed in subsequent parts of this thesis.

This chapter begins with an historical overview of revenue management in the airline industry. The next section of this chapter considers the effects of LFCs and unrestricted fare structures on traditional revenue management methods. The final section describes 2 revenue management techniques involving forecasting. Since these 2 forecasting methods make use of sell-up estimation, they also provide an appropriate segue into the description of the sell-up estimation methods in the next chapter.

2.1 Airline Revenue Management

As already mentioned in the beginning of Section 1.1, revenue management involves using inventory control and differential pricing to maximize revenues. Although this maximization goal has remained the aim of revenue management, the techniques have changed over time.

Both Barnhart, Belobaba and Odoni [1], as well as McGill and van Ryzin [27] provide excellent coverage on the history of RM in the airline industry. The evolution of RM systems has traced the development of RM practices; the first of such systems appeared in the 1980s and mainly relied on overbooking (as described in Thompson [33] and Littlewood [26]) so as to increase revenue. By intentionally accepting more reservations than seats available

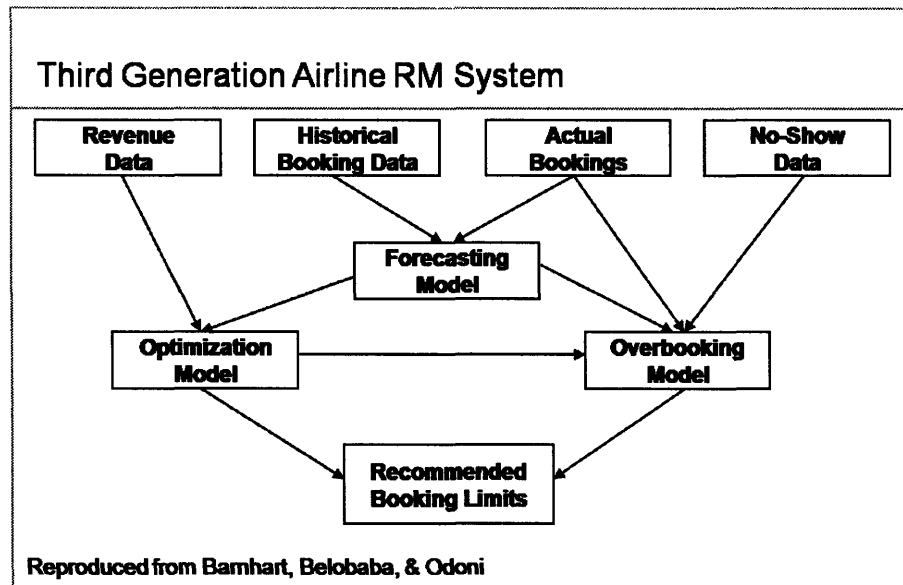


Figure 2-1: 3rd Generation Airline Revenue Management System

for a given flight, an airline reduces revenue loss from passenger no-shows and last-minute cancellations. Second-generation systems added the capability to control fare-class inventory using threshold methods. The most recent version of such systems have been used by airlines since the 1990s. As shown in Figure 2-1, these third-generation systems utilize additional RM components such as forecasting and optimization to generate recommended booking limits for each flight and fare class.

2.1.1 Forecasting

For a given flight, forecasting involves using unconstrained historical booking data to estimate the number of future bookings-to-come for each fare product. As already mentioned in Section 1.1, higher-fare (e.g. business-oriented) passengers tend to book closer to the departure date than lower-fare (e.g. leisure-oriented) passengers. Accurately predicting the expected number of passengers in these 2 categories thus enables the airline to set aside sufficient seats early on in the booking process by not selling too many of them to the low-fare passengers. These unsold seats can later be made available in the higher-priced fare classes to business travelers closer to the departure date.

Various forecasting methods exist, the most popular of which are simple time series forecasting and pick-up forecasting. Simple time series forecasting involves averaging the

number of bookings at every historical time period to forecast the expected number of bookings at the next time period. In contrast, pick-up forecasting involves first averaging the number of incremental bookings at every historical time period. This value is then added to the number of bookings at the most recent time period to give the expected number of bookings at the next time period. As noted in Wickham [34], pick-up forecasting is more commonly used in the airline industry than time-series forecasting, and will thus consequently be used in the rest of this thesis as well. Further information on other forecasting techniques such as regression forecasting can also be found in Zickus [38] and Gorin [19].

2.1.2 Inventory Seat Control

Inventory seat control plays a crucial role in setting the booking limits among the various fare classes on a given flight so as to maximize revenues for the airline. There are 2 main types of inventory control algorithms, each providing a different level of optimization at the flight leg level or the OD path level.

Leg-based Control (Fare Class Yield Management)

Leg-based controls allocate seats among the various fare classes on each leg within an airline's network. The fare class allocation most popularly used by airlines involves the serial "nesting" of fare classes - a technique first developed by Littlewood [26] at BOAC for the case of two fare classes. Instead of allocating seats separately among the two fare classes, this nesting approach limits the number of seats sold in the lower fare class so as to protect the number of seats sold in the higher fare class. This allocation is based on the forecasted demand in each class, as well as on the expected seat revenue.

Belobaba (see [2] and [3]) subsequently extended the nested seat allocation method from two classes to multiple fare classes by creating the Expected Marginal Seat Revenue (EMSR) heuristic. EMSR sets booking limits using the expected marginal revenue of each additional seat i.e. the the probability of selling that additional seat in a given fare class multiplied by the average price of a ticket in that fare class. The probability of selling each additional seat decreases as more and more seats become protected in a particular fare class. The protection level for that fare class is thus obtained when the EMSR equals the average price of a ticket in the next lower-priced fare class.

In 1992, Belobaba [4] updated his heuristic to protect joint upper fare classes from the

next lower-priced fare class. This approach, known as EMSRb, has become one of the most popularly-used fare class allocation methods used by airlines. Further information on the EMSRb heuristic can be found in Belobaba [10], Lee [25] and Williamson [35].

Optimal formulations to the multiple nested fare class problem exist. However, these approaches, as proposed by Brumelle and McGill [14], Curry [17], and Wollmer [37] require substantially more computation effort. More importantly, they do not provide significant revenue advantages over the EMSRb heuristic.

Network OD Control

By only considering bookings at the flight leg level, leg-based control optimizers do not fully account for O-D path bookings that can either comprise of multiple legs or a single leg (for a non-stop service). As such, they have a greater tendency of favoring local passengers on a particular flight leg instead of higher revenue O-D passengers on paths that include that particular leg. This spurred research into developing algorithms that determined booking limits at the path level.

Developed by Smith and Penn [31], one such approach uses “virtual buckets” to compare the network value of local and connecting fare classes. This virtual bucket method, specifically known as Displacement Adjusted Virtual Nesting (DAVN) , is used during simulations later in this thesis. DAVN combines O-D forecasting with leg-based seat inventory control, and uses a deterministic linear program (LP) to maximize the total network revenue. This involves first calculating a “pseudo fare” for each fare class on each leg in the network. If a multiple-leg itinerary is accepted at the expense of denying one or more single-leg itineraries on the constituent legs, then summing the “pseudo fares” on the affected legs provides a value of the displacement costs of accepting the multiple-leg itinerary. This displacement value is then subtracted from the fare of the multiple-leg itinerary to correct the revenue value for network displacement effects. Detailed information on virtual bucketing and DAVN can be found in Lee [25] and Williamson [35].

An alternative approach developed for O-D control involves the use of bid-prices, where a booking is accepted as long as the associated fare is not less than the bid price set for that particular itinerary. These bid-price algorithms include the Network Bid Price method, the Heuristic Bid Price method developed by Belobaba [5], and the Prorated Bid Price method (see Bratu [13]).

2.2 Low-fare Airlines and Less-Restricted Fare Environments

As already mentioned in Section 1.2, the rise of low fare carriers significantly changed the airline industry landscape. Such airlines not only charge fares that are lower than comparable ones from traditional carriers, but also adopt simplified fare structures by eliminating restrictions such as Saturday-night stays. (Gorin [20], and Dunleavy and Westerman [18] provide additional information about LFCs.) In addition, the popularity of online booking sites also increased pricing transparency for consumers. According to Kuhlmann [24], this compelled airlines to charge the same fares. The combination of these two factors led traditional airlines to similarly lower their fares and simplify their fare structures (see Windle and Dresner [36]). As a result, they reduced their ability to segment passenger demand into business and leisure types since there now exist fare classes that differ only by price and not by restrictions.

As pointed out by Ratliff and Vinod [28], the assumption that the demand of each fare class is independent thus no longer applies because passengers will always buy the lower possible fare class, which may be priced below their WTP. This “buying-down” behavior causes revenues to spiral down in airlines that use traditional forecasting systems.

As shown in Figure 2-2, traditional forecasting systems use historical bookings to produce future booking forecasts. When “buy-down” happens, there are fewer bookings observed in the higher-priced fare classes. The forecaster thus predicts fewer future bookings for these fare products, causing the optimizer to make more seats available in the lower fare classes. This creates a cycle where there are fewer and fewer bookings in the higher-priced fare classes because an increasing number of seats are sold in the lower fare classes. Cooper et al [16] provides additional details about this spiral-down effect.

2.3 Revenue Management Tools for the New Environment

To counteract the spiral down effect, new forecasting methods have been developed. These methods rely on the estimated WTP of passengers instead of assuming the independence of fare class demand in forecasting future bookings.

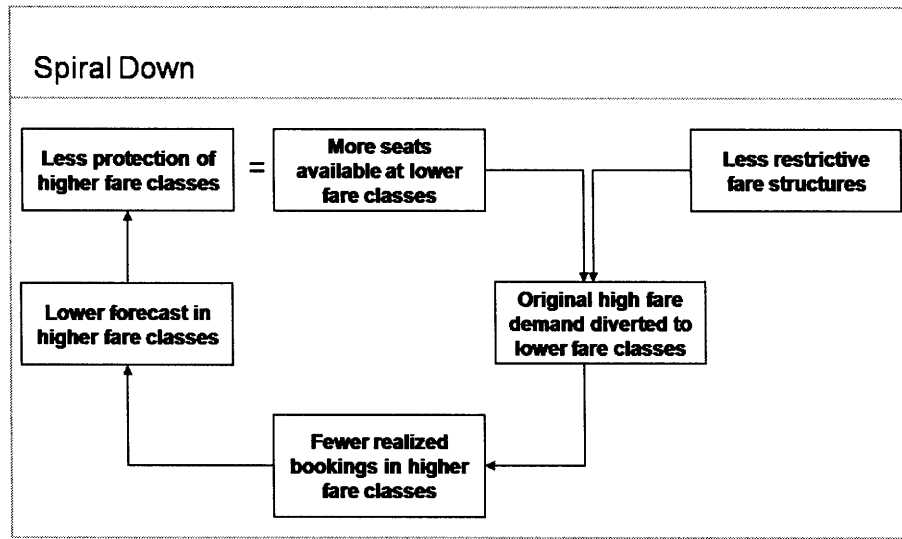


Figure 2-2: Spiral Down Effect on Revenue

2.3.1 Q-forecasting

Belobaba and Hopperstad (see [7] and [9]) developed the “Q-Forecasting” method for use in fully unrestricted fare environments. This approach first forecasts demand in the lowest fare class (also known as Q-class) before using estimated WTP values of passengers to compel sell-up by closing lower fare classes, thereby increasing revenues for the airline.

As depicted in Figure 2-3, historical bookings are first converted into an equivalent number of “Q-bookings” before being summed across all fare-classes. By strategically repartitioning fare class demand and closing down lower fare classes based on the estimated passenger WTP, a fraction of future bookings will be forced to sell-up to higher fare classes.

Cleaz-Savoyen [15] concluded from simulations that Q-forecasting enables airlines to regain some lost revenues when used in unrestricted fare environments.

2.3.2 Hybrid-forecasting

Many markets continue to offer semi-restricted fare environments. Such environments typically consist of undifferentiated (in terms of restrictions) lower-price fare classes and several differentiated higher-priced fare classes. The presence of these restrictions in the higher fare classes prevents all passengers from buying down into the lowest possible fare class, making Q-forecasting not the most appropriate method to be used under such circumstances. At the same time, the possibility of buy-down among the lower-price fare classes makes

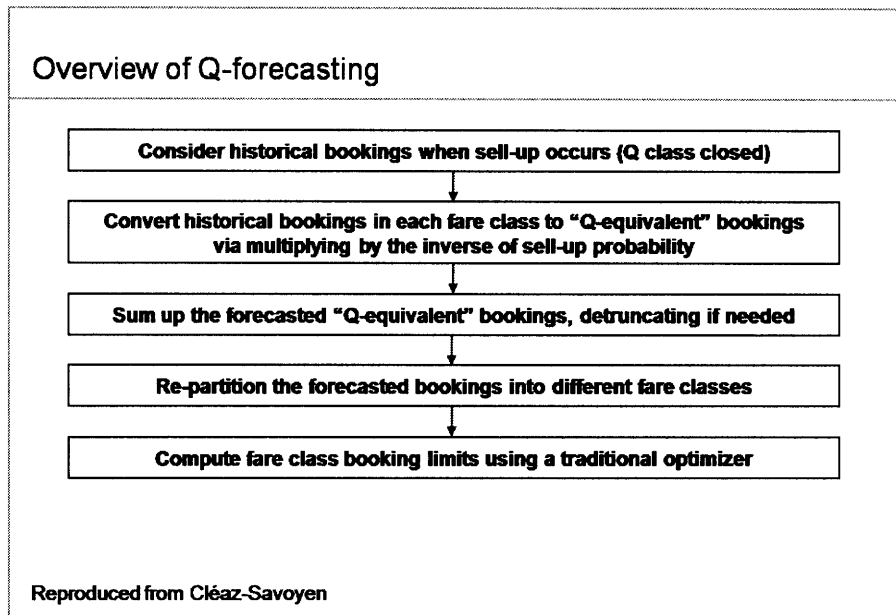


Figure 2-3: Overview of Q-Forecasting

traditional forecasting methods ineffective since the fare class demand is not independent.

Boyd and Kallesen [12] thus developed hybrid forecasting for use in semi-restricted fare environments. Hybrid forecasting first classifies all bookings into price-oriented and product-oriented bookings, before separately forecasting future demand for each group. Since price-oriented bookings are made by passengers who always buy the lowest-possible fare class, Q-forecasting can be used for this group. Similarly, as product-oriented passengers choose a fare class based on the different restrictions of each fare class, this allows traditional forecasting to be used on this group. More information on hybrid forecasting can be found in Reyes [30], who found that hybrid forecasting contributed to a 3% rise in revenue over standard pick-up forecasting.

2.4 Chapter Summary

This chapter has provided the historical development of revenue management and forecasting methods so as to illustrate the importance of estimating sell-up potential. As mentioned in Section 2.3.1 and Section 2.3.2, modern forecasting techniques that rely on estimating the sell-up potential through the WTP of passengers increase airline revenues. There is thus a need to examine the revenue impacts of using different sell-up estimation methods.

However, before such revenue comparisons can be made, a detailed explanation of these 3 sell-up estimation methods evaluated in this thesis needs to be made. As such, the complete description of these methods is found in the next chapter.

Chapter 3

Methods for Estimating Sell-up Potential

This chapter describes in detail 3 methods of estimating sell-up. The primary aim of this chapter is to illustrate how each method estimates sell-up using the historical booking database of a typical RM system. The secondary aim of this chapter is to explain the fare-ratio sell-up model used by PODS, which relates sell-up probabilities with fare ratios. As such, the first half of this chapter describes the general concept of how the methods use historical booking data from the RM system to estimate sell-up. The second half of this chapter discusses the fare-ratio model that PODS uses in conjunction with the historical booking data and one of the three methods in order to estimate sell-up. A further explanation of the PODS environment is found in the next chapter together with the simulation description and results.

3.1 Price-Oriented Behavior, Sell-up and Historical Bookings

As already mentioned in Section 2.3.2, price-oriented behavior involves passengers comparing fare classes solely using fare prices. Such behavior is typically exhibited in semi-restricted fare environments by leisure passengers, who buy the lowest priced fare class, Q-class, whenever possible. However, price-oriented passengers can include both business and leisure passengers in unrestricted fare environments since the fare classes differ only

by price. Sell-up occurs when price-oriented passengers, upon being denied their intended booking in Q-class for a given O-D market, pay more for the next available higher fare class in the same O-D market as long as the fare does not exceed their maximum WTP.

Since passengers have different maximum WTP values, a randomly chosen price-oriented passenger has a different probability value associated with paying at most each of the higher-priced fare classes. Consequently, an alternative to calculating the expected WTP value for the random passenger is calculating the conditional probability values of the passenger paying at most the price of each of the fare class products, given that the passenger would have bought Q-class. These values are thus the probabilities of the airline selling up a booking from Q to each of the fare class products j , when j is the lowest open available fare class.

The sell-up probabilities change over time during the booking period, which starts when seats on the flight serving the O-D market are first open for sale and ends on the departure day of the flight. As mentioned in Section 1.1, business passengers both book closer to the departure date than leisure passengers, and have greater WTP values than them. The overall sell-up probabilities for a flight thus increase as the booking period progresses towards the departure date.

As earlier shown in Figure 2-1, current third-generation RM systems have a historical booking database that stores past booking records of flights. For a particular flight number h on a given day of the week, the database contains fare class booking numbers of past such occurrences of h at fixed time intervals during the booking period. Over time, more booking data is accumulated in the database as the number of occurrences of h increases.

One way in which such booking data can be presented is shown in Table 3.1. Consider the case of flight number h at a particular booking time interval t that offers 4 fare classes (labeled in order of decreasing fare value as Y, B, M and Q) that are undifferentiated in terms of restrictions. To aid in developing subsequent examples, we further assume that fare classes can only be closed at the beginning of any given booking time interval. This results in bookings being made only in the lowest open fare class j during the entire duration of t . Each of these past occurrences of h at t results in a booking sample.

Since the number of samples can differ among the fare classes, any meaningful comparison using fare class booking numbers requires the value of the average bookings $b_{j,k}$ to be calculated for each possible fare class j given that there are k samples in total. $b_{j,k}$ is

Lowest Open Fare Class j	Fare Class Name	Bookings in Sample No.										Number of Samples in j	Bookings	
		1	2	3	4	5	6	7	8	9	10		Total	Average $b_{j,k}$
1	Y			10								1	10	10
2	B				17		13					2	30	15
3	M		22			16				22		3	60	20
4	Q	44						38	38		40	4	160	40

Table 3.1: Flight h at booking interval t with $k = 10$ samples and $n = 4$ fare classes

recalculated whenever a new sample is added.

3.2 Estimation Methods

The 3 methods of estimating sell-up that are covered in this thesis are i) Direct Observation (DO), ii) Inverse Cumulative (IC), and iii) Forecast Prediction (FP). As noted by Belobaba [8], IC and FP have been available in PODS since 2006. Based on feedback from a PODS Consortium airline member, DO was implemented in PODS by the end of 2007 (See Guo [22]).

3.2.1 Direct Observation

For classes $j = 1, 2, \dots, n$ where n refers to fare class Q, Direct Observation (DO) uses the average bookings $b_{j,k}$ for each j to compute the sell-up probabilities from n to j . DO views each $b_{j,k}$ as an independent observation of the average bookings in j received by the airline during interval t ; there are thus n such observations (1 for each fare class). Since passengers always buy the cheapest fare class whenever possible, the probability $p_{j,k}$ of selling up a booking from n to j is given by:

$$p_{j,k} = \frac{b_{j,k}}{b_{n,k}} \quad (3.1)$$

Observe that $p_{l,k} \leq p_{m,k}$ for $l, m \in [1, n], l < m$ as passengers who are willing to buy l must also be willing to buy m . Intuitively, this will occur as long as $b_{l,k} \leq b_{m,k}$. When a new sample is added, $b_{j,k}$ and thus $p_{j,k}$ are recalculated for the new value of k .

Lowest Open Fare Class j	Fare Class Name	Number of Samples in j	Bookings		$p_{j,k}$
			Total	Average $b_{j,k}$	
1	Y	1	10	10	0.250
2	B	2	30	15	0.375
3	M	3	60	20	0.500
4	Q	4	160	40	1.000

Table 3.2: Direct Observation Example with $k = 10$ samples

Lowest Open Fare Class j	Fare Class Name	Bookings in Sample No.				Number of Samples in j	Bookings	
		1	3	4	5		Total	Average $b_{j,k}$
1	Y		10			1	10	10
2	B			17		1	17	17
3	M				16	1	16	16
4	Q	44				1	44	44

Table 3.3: Flight h at booking interval t with $k = 4$ samples and $n = 4$ fare classes

Using the same data from Table 3.1, the sell-up probabilities calculated using DO are shown in Table 3.2.

3.2.2 Inverse Cumulative

There may be instances when $b_{l,k} > b_{m,k}$ for $l, m \in [1, n], l < m$. An example of this is shown in Table 3.3 where only the 1st, 3rd, 4th and 5th samples from Table 3.1 are extracted (i.e. $k = 4$). If DO is used to calculate $p_{j,k}$, then $p_{l,k} > p_{m,k}$ for $l = 2$ and $m = 3$. Such probability value inversions do not agree with the price-oriented behavior of passengers, which states that passengers who are willing to buy l must also be willing to buy a cheaper fare class m ; there cannot be more passengers willing to buy up to l than those willing to buy up to m . An alternative approach to computing sell-up is thus needed.

The Inverse Cumulative (IC) approach views each $b_{j,k}$ as concurrent observations during t in computing the sell-up probabilities from n . Since each booking corresponds to a passenger, the total number of passengers is $\sum_{i=1}^n b_{i,k}$. This is also the number of passengers who are willing to buy n . The number of passengers who are willing to buy l is thus $\sum_{i=1}^l b_{i,k}$. The probability $p_{j,k}$ of selling up a booking from n to j is given by:

Lowest Open Fare Class j	Fare Class Name	Number of Samples in j	Bookings			
			Total	Average $b_{j,k}$	$\sum_{i=1}^j b_{i,k}$	$p_{j,k}$
1	Y	1	10	10	10	0.118
2	B	2	30	15	25	0.294
3	M	3	60	20	45	0.529
4	Q	4	160	40	85	1.000

Table 3.4: Inverse Cumulative Example with $k = 10$ samples

$$p_{j,k} = \frac{\sum_{i=1}^j b_{i,k}}{\sum_{i=1}^n b_{i,k}} \quad (3.2)$$

Using the same data from Table 3.1, the sell-up probabilities calculated using IC are shown in Table 3.4. As in DO, when a new sample is added, $b_{j,k}$ and consequently $p_{j,k}$ are recalculated for the new value of k .

3.2.3 Forecast Prediction

Up to this point, calculating $b_{n,k}$ requires actual bookings in n to exist among the k samples. However, there may be instances when few if not none of the k samples have bookings in n . This results in the lack of information about the number of passengers who are willing to pay for n , thereby affecting the calculation of $p_{j,k}$.

Forecast Prediction (FP) addresses this potential issue by making use of the sell-up rates calculated in the previous iteration when there were $k - 1$ samples, $p_{j,k-1} \forall j$, to scale the current **total** bookings in each j , $d_{j,k} \forall j$, to an **equivalent** value for n . If there does not exist a previous iteration, $p_{j,k-1}$ are replaced with arbitrarily-chosen input sell-up probabilities. The sum of all the scaled values is then divided by the number of current samples k to provide the average forecasted booking values for n , $\hat{b}_{n,k}$. Mathematically, this is given by:

$$\hat{b}_{n,k} = \frac{\sum_{i=1}^n (d_{i,k})(p_{i,k-1})}{k}$$

Similar to DO, FP views each $b_{j,k}$ as an independent observation of the average bookings received by the airline in j during interval t . The probability $p_{j,k}$ of selling up a booking from n to j is thus given by:

Lowest Open Fare Class j	Fare Class Name	Number of Samples in j	$p_{j,k-1}$	Total $d_{j,k}$	Bookings		$p_{j,k}$
					$(d_{j,k})/(p_{j,k-1})$	Average $b_{j,k}$	
1	Y	1	0.200	10	50	10	0.260
2	B	2	0.400	30	75	15	0.390
3	M	3	0.600	60	100	20	0.519
4	Q	4	1.000	160	160	40	1.000
				Sum	385		
					$\frac{\sum_{j=1}^4 (d_{j,k})(p_{j,k-1})}{k}$	38.5	

Table 3.5: Forecast Prediction Example with $k = 10$ samples (Iteration 1)

Lowest Open Fare Class j	Fare Class Name	Number of Samples in j	$p_{j,k-1}$	Total $d_{j,k}$	Bookings		$p_{j,k}$
					$(d_{j,k})/(p_{j,k-1})$	Average $b_{j,k}$	
1	Y	1	0.260	10	38.5	10	0.261
2	B	2	0.390	30	77	15	0.392
3	M	3	0.519	60	115.5	20	0.523
4	Q	5	1.000	190	190	38	1.000
				Sum	421		
					$\frac{\sum_{j=1}^4 (d_{j,k})(p_{j,k-1})}{k}$	38.3	

Table 3.6: Forecast Prediction Example with $k = 11$ samples (Iteration 2)

$$p_{j,k} = \begin{cases} \frac{b_{j,k}}{b_{n,k}} & j \neq n \\ 1 & j = n \end{cases} \quad (3.3)$$

Using the same data from Table 3.1, the sell-up probabilities calculated using FP are shown in Table 3.5. Since this is the first iteration, input values for $p_{j,k-1}$ are needed. A second iteration (when $k = 11$) is continued in Table 3.6 with an additional sample of 30 bookings in Q. This causes $d_{4,11}$ to increase to 190, thereby changing the computed sell-up probabilities.

As will be seen in the simulation results found in the next chapter, forecast prediction generally results in the highest revenues among the 3 methods estimation methods.

3.3 Sell-up Estimation in PODS

Up to now, our discussion on the 3 conceptual sell-up estimation models has yet to incorporate fare values in the calculations. As described by Hopperstad [23], PODS incorporates fare values in estimating sell-up by considering fare ratios between Q-class and the higher-priced fare classes. Since this fare-ratio sell-up model in PODS contains a parameter (known as the sell-up constant) whose value is intended to be representative of the historical booking data, data-fitting is performed on the model so as to calculate the parameter value that most closely matches the historical data.

3.3.1 Fare-Ratio Sell-up Model

As before, sell-up is considered as occurring from class n (i.e. Q) to the next lowest available class j . However, for a given j , the number of combinations involving the fare value of n and that of j , f_n and f_j respectively, remains large since there are no restrictions on the values of f_n and f_j . To facilitate comparisons involving f_n and f_j , the fare ratio of j to n , $f_{j,n}$, is used instead. Mathematically, this is given to be

$$f_{j,n} = \frac{f_j}{f_n}, f_{j,n} \geq 1$$

To further facilitate our discussion, we also consider $f_{j,n,t}^{50\%}$ to be the fare ratio that causes 50% of the demand in n to sell-up to j during interval t . $f_{j,n,t}^{50\%}$ is abbreviated as **FRAT5** in PODS, and serves as a measure of the WTP of passengers during t . A larger $f_{j,n,t}^{50\%}$ value implies that passengers are willing to pay more to sell up from n to j during t .

From past observations of PODS simulations (as noted by Reyes [30]), the probability $p_{j,t}^f$ of selling up a booking from n to j during t (the superscript f indicates the fare ratio model) is modeled as having the following negative exponential relationship with $f_{j,n}$:

$$p_{j,t}^f = e^{-\left(\frac{\ln(0.5)}{f_{j,n,t}^{50\%} - 1}\right)(f_{j,n} - 1)} \quad (3.4)$$

The first expression in the exponent of Equation 3.4:

$$-\frac{\ln(0.5)}{f_{j,n,t}^{50\%} - 1}$$

Time Frame	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
From Days before Departure	63	56	49	42	35	31	28	24	21	17	14	10	7	5	3	1
To Days before Departure	57	50	43	36	32	29	25	22	18	15	11	8	6	4	2	0
Days in Time Frame	7	7	7	7	4	3	4	3	4	3	4	3	2	2	2	1

Table 3.7: Time Frames in PODS

is referred to as the **sell-up constant**, $c_t^{50\%}$, and is the model parameter whose value is fitted with the data.

Additionally, observe that $f_{j,n} = f_{j,n,t}^{50\%}$ when $p_{j,t}^f$ is set to be 0.5 in Equation 3.4. This intuitively makes sense because as already stated above, $f_{j,n,t}^{50\%}$ to be the fare ratio that causes 50% of the demand in n to sell-up to j during interval t .

While PODS has the ability to accept user-defined values of $p_{j,t}$, **sell-up estimation under the fare-based model involves using $f_{j,n}$ and $f_{j,n,t}^{50\%}$ as inputs to calculate $p_{j,t}^f$** . $f_{j,n}$ is user-defined whereas $f_{j,n,t}^{50\%}$ is expressed in terms of $c_t^{50\%}$, which results in the following expression for $f_{j,n,t}^{50\%}$:

$$f_{j,n,t}^{50\%} = \frac{-\ln(0.5)}{c_t^{50\%}} + 1$$

As will be seen later in Section 3.3.3, the fare-ratio sell-up model **uses historical booking data to estimate $c_t^{50\%}$ so as to calculate $f_{j,n,t}^{50\%}$ and consequently $p_{j,t}^f$** .

3.3.2 Booking Period in PODS

In PODS, the booking period consists of 16 consecutive intervals or timeframes. As shown in Table 3.7, these timeframes are defined to begin 63 days before the departure day of the flight and end on the departure day. Just like the assumptions made in the earlier sell-up model discussed in Section 3.1, passenger-related events such as bookings occur anytime within a timeframe whereas inventory control-related events like fare class closures usually occur only at the immediate beginning of the timeframe.

Lowest Open Fare Class j	Fare Class Name	$f_{j,4}$	Number of Samples in j	Bookings		
				Total	Average $b_{j,k,t}$	$p_{j,k,t}^d$
1	Y	4.00	1	10	10	0.250
2	B	3.00	2	30	15	0.375
3	M	2.00	3	60	20	0.500
4	Q	1.00	4	160	40	1.000

Table 3.8: Direct Observation Example with $f_{j,n}$ added for $n = 4$ at interval t

3.3.3 Estimation Using The Fare-Ratio Model

To estimate sell-up using the fare-ratio model, fare ratios $f_{j,n} \forall j$ are first defined for the fare classes and added to the booking data from the historical database. One of the three sell-up estimation methods is also selected.

Table 3.8 contains an example at interval t where $f_{j,n}$ values are added to the original data from Table 3.2. Additionally, the appropriate subscript is added to the average bookings and the sell-up probabilities in order to keep track of the interval t . The superscript d is also added to the sell-up probabilities to indicate that they were calculated using from the booking data.

The next stage involves estimating the sell-up constant, $c_t^{50\%}$ using linear least squares minimization. Since the value of $c_t^{50\%}$ is supposed to be indicative of the historical data, data-fitting is performed on it so as to obtain the value of $c_t^{50\%}$ that most closely represents the booking data. **It thus follows that the rationale for performing least squares minimization is that the value of $c_t^{50\%}$ that minimizes the square of the sum of difference between the model and the data is the value that best fits the historical data.** To make full use of the booking data across the 16 timeframes, PODS conducts least squares minimization twice: (i) within each timeframe and (ii) across all 16 timeframes. This two-step process is as follows:

1. Within each t , solve for $c_t^{50\%}$ in the least squares minimization of

$$\sum_{i=1}^n (p_{i,k,t}^d - p_{i,t}^f)^2$$

As an illustration, the data from Table 3.8 is plotted in Figure 3-1; the points correspond to $p_{i,k,t}^d$ and the fitted exponential line corresponds to $p_{i,t}^f$. As already men-




Figure 3-1: Step 1: Estimation of $c_t^{50\%}$ at interval t

tioned in Section 3.3.1, $f_{j,n} = f_{j,n,t}^{50\%}$ when $p_{j,t}^f$ is set to be 0.5 in Equation 3.4. The corresponding value of $c_t^{50\%}$ is thus computed to be 0.495 from $f_{j,n,t}^{50\%} = 2.4$.

2. Once Step 1 has been done for all 16 intervals, solve for the 2 constants, α and β , in the linear least squares minimization of

$$\sum_{i=1}^{16} \left((\alpha + \beta \cdot i) - c_t^{50\%} \right)^2$$

As an example, suppose we have already completed step 1 for some interval t . We then proceed to repeat step 1 for the 15 remaining intervals and obtain the corresponding values of $c_t^{50\%}$. The result of the linear least squares minimization will thus be similar to that shown in Figure 3-2.

The final estimated value of $c_t^{50\%}$ in t is thus $(\alpha + \beta \cdot t)$. For each of the 16 values of t , corresponding values of $f_{j,n,t}^{50\%}$ and $p_{j,t}^f$ are then calculated. These values of $f_{j,n,t}^{50\%}$ all the 16 timeframes thus give rise to the FRAT5 curve as shown in Figure 3-3. Note that $f_{j,n,t}^{50\%}$ values closer to the departure date are typically larger than those further away (i.e. in the earlier time-frames). This arises from the tendency of business travelers to book later in the booking period than leisure passengers.




Figure 3-2: Step 2: Estimation of $c_t^{50\%} = \alpha + \beta.t$ across all 16 intervals to obtain α and β

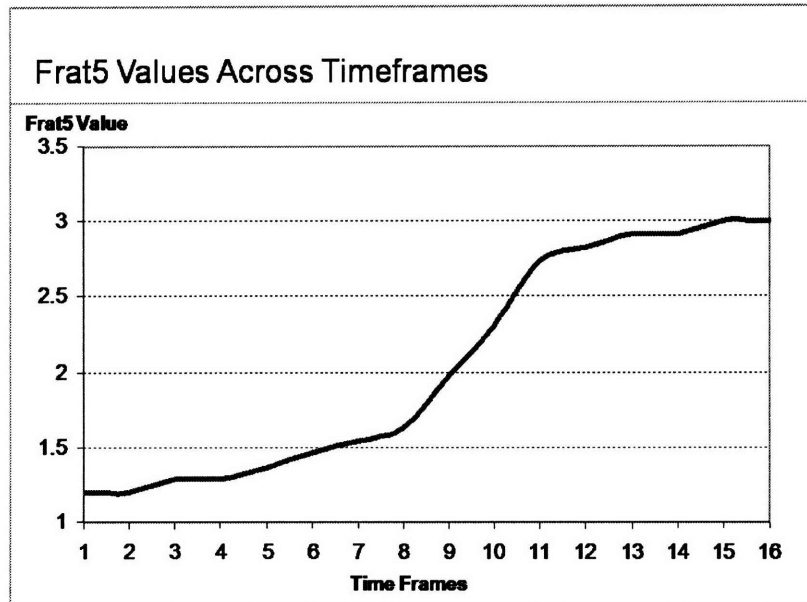


Figure 3-3: Example - Plot of $f_{j,n,t}^{50\%}$ over time

3.4 Chapter Summary

This chapter has described in detail three methods of estimating sell-up; it has illustrated how each method estimates sell-up using historical booking data. It has also described the fare-based model that PODS uses in conjunction with any of these three methods to estimate sell-up. The next chapter contains a further explanation of the PODS environment, together with the simulation description and results.

Chapter 4

PODS Simulation Environment and Simulation Results

The main aim of this chapter is to examine and compare the effects that each of the 3 sell-up estimation methods has on an airline's revenues. As described by Gorin and Belobaba [21], such analysis is best done in a simulation environment where both (i) the competitive dynamics between airlines as well as (ii) the interactions between passenger booking decisions and RM systems can be modeled. PODS thus serves as an appropriate tool to evaluate the sell-up estimation methods. Consequently, the first half of this chapter provides a brief overview of the PODS simulation environment. The second half of this chapter discusses the particular network setup used for the simulation before explaining the simulation results.

4.1 PODS General Architecture

As already alluded to in the introduction to this chapter, the PODS environment simulates the interactions of two groups - airlines and passengers - leading up to flight departures for a single day (represented by a sample). The airline side is represented by one or more third-generation RM systems similar to those described in Section 2.1. The passenger side that provides booking demand is based on an improved version of the Decision Window Model (DWM) for passenger choice, which was originally developed at Boeing. This enhanced version of the DWM models passenger preferences based on fare values and fare restrictions, in addition to factors such as flight schedules already present in the original version. A

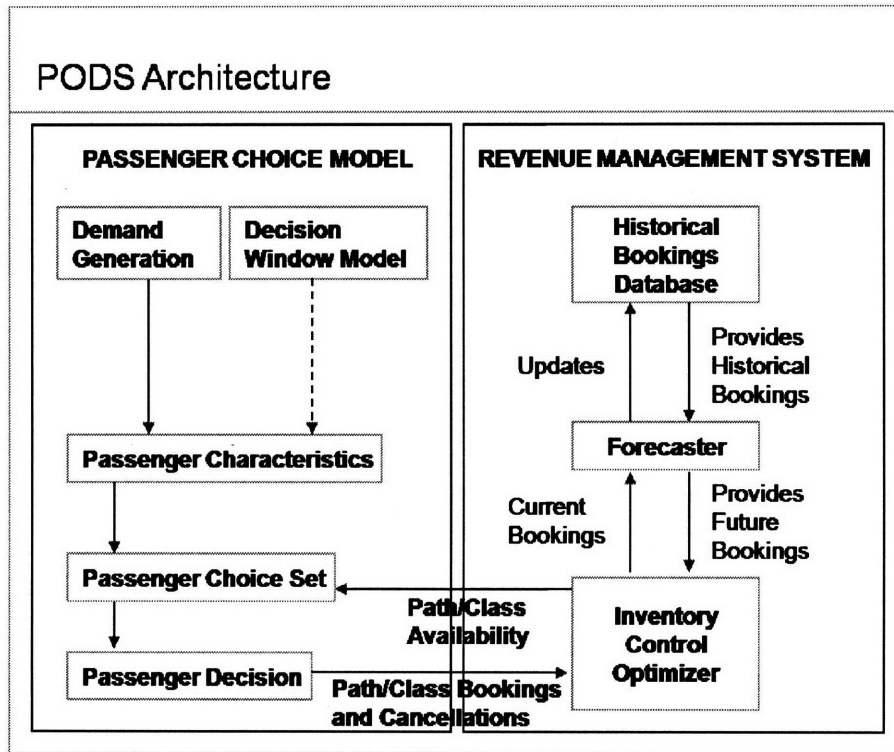


Figure 4-1: PODS Structure

schematic representation of the inter-relationship between passengers and airlines is shown in Figure 4-1.

As has already been described in greater detail in Section 3.3.2, bookings for each sample (a single departure day) arrive over 16 consecutive timeframes in PODS. A typical simulation run in PODS consists of 5 trials, each comprising of 600 samples. Since user-defined inputs are used to initialize the first sample in the first trial, the first 200 samples in every trial are discarded to reduce any biasing effects of the input values. The simulation results are thus obtained by averaging the remaining 2000 samples (400 from each of the 5 trials).

4.1.1 Passenger Choice Model

Given a set of fare products with associated fare values and restrictions, the resultant demand depends on the passengers' preferences or behavior towards these fare products. In PODS, such preferences are taken into account by the Passenger Choice Model, which ultimately determines whether or not a particular fare product will be accepted by a passenger.

The outcome of such an event in the model involves a series of four sequential steps that are outlined below: (i) Demand Generation, (ii) Passenger Characteristics Assignment, (iii) Passenger Choice Set Definition, and (iv) Passenger Decision. These four steps are briefly described below; a complete description and justification of the Passenger Choice Model can be found in Carrier [29].

Demand Generation

The Passenger Choice Model first generates an average value for air travel demand in terms of passenger numbers for each of the O-D markets declared in the user-defined network at the base fare (Q-class). Variability is then introduced by randomly generating passenger numbers around this average value to provide a range of passenger demand values. To model scenarios with varying passenger demand, PODS also accepts as input a Demand Multiplier (DM) value. Typical DM values under low and moderate passenger demand scenarios are 0.8 and 1.0 respectively.

Passenger Characteristics Assignment

Three characteristics are assigned to each passenger. Firstly, a decision window that defines the earliest acceptable departure time and latest acceptable arrival time is generated for every passenger so as to represent his or her time preferences. The second characteristic assigned to each passenger is a maximum WTP value. The third characteristic assigned is a value for the disutility costs, which quantifies the amount of disutility the restrictions (if present) exert on the passenger.

Passenger Choice Set Definition

Based on the passenger's WTP and Decision Window and disutility values, the applicable fare products are then selected from the list of available fare products provided by the RM system. A booking will not be made if there are no applicable or available fare products.

Passenger Decision

For each of the applicable fare products, the passenger sums the associated fare value and disutility costs so as to choose the path/fare option with the lowest total value. This booking

decision regarding the chosen path/fare option is then recorded by the historical database in the revenue management system.

4.1.2 PODS Revenue Management System

As already mentioned in Section 2.1, a typical 3rd-generation RM system consists of three components: (i) Historical Booking Database, (ii) Forecaster, and (iii) Seat Allocation Optimizer.

Historical Booking Database

Each airline has a historical booking database that stores booking data by fare class and path. (At the start of each trial, it is initialized with default values.) For each flight in PODS, booking data for the previous 26 departures is accumulated in the database.

Forecaster

As previously described in Section 2.1.1, the forecaster uses pick-up forecasting on the historical booking data to estimate future demand for each fare class. Should the airline decide to account for sell-up when forecasting future bookings in an unrestricted fare environment, then the Q-forecasting method (as described in Section 2.3.1) is also used in conjunction with pick-up forecasting.

Seat Allocation Optimizer

There exist several seat allocation optimizers in PODS that exhibit varying levels of sophistication in assigning a fixed inventory of seats to different fare classes. For this thesis simulation, three optimizers are used: (i) Adaptive Threshold (AT90), (ii) Expected Marginal Seat Revenue b (EMSRb) and (iii) Displacement Adjusted Virtual Nesting (DAVN). **Adaptive Threshold does not consider sell-up, whereas the other two methods for the simulations in this thesis are set to account for sell-up** in optimizing seat allocation among the different fare classes.

4.1.3 PODS Seat Allocation Optimizer Methods

Other than the Adaptive Threshold method, the other two methods have already been described in Section 2.1.2 and Section 2.1.2 respectively.

Fare Class j	Initial Load Factor Threshold (%)
1	100
2	90
3	70
4	60
5	50
6	40

Table 4.1: Initial Load Factor Thresholds for AT90

Adaptive Threshold

The Adaptive Threshold algorithm is a leg-based allocation method that progressively closes down fare classes on a leg as they approach a target load factor. For the simulation, this method is abbreviated as AT90 since the target load factor is set at 90%. This is done by assigning the baseline load factor thresholds for each of the fare classes at the beginning of the simulation as shown in Table 4.1. More information on AT90 can be found in Cleaz-Savoyen [15].

EMSRb

Please refer to Section 2.1.2.

DAVN

Please refer to Section 2.1.2.

4.2 Simulation Environment

Since sell-up is based on price-oriented behavior, a simulation environment with an unrestricted fare structure, known as Network D6 (unrestricted), is used for this thesis. A schematic representation of this network can be found in Figure 4-2.

Network D6 (unrestricted) is a simplified representation of the US domestic airline network consisting of two hub-and-spoke airlines: AL1 and AL2 with hubs $H1$ (Minneapolis-Saint Paul International Airport) and $H2$ (Dallas-Fort Worth International Airport) respectively. Each airline flies from its own hub to twenty-one other cities: twenty western cities (numbered 1 to 20) and its competitor's hub. Each airline also flies to its own hub




Figure 4-2: Network D6 (unrestricted) Network

from twenty-one other cities: twenty eastern cities (numbered 21 to 40) and its competitor's hub. This results in a total of 252 legs in Network D6 (unrestricted) all operated by aircraft with identical capacity.

Such a network layout provides both local west-hub and hub-east service, as well as connecting west-east service with three connecting banks of flights. As such, there are a total of 482 OD markets in each bank, 42 of which are local markets and the remaining 440 being connecting markets.

The fare structure and (lack of) restrictions for the 6 fare classes in Network D6 (un-

Fare Class j	Average Fare (\$)	Restrictions			
		Advance Purchase	R1	R2	R3
1	412.85	None	None	None	None
2	293.34	None	None	None	None
3	179.01	None	None	None	None
4	153.03	None	None	None	None
5	127.05	None	None	None	None
6	101.06	None	None	None	None

Table 4.2: Fare Structure and Restrictions in Network D6 (unrestricted)

restricted) are displayed in Table 4.2. R1 refers to the Saturday night stay restriction, R2 refers to the itinerary change fee, and R3 refers to the non-refundability of the ticket.

4.3 Simulation Setup

The main motivation for performing the simulations is to compare the revenue performance of an airline (AL1) under different seat allocation optimization scenarios when it uses the 3 sell-up estimation methods previously discussed in Chapter 3: **(i) Direct Observation (DO)**, **(ii) Forecast Prediction (FP)** and **(iii) Inverse Cumulative**. The performance of these 3 sell-up estimation methods thus needs to be bench-marked against some other method available in PODS that calculates the sell-up probabilities. As earlier mentioned in Section 3.3.1, PODS has the ability to accept user-defined values of the sell-up probabilities $p_{j,t}$ by having the user specify a particular FRAT5 ($f_{j,n,t}^{50\%}$) value for each of the 16 time frames. Among the various sets of input FRAT5 sell-up values available in PODS, using a particular set known as FRAT5C has typically resulted in the greatest revenues for the practicing airline. Even though no airline in the real world has access to a set of sell-up values that would maximize its revenues, **an appropriate reference for comparing the revenues of the 3 estimation methods in AL1 would still be for AL1 to use FRAT5C sell-up values** since it allow the comparison of how well the 3 sell-up estimation methods perform relative to the best-performing input set of FRAT5 values. Table 4.3 contains the particular $f_{j,n,t}^{50\%}$ value associated with the FRAT5C set for each of the 16 time frames. The corresponding plot for FRAT5C is also depicted in Figure 4-3.

Since Network D6 (Unrestricted) is a 2-airline network, a decision also needs to be made regarding the sell-up method (i.e estimated or input values) when a suitable seat allocation optimizer is used (EMSRb or DAVN) i.e. whether it uses estimated or input values. Since the primary objective of the simulations is to observe the revenue performance of the 3 sell-up estimation methods used by AL1, the sell-up method practiced by AL2 should (i) rely on input values and not on estimation and (ii) be kept constant throughout the simulations. Again, **having AL2 use input FRAT5C values for sell-up** is an appropriate reference setup since it also facilitates the comparison of how the 3 sell-up estimation methods (used by AL1) fare against a competitor (AL2) that uses the best-performing input set of FRAT5 values.

Time Frame	FRAT5 Value
1	1.20
2	1.20
3	1.29
4	1.29
5	1.37
6	1.46
7	1.54
8	1.63
9	1.97
10	2.31
11	2.74
12	2.83
13	2.91
14	2.91
15	3.00
16	3.00

Table 4.3: FRAT5C Input Values




Figure 4-3: Plot of FRAT5C Input Values

Test Case	AL1	AL2
0	AT90	AT90
1	EMSRb (FRAT5C, DO, FP, IC)	AT90
2	EMSRb (FRAT5C, DO, FP, IC)	EMSRb (FRAT5C)
3	DAVN (FRAT5C, DO, FP, IC)	AT90
4	DAVN (FRAT5C, DO, FP, IC)	EMSRb (FRAT5C)
5	DAVN (FRAT5C, DO, FP, IC)	DAVN (FRAT5C)

Table 4.4: Simulation Test Cases (DM values of both 0.8 and 1.0)

Table 4.4 lists 5 main simulation test cases (numbered 1 to 5) that represent a variety of possible leg-based/OD-based optimization combinations existing in airline markets. To model low and moderate passenger demand scenarios, DM values of both 0.8 and 1.0 are used for each test case.

As already described in Section 2.3.1, compelling passengers to sell-up (using Q-forecasting) in unrestricted fare environments results in greater revenues for airlines. Since there exist 2 test cases (1 and 3) where AL2 uses a seat allocation optimizer (AT90) that does not consider sell-up, there needs to be a reference test case (known as Test Case 0) where both airlines do not practice sell-up. **This allows the revenues of AL1 achieved by estimating sell-up (in Test Cases 1 and 3) to be compared against that in the absence of sell-up (in Test Case 0).**

4.4 Simulation Results

In each of the 5 main test cases (numbered 1 to 5), the revenues obtained by each of the 3 sell-up estimation methods by AL1 are compared against those achieved using input FRAT5C sell-up values by AL1. The other metrics that are closely-related to revenues are Load Factor (expressed as a percentage) and Yields (measured in dollars per Revenue Passenger Mile (RPM)). The Load Factor value is an indication of how full the airline's flights are across its entire network; it is calculated by finding the ratio of total RPMs to total Available Seat Miles (ASMs). Yields in turn measure the average fare paid by passengers, per mile flown. In general, there exists an inverse relationship between Load Factor values and Yields since cheaper seats attract more passengers and vice-versa.

Additional metrics that are relevant to this simulation are the Fare Class Load (expressed as a percentage), FRAT5 curves, and the Sell-Up Probability curves. The Fare

D.M.	Revenue - AL1
0.8	854808
1.0	1110128

Table 4.5: Revenue values of AL1 in Test Case 0

Class Load indicate the distribution of passengers among the 6 fare classes found in Network D6 (unrestricted); higher yields are typically due to a greater percentage of passengers having bought tickets in the higher-priced fare classes. As described in much greater detail in Section 3.3.3, the FRAT5 curves serves as a measure of the WTP of passengers since they indicate the fare ratio that result in 50% of the passengers selling up from Q-class (fare class 6 in the simulations) to a higher fare class. The Sell-Up Probability curves provide even more insight into the WTP of passengers since they track the values of $p_{j,t}^f$, which as described in Section 3.3.1, are the probability values of a passenger selling up from Q-class to class j during time frame t (where $j =$ fare class 1, ... , 5 in the simulations).

4.4.1 Test Case 0 - AT90 (AL1) vs. AT90 (AL2)

As stated in Section 4.3, the primary purpose of this test case is to serve as a reference for Test Cases 1 and 3. As such, only the revenue values of AL1 are utilized for comparison with Test Cases 1 and 3. These revenue values are shown in Table 4.5.

4.4.2 Test Case 1 - EMSRb (AL1) vs. AT90 (AL2)

Unlike AL2, AL1 considers sell-up in optimizing its seat allocation among the different fare classes. As expected, AL1's revenues are thus on average about 10% to 15% greater than those of AL2. This result applies for both DM values, as shown in Figures 4-4 and 4-5.

When compared with Test Case 0, the consideration of sell-up using either estimation or input FRAT5C values generally leads to a revenue increase of approximately 2% for both DM values. The only exception to this is when AL1 uses DO for $DM = 0.8$. As shown by the lower yields in Figure 4-6, this is due to DO causing AL1 to carry a greater proportion of its passengers in cheaper fare classes.

In comparing the performance of the 3 different sell-up estimation methods within AL1, we observe that FP results in slightly less revenues than that obtained using input FRAT5C when $DM = 0.8$. Closely examining the corresponding sell-up probabilities and FRAT5




Figure 4-4: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 1, DM = 0.8)




Figure 4-5: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 1, DM = 1.0)

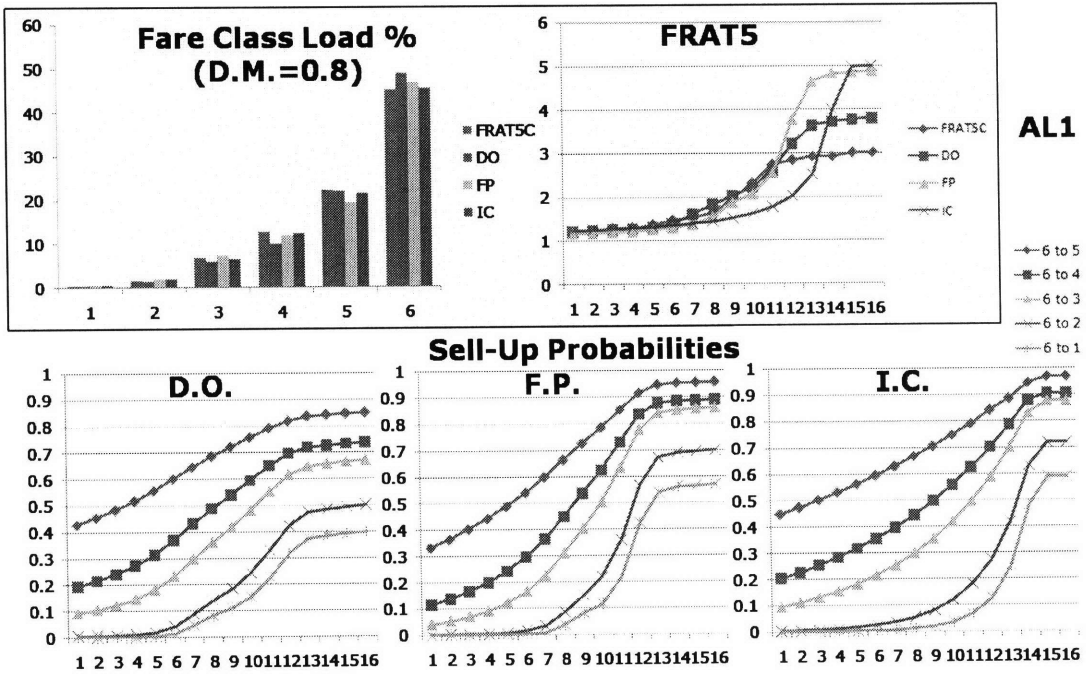


Figure 4-6: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 1, DM = 0.8)

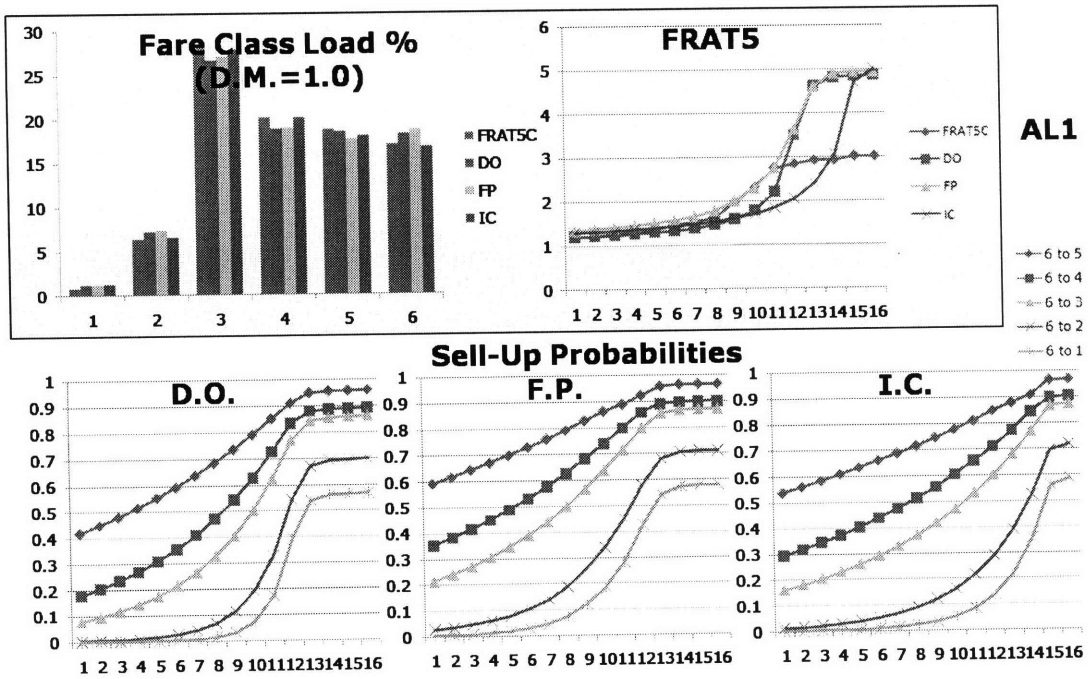


Figure 4-7: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 1, DM = 1.0)

curves in Figure 4-6 suggests that FP is over-aggressive relative to the demand environment in estimating sell-up. By closing down seat availability in the lower-priced fare classes more than necessary, AL1 forgoes revenue from passengers whose maximum WTP values fall within the ticket price of one of the closed fare classes. This claim is supported by comparing load factor and yield metrics with those of the FRAT5C reference scenario; FP has the greatest percentage decrease in load factor values (0.9%) and similarly the greatest percentage increase in yields (0.6%) among all 3 sell-up estimation methods. The load factor decreases due to the forgone passengers in the closed fare class products. At the same time, yields increase because the average fare paid by all carried passengers has risen.

When $DM = 1.0$, the greater passenger demand appears to result in more sell-up opportunities, thereby making FP the best performing method. Although the corresponding FRAT5 curve remains higher (which in turn results in more aggressive sell-up estimates) as shown in Figure 4-7, both load factor values and yields are greater than those for the $DM = 0.8$ case. In fact, when compared to the corresponding FRAT5C reference scenario, the percentage decrease in load factor values (0.5%) is smaller than the percentage increase in yields (1.1%). This allows AL1 to achieve greater revenues using FP than using input FRAT5C sell-up values.

4.4.3 Test Case 2 - EMSRb (AL1) vs. EMSRb (AL2)

Since both airlines now use an identical leg-based method (EMSRb) that considers sell-up in optimizing seat allocation among the different fare classes, any significant differences between the revenues of the airlines can be primarily attributed to the sell-up estimation method used by AL1 (since AL2 uses input FRAT5C sell-up).

The first test case illustrated the negative effects on revenue when over-estimating sell-up, when FP was over-aggressively closing down cheaper fare classes during $DM = 0.8$. This test case showcases the other extreme by illustrating the negative effects on revenue when under-estimating sell-up. Specifically we observe that DO now results in the lowest revenues among all 3 sell-up estimation methods when $DM = 0.8$ as shown in Figure 4-8. Closely examining the corresponding sell-up probability curves in Figure 4-9 indicates that the estimated sell-up probability values are too low relative to the demand environment, especially in the later timeframes closer to the departure date of the flight.

By making more seats than necessary available in the lower-priced fare classes, AL1

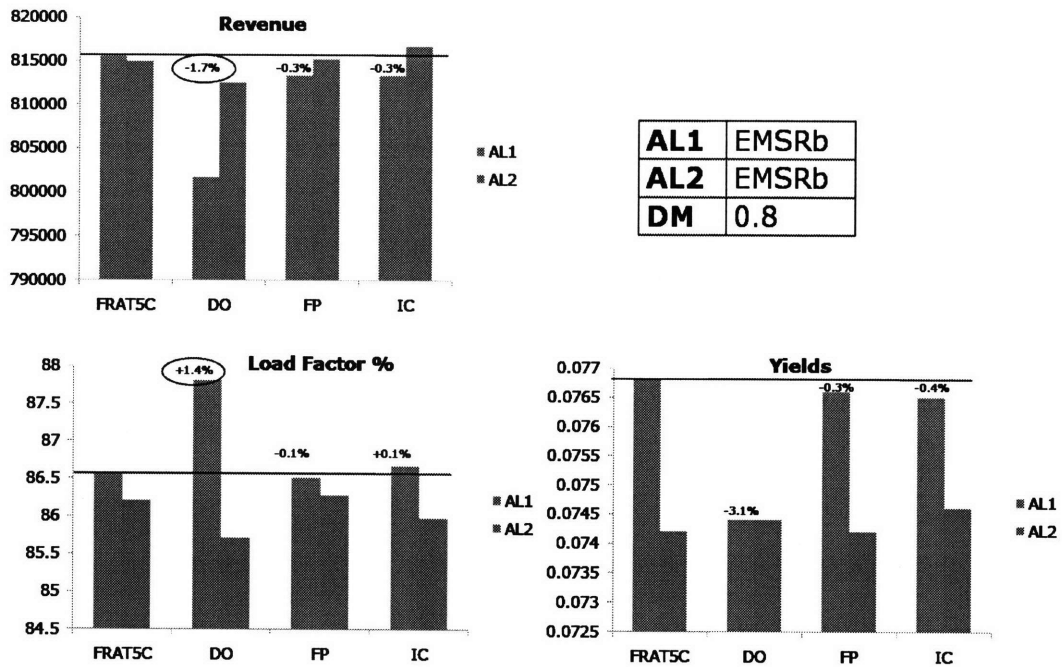


Figure 4-8: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 2, DM = 0.8)

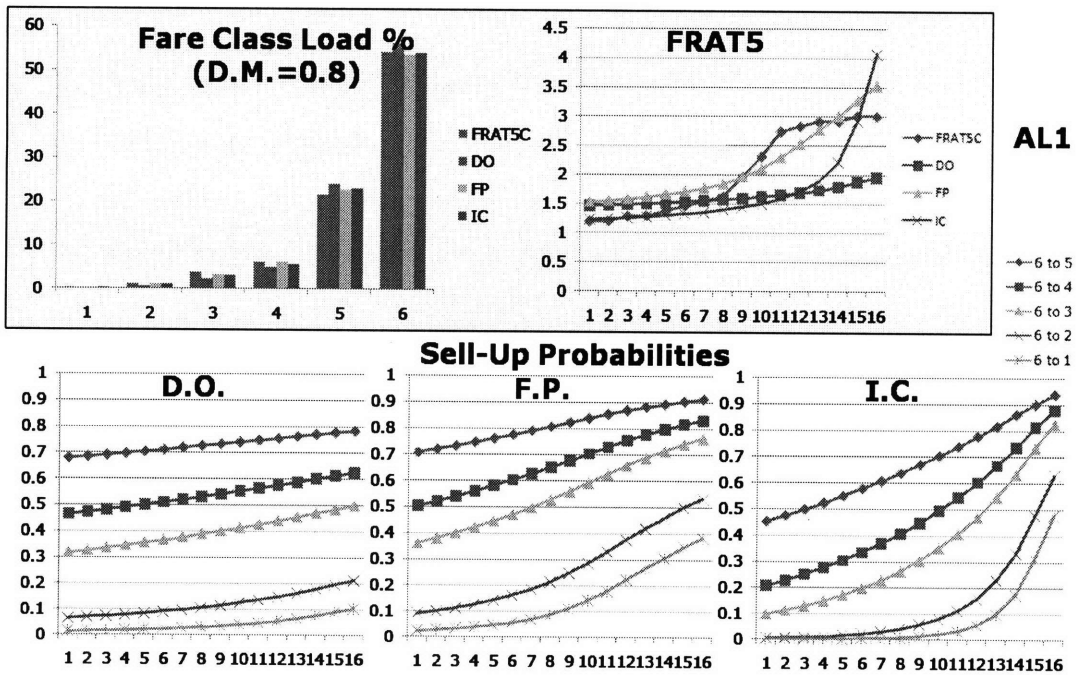


Figure 4-9: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 2, DM = 0.8)

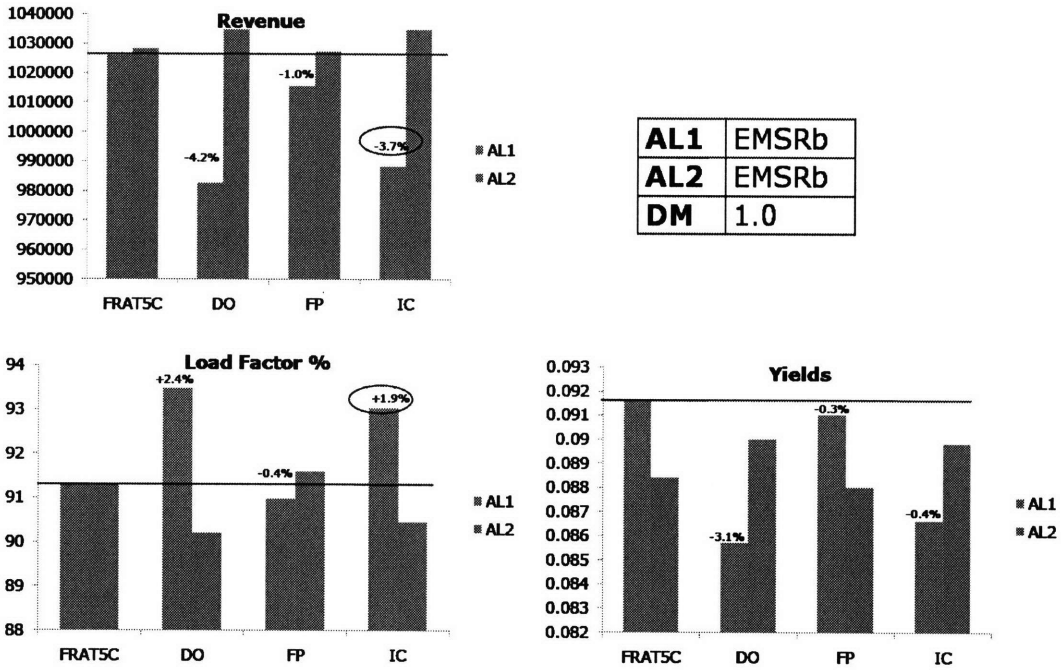


Figure 4-10: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 2, DM = 1.0)

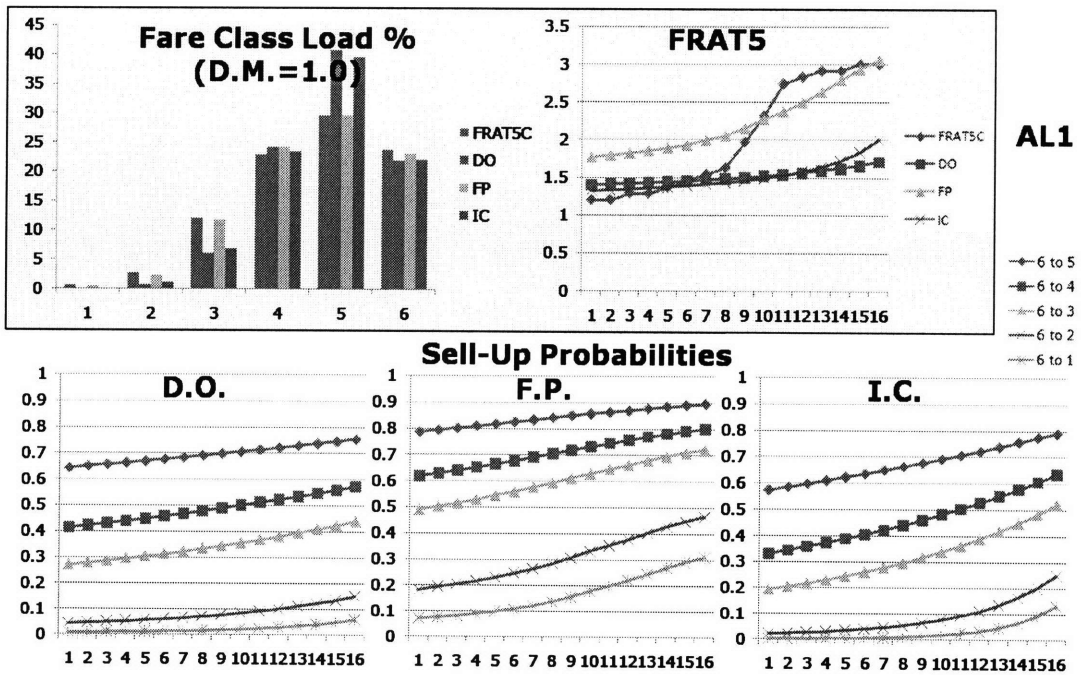


Figure 4-11: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 2, DM = 1.0)

forgoes revenue from passengers whose maximum WTP values are greater than the ticket price of the lowest open fare class. This claim is supported by comparing load factor and yield metrics with those of the other sell-up methods; DO has the greatest load factor values (about 88%) and similarly the lowest yields (about \$ 0.081 per passenger mile) among all 3 sell-up methods. The load factor increase since cheaper fare classes are able to attract more passengers whose maximum WTP values are at least equal to the ticket prices. At the same time, yields decrease because the average fare paid by all carried passengers has dropped.

When $DM = 1.0$, all 3 methods appear to exhibit an appropriate level of aggressiveness in estimating sell-up. From Figure 4-11, the sell-up probability curves of all 3 methods appear to be relatively flat and monotonically increase closer to the departure date. As shown in Figure 4-10, all 3 methods have lower load factors and higher yields than the FRAT5C reference scenario, resulting in revenues equivalent to or greater than that when AL1 uses input FRAT5C sell-up values.

At each time frame, FP has the greatest sell-up probability values among the sell-up methods, resulting in the greatest revenue increase for AL1.

4.4.4 Test Case 3 - DAVN (AL1) vs. AT90 (AL2)

Similar to Test Case 1, AL1 considers sell-up in optimizing its seat allocation among the different fare classes whereas AL2 does not. As expected, AL1's revenues are on average about 10% to 15% greater than those of AL2. This result holds true for both DM values across all sell-up methods and demand environments, as shown in Figures 4-12 and 4-13. As in Test Case 1, the revenue benefits of considering sell-up either using estimation or input FRAT5C values are also illustrated in the same figures; depending on the demand environment and estimation method used, revenues can increase by approximately up to 2%.

However, the main insight from this test case is that it is difficult to compensate for inappropriate sell-up estimates early on in booking period by trying to get as much revenue closer to the departure date, even though a greater proportion of bookings arrive in the later time frames. Specifically, spurring sell-up during the later booking timeframes may not make up for under-estimating sell-up in the earlier booking periods. There are two diametrically opposite causes for this: (i) the degree of sell-up in the later timeframes might




Figure 4-12: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 3, DM = 0.8)




Figure 4-13: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 3, DM = 1.0)

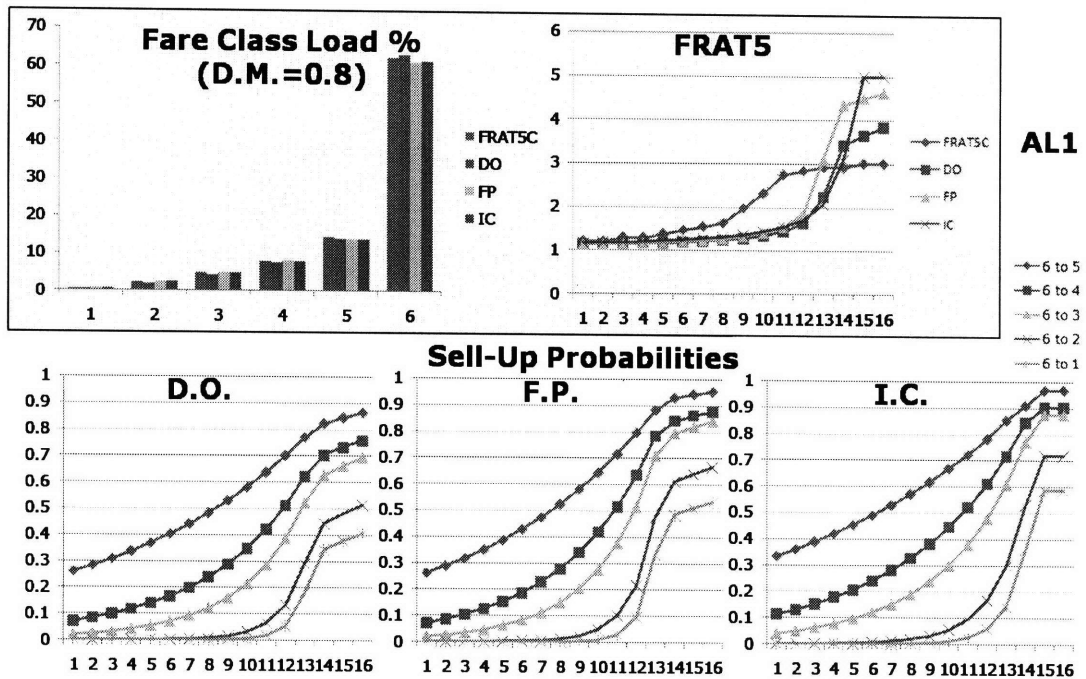


Figure 4-14: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 3, DM = 0.8)

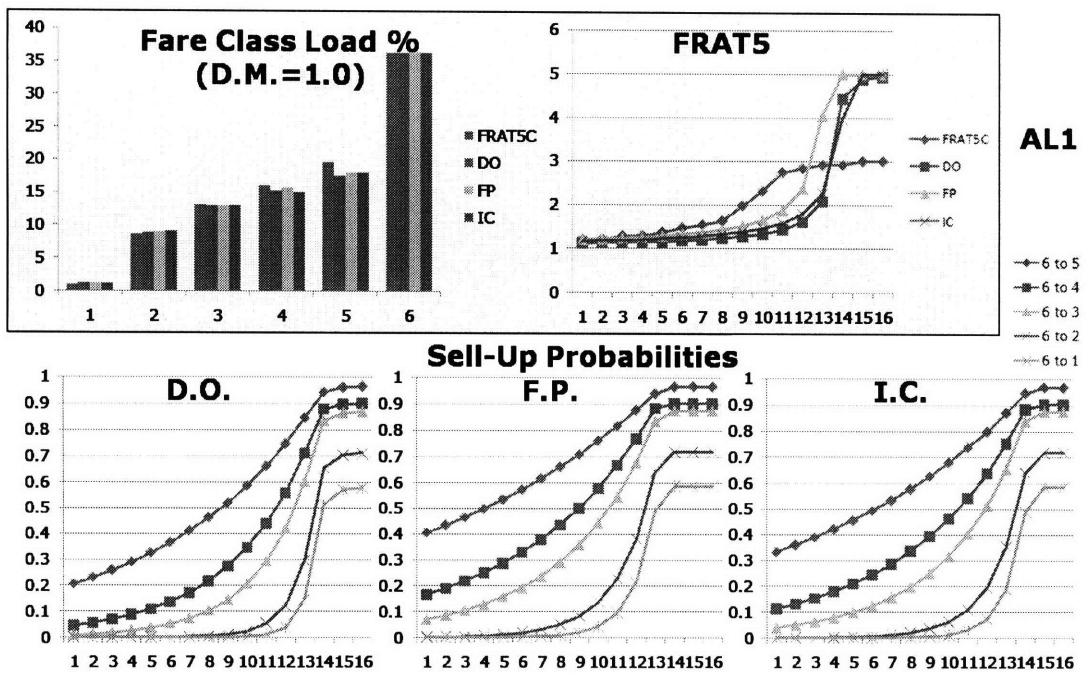


Figure 4-15: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 3, DM = 1.0)

still be underestimated (as shown by the less steep FRAT5 curve and lower-than-average yields of DO when $DM = 0.8$) and (ii) the degree of sell-up in the later timeframes might be overly-aggressive (as shown by all other 5 method/ demand-environment combinations) where the percentage increase in yields is less than the percentage increase in load factor.

As indicated in greater detail in Figures 4-14 and 4-15, all 3 methods start off by with their FRAT5 curves remaining significant flatter than that of the FRAT5C reference case; they only become noticeably steeper after timeframe 12. This increase in aggressiveness later on in the booking process is mainly attributed to the significant increase in the probabilities of selling up from fare class 6 (Q-class) to the most expensive fare classes such as 1 and 2.

As also mentioned at the beginning of this section, although the degree of sell-up realized closer to the departure date is sufficient to cause average yields to rise by about 1% over those of FRAT5C, the level of aggressiveness causes load factor values to decrease by a greater percentage value of up to 2%. This results in revenues for the 3 methods to all fall below that of the FRAT5C reference scenario.

Among these methods, using FP results in one of the highest revenues in both demand environments because it starts becoming aggressive earlier than DO or IC as shown by the relevant FRAT5 curves, thereby capturing additional sell-up revenue.

4.4.5 Test Case 4 - DAVN (AL1) vs. EMSRb (AL2)

Previous simulations by Cleaz-Savoyen [15] and Reyes [30] have shown that when airlines using network-OD optimization methods compete against airlines employing leg-based heuristics, the former category typically earns greater if not equivalent revenues. This observation is corroborated by the revenues of AL1 and AL2 when both adopt the input FRAT5C sell-up method, as shown in Figures 4-16 and 4-17.

The main insight from this test case is thus to be mindful of the effects of having only 1 competitor (AL2) present in Network D6 on the revenue performance of the sell-up estimation methods, namely the seemingly counter-intuitive result that leg-based heuristics (used by AL2) result in more revenues than network-OD optimization methods using sell-up estimation (used by AL1). The reason for this phenomenon is that all 3 sell-up estimation methods in both demand environments underestimate sell-up relative to the corresponding FRAT5C reference scenario as shown by the higher load factor values and lower yields. As shown by the fare class load distribution, the 3 sell-up estimation methods consequently




Figure 4-16: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 4, DM = 0.8)




Figure 4-17: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 4, DM = 1.0)




Figure 4-18: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 4, DM = 0.8)

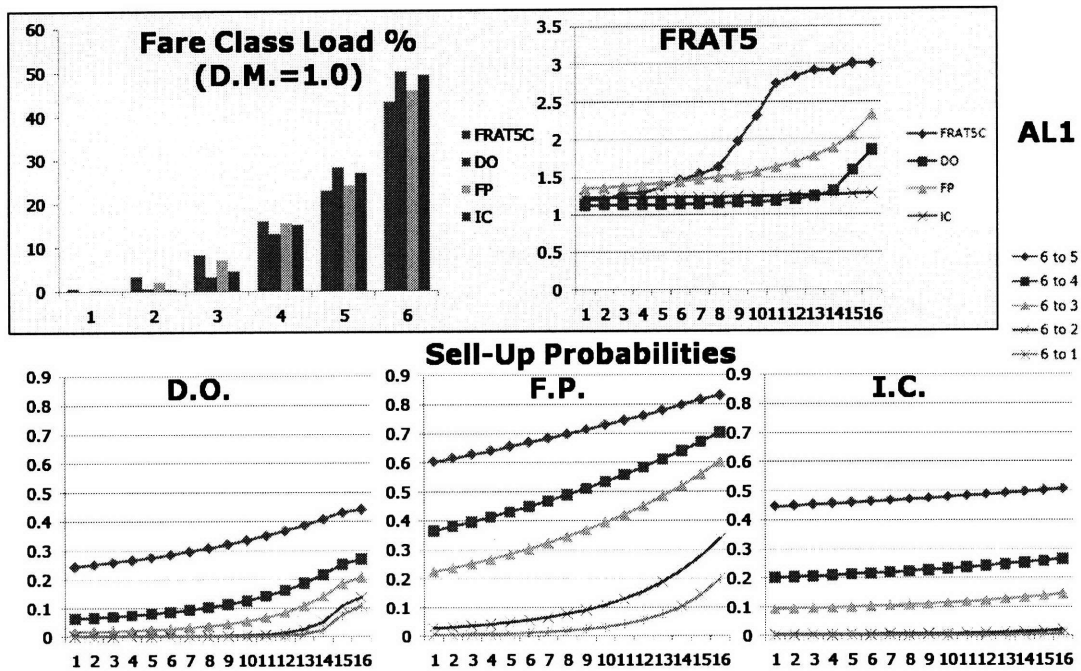


Figure 4-19: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 4, DM = 1.0)




Figure 4-20: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 5, DM = 0.8)

result in AL1 carrying more passengers in the cheaper-priced fare classes than it would have under FRAT5C. More crucially, as suggested in Figures 4-18 and 4-19 where the FRAT5 and sell-up probability curves remain relatively flat in the later time frames, having cheaper fare classes made available results in insufficient seats to accommodate passenger demand. Some of the higher-yield business passengers are unable to purchase seats with AL1 and thus spill-over to AL2. This causes both yields and revenues at AL2 to be greater than those at AL1 - a result unique to Network D6 since there is only other 1 competitor present.

As in Test Case 3, the relative performance of FP in terms of estimating sell-up (and the resultant revenues) remain superior to the other 2 methods throughout both demand environments.

4.4.6 Test Case 5 - DAVN (AL1) vs. DAVN (AL2)

The main insight from this test case is similar to that of Test Case 4, i.e. having only 1 other competitor present in the network leads to yields and revenue results for the competitor that may not necessarily be generalized across all possible airline networks.

As in the previous test case, all 3 sell-up estimation methods continue to underestimate

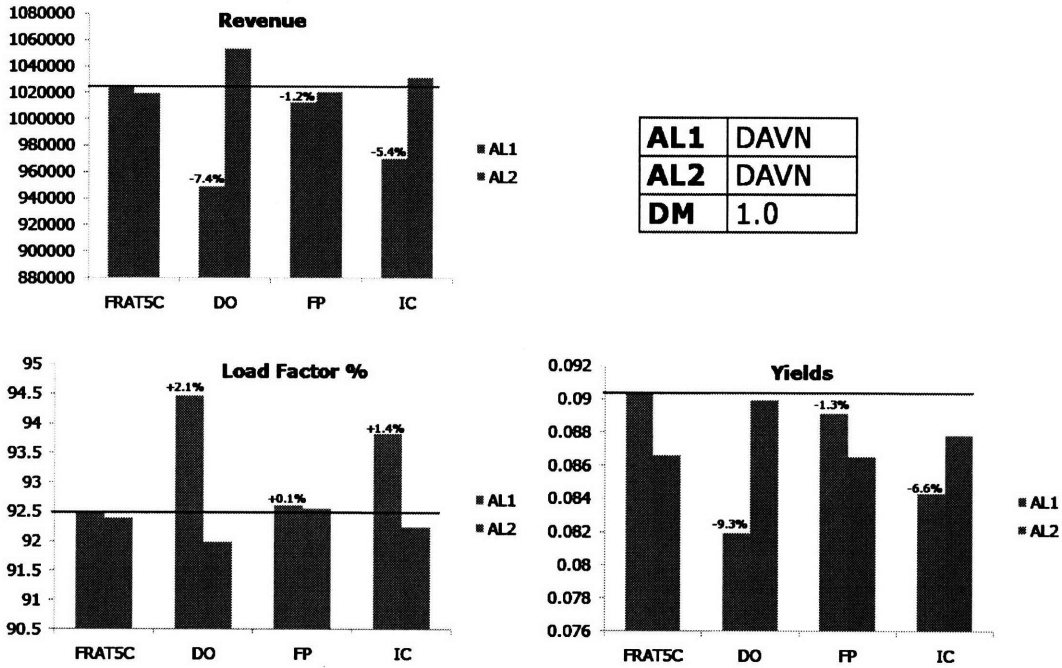


Figure 4-21: AL1 vs. AL2 - Revenues, Load Factor, Yields (Test Case 5, DM = 1.0)

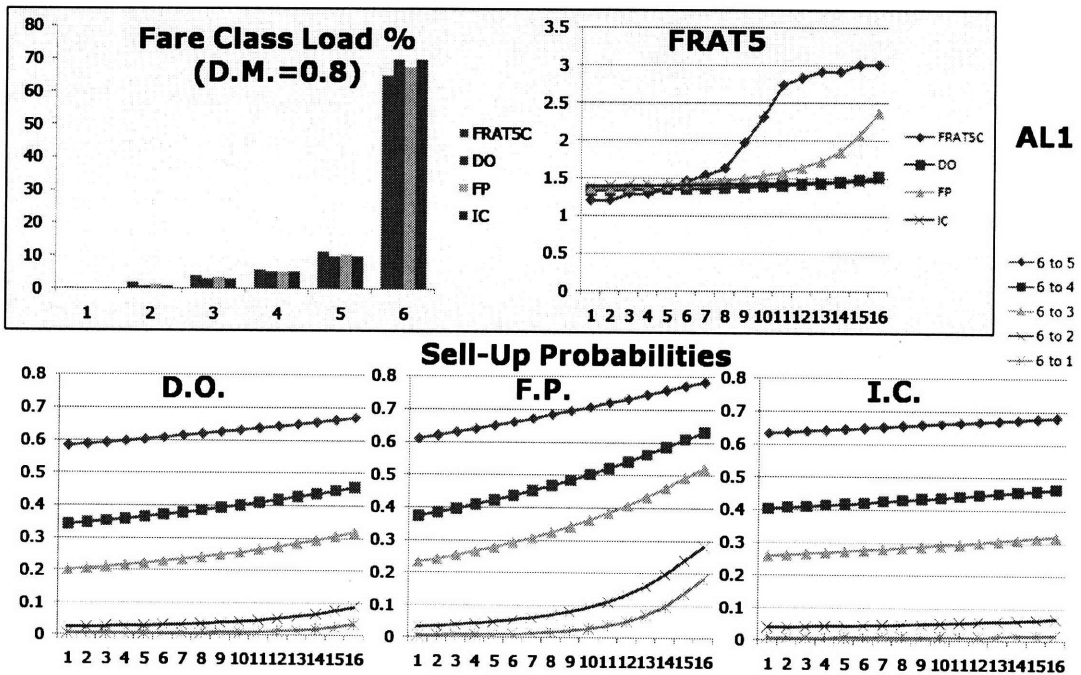


Figure 4-22: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 5, DM = 0.8)

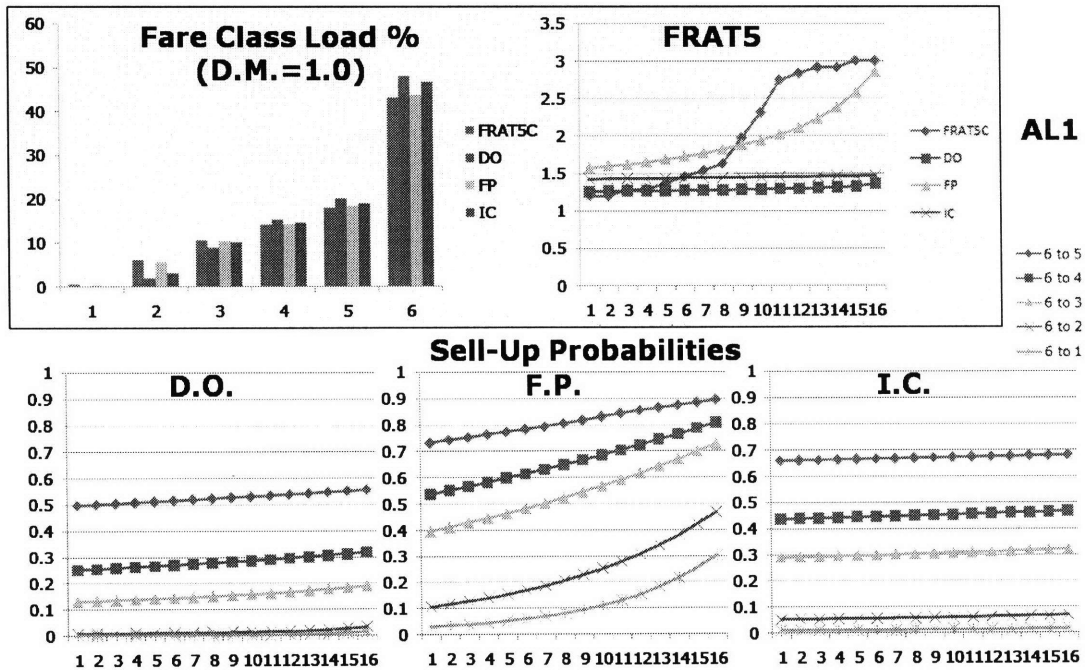


Figure 4-23: AL1 - Fare Class Load, FRAT5, Sell-up Probabilities (Test Case 5, DM = 1.0)

sell-up relative to the corresponding FRAT5C reference scenario as indicated by their higher load factor values and lower yields. As indicated in Figures 4-22 and 4-23, their FRAT5 and sell-up probability curves remain relatively flat throughout the entire booking period, causing AL1 to again spill some of the higher-yield business passengers over to AL2. The resultant revenue decreases for AL1 are such that it now earns up to 7% less revenue than AL2 (for DO when DM = 1.0), a result that is again unique to Network D6 since there is only 1 competitor present.

Among the 3 sell-up estimation methods, FP's ability to result in the greatest revenues arises from its moderate sell-up aggressiveness; its FRAT5 slopes upwards, especially in the later time frames.

4.5 Chapter Summary

This chapter has evaluated the 3 sell-up estimation methods by comparing their simulated effects on revenues. As can be inferred from the series of test cases, as AL2 becomes increasingly sophisticated in shifting from threshold to leg-based heuristics and finally network-OD optimization methods, the degree of sell-up estimated by AL1 falls. Intuitively, this makes

sense since AL2 would be better positioned to capture passengers from AL1. By managing its seat inventory better, AL2 has a better chance of offering the appropriate fare-class products to passengers; passengers thus have a wider set of applicable fare products from which to choose.

The main insight after conducting the series of test cases with regards to AL1 is that having an appropriate level of sell-up aggressiveness is crucial to increase revenues; under- and/or over-estimating sell-up relative to the actual WTP values of the passengers is often detrimental to the airline. **FP has resulted in the highest revenues among all 3 methods because it has been able to provide estimated sell-up levels that best represent the WTP values of passengers.**

The next chapter will discuss further the implications of these test case results before concluding with directions for future research.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Conclusion

As mentioned in the introduction, the popularity of Low Fare Carriers has led to the removal of fare restrictions in many markets. The most competitive of these markets have unrestricted fare structures; the fare products all lack restrictions and thus differ only by price. This results in passengers buying the lowest-priced fare product available since there are no longer any restrictions that prevents them from doing so.

To forecast the demand for each of the fare products in such markets, airlines can use a forecasting technique known as “Q-forecasting” that considers the sell-up potential of passengers. Sell-up occurs when passengers upon being denied booking their original fare class choice, choose to pay more for the next available higher-priced fare class provided their maximum willingness to pay value has not been exceeded. Airlines can increase their revenues in unrestricted markets if they quantify the sell-up potential of passengers either using estimated or input values. Estimation of sell-up potential based on the maximum willingness to pay of passengers is thus critical to helping airlines raise their revenues when competing in unrestricted fare markets.

This thesis has consequently investigated and compared the revenue performance of the following 3 sell-up estimation methods available in the PODS simulation model: (i) Direct Observation (DO), (ii) Forecast Prediction (FP) and (iii) Inverse Cumulative (IC). Since sell-up is based on price-oriented behavior, the simulations were conducted in a network with an unrestricted fare structure consisting of 2 airlines, AL1 and AL2, under a variety of fare class optimization scenarios. Both estimated sell-up values and input sell-up values were tested on AL1 whereas only input sell-up values were tested on AL2. Additionally, to

Sell-Up Method	Percentage Changes in Revenues over Base Case	
	DM = 0.8	DM = 1.0
Input FRAT5C	2.3%	2.4%
Estimated DO	0.0%	2.5%
Estimated FP	2.0%	3.0%
Estimated IC	2.6%	2.6%

Table 5.1: Revenues Changes for AL1 in switching from AT90 to EMSRb (Q-forecasting)

Sell-Up Method	Percentage Changes in Revenues over Base Case	
	DM = 0.8	DM = 1.0
Input FRAT5C	2.5%	0.2%
Estimated DO	0.0%	-0.9%
Estimated FP	1.6%	0.1%
Estimated IC	1.8%	-0.6%

Table 5.2: Revenues Changes for AL1 in switching from AT90 to DAVN (Q-forecasting)

verify the revenue gains achieved by airlines from practicing sell-up using “Q-forecasting”, a simulation base case was conducted whereby both airlines used a threshold RM method in the absence of sell-up. The results from the other test cases in the simulation could then be compared with those of the base case.

5.1 Summary of Findings

The revenue increases arising from using “Q-forecasting” can be readily observed from comparing the percentage changes in revenues experienced by AL1 as it switched from the threshold RM method to one that incorporates sell-up. Table 5.1 contains the percentage changes in revenues for AL1 as it switched from a threshold method (AT90) to a leg-based heuristic that uses “Q-forecasting” (EMSRb); this corresponds to comparing Test Case 0 with Test Case 1. As indicated by the percentage values, AL1 increased its revenues when it incorporated sell-up in optimizing its fare class distribution.

In switching from a threshold method (AT90) to a network O-D optimizer that uses “Q-forecasting” (DAVN), AL1 also benefited from sell-up if it used either Input FRAT5C or FP. Table 5.2 contains the percentage changes in revenues for AL1 for this comparison, which corresponds to comparing Test Case 0 with Test Case 3. Due to the specific network layout of Network D6 used for the simulation, the percentage changes in revenues were

Test Case	Sell-Up estimation method with the highest revenues	
	DM = 0.8	DM = 1.0
1	IC	FP
2	FP	FP
3	IC	FP
4	FP	FP
5	FP	FP

Table 5.3: Summary of the best-performing sell-up estimation method

greater when it switched to EMSRb than when it switched to DAVN. Nonetheless, using DAVN (with either input FRAT5C values or FP) instead of AT90 always resulted in revenue gains for AL1.

In summary, AL1 typically achieved the greatest revenues out of these 3 sell-up estimation methods when it used FP as indicated by Table 5.3. Further details on the revenue values of the 3 sell-up estimation methods as well as their percentage differences from using input FRAT5C values for AL1 can be obtained from the various figures in Chapter 4 that depict the simulation results.

The robustness of FP as the best-performing sell-up estimation method in most of these test cases arises primarily because the model sell-up probabilities for the current time frame are calculated using the associated sell-up probabilities from the immediately preceding time frame as indicated by Equation 3.3 in Section 3.2.3. By using such information from the previous time frame, the resultant calculated sell-up probability values for the current time frame has been shown in the simulations to be less prone to under- or over-estimated relative to the actual WTP values of the passengers in that particular time frame. When the estimated sell-up probability values best represent the passengers' maximum willingness to pay for each of the time frames, the airline thus stands to gain revenues.

Based on this outcome, the simulations have shown that sell-up estimators such as FP produce more revenues than threshold methods. More importantly, they can even perform better than input FRAT5 sell-up values under some of the seat allocation optimizers. The implication for airlines is that they should considering switching from threshold methods to sell-up estimators since the latter are robust enough to operate under seat allocation optimizers when used in unrestricted fare environments to raise revenues.

The simulations have also shown that the revenue performance of the sell-up estimators

are also affected by the competitive environment. As noted by the test cases, when the competitor (AL2) becomes increasingly sophisticated in shifting from threshold to leg-based heuristics and finally network-OD optimization methods, the degree of sell-up estimated by AL1 decreases. Intuitively, this makes sense since as AL2 uses more sophisticated seat allocation optimizers, it is better positioned to capture passengers from AL1.

5.2 Future Research Directions

Based on the findings of this thesis, there are two potential research directions for further investigation.

The first suggestion is a direct follow-up to the simulation results of this thesis - that the revenue performance of the 3 sell-up estimation methods be tested in simulated airline networks having more than 2 airlines. The main motivation of this would be to obtain a more detailed understanding of the effects that each of the sell-up estimation methods has on the revenues of the practicing airline's competitors. Using multi-airline networks would especially be beneficial in evaluating the sell-up estimation methods when used in conjunction with network-OD optimization methods since such networks will probably not exhibit the misleadingly large competitor revenue gains as depicted in Test Cases 4 and 5 from Chapter 4.

The second suggestion requires significantly more effort since it involves developing and testing an additional definition of sell-up that is applicable to price-oriented passengers. As stated in Section 3.1, the current definition of sell-up is when price-oriented passengers, upon being denied their intended bookings in Q-class, pay more for the next available higher-priced fare class. However, it can be argued that an alternative form of sell-up can also be defined for product-oriented passengers who, as described in Section 2.3.2, choose a fare class based on the associated restrictions. Since higher-priced fare classes generally have fewer restrictions than lower-priced fare classes, product-oriented passengers may decide to pay more for the next applicable higher-priced fare class if their original fare class choice is no longer available. Some preliminary investigation into such price-oriented sell-up (also known as pairwise sell-up) has been done by the PODS Consortium as indicated by Soo [32]. However, further development of the definition followed by testing in both semi-restricted and mixed-fare structures need to be performed before any conclusive remarks can be made.

Bibliography

- [1] C. Barnhart, P. P. Belobaba, and A. R. Odoni. Applications of operations research in the air transport industry. *Transportation Science*, 37(4):368–391, November 2003.
- [2] P. P. Belobaba. *Air Travel Demand and Airline Seat Inventory Management*. PhD thesis, Massachusetts Institute of Technology, May 1987.
- [3] P. P. Belobaba. Applications of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37:183–197, 1989.
- [4] P. P. Belobaba. Optimal versus heuristic methods for nested seat allocation. *AGIFORS Reservations Control Study Group Meeting*, 1992.
- [5] P. P. Belobaba. The evolution of airline yield management: Fare class to origin-destination seat inventory control. pages 285–302, 1998.
- [6] P. P. Belobaba. Airline differential pricing for effective yield management. pages 349–361, 1999.
- [7] P. P. Belobaba. Q investigations - algorithms for unrestricted fare classes. In *PODS Presentation, Amsterdam*, 2004.
- [8] P. P. Belobaba. Methods for estimating sell-up. *PODS Consortium Summit : Sep 25 - 26, Cambridge, MA*, 2006.
- [9] P. P. Belobaba and C. Hopperstad. Algorithms for revenue management in unrestricted fare markets. In *Presented at the Meeting of the INFORMS Section on Revenue Management, Cambridge, MA*, 2004.
- [10] P. P. Belobaba and A. R. Weatherford. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences*, 27(2):343–363, 1996.
- [11] C. H. Bohutinsky. The sell-up potential of airline demand. Master’s thesis, Massachusetts Institute of Technology, May 1990.
- [12] E. A. Boyd and R. Kallesen. The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*, 3(2):171–177, 2004.
- [13] S. J-C. Bratu. Network value concept in airline revenue management. Master’s thesis, Massachusetts Institute of Technology, May 1998.

- [14] S. L. Brumelle and J. I. McGill. Airline seat allocation with multiple nested fare classes. In *Paper presented at the Fall ORSA/ TIMS Conference, Denver, CO*, 1987.
- [15] R. L. Cleaz-Savoyen. Airline revenue management for less restricted fare structures. Master's thesis, Massachusetts Institute of Technology, May 2005.
- [16] W. L. Cooper, T. Homem-de Mello, and A. J. Kleywegt. Models of the spiral-down effect in revenue management. *Operations Research*, 54(5):968–990, 2006.
- [17] R. E. Curry. Optimal airline seat allocation with fare classes nested by origin and destinations. *Transportation Science*, 24(3):193–204.
- [18] H. Dunleavy and D. Westermann. Future of airline revenue management. *Journal of Revenue and Pricing Management*, 3(4):280–282, 2005.
- [19] T. O. Gorin. Airline revenue management: Sell-up and forecasting algorithms. Master's thesis, Massachusetts Institute of Technology, May 2000.
- [20] T. O. Gorin. *Assessing Low-Fare Entry in Airline Markets: Impacts of Revenue Management and Network Flows*. PhD thesis, Massachusetts Institute of Technology, May 2004.
- [21] T. O. Gorin and P. P. Belobaba. Revenue management performance in a low-fare airline environment: Insights from the passenger origin-destination simulator. *Journal of Revenue and Pricing Management*, 3(3):215–236, 2004.
- [22] J. C. Guo. Review: Methods for estimating sell-up. *PODS Consortium Summit : Jan 31 - Feb 1, Los Angeles, CA*, 2008.
- [23] C. Hopperstad. Technical specifications. *Passenger O/D Simulator PODS v.10*, 2005.
- [24] R. Kuhlmann. Why is revenue management not working? *Journal of Revenue and Pricing Management*, 2(4):378–387, 2004.
- [25] A. Y. Lee. Investigation of competitive impacts of origin-destination control using pods. Master's thesis, Massachusetts Institute of Technology, May 1998.
- [26] K. Littlewood. Forecasting and control of passenger bookings. *12th AGIFORS Annual Symposium Proceedings, Nathanya, Israel*, pages 95–117, 1972.
- [27] J. I. McGill and G. J. Van Ryzin. Revenue management: Research overview and prospects. *Transportation Science*, 33(2):233–256, May 1999.
- [28] R. Ratliff and B. Vinod. Airline pricing and revenue management: A future outlook. *Journal of Revenue and Pricing Management*, 3(4):302–307, 2005.
- [29] E. Reyes. Modeling airline passenger choice: Passenger preference for schedule in the passenger origin-destination simulator (pods). Master's thesis, Massachusetts Institute of Technology, May 2003.
- [30] M. H. Reyes. Hybrid forecasting for airline revenue management in semi-restricted fare structures. Master's thesis, Massachusetts Institute of Technology, May 2006.

- [31] B. C. Smith and C. W. Penn. Analysis of alternate origin-destination control strategies. In *AGIFORS Annual Symposium Proceedings*, volume 28, pages 123–144, 1988.
- [32] S. Soo. Leg-based pairwise sell-up estimation. *PODS Consortium Summit : May 30 - 31, Minneapolis, MN*, 2007.
- [33] H. R. Thompson. Statistical problems in airline reservations control. *Operations Research Quarterly*, 12:167–185, 1961.
- [34] R. R. Wickham. Evaluation of forecasting techniques for short-term demand of air transportation. Master's thesis, Massachusetts Institute of Technology, May 1995.
- [35] E. L. Williamson. *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*. PhD thesis, Massachusetts Institute of Technology, May 1992.
- [36] R. Windle and M. Dresner. Competitive responses to low cost carrier entry. *Transportation Research Part E*, 35:59–75, 1999.
- [37] R. D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40(1):26–37.
- [38] J. S. Zickus. Forecasting for airline network revenue management: Revenue and competitive impacts. Master's thesis, Massachusetts Institute of Technology, May 1998.